EXHIBIT C

Nantworks, LLC v. Bank of America Corporation
Central District of California, No. 2:20-cv-7872
Invalidity Chart for U.S. Patent 9,324,004 based on
U.S. Patent U.S. Patent Publication No. US2002/004786 to Mault et al. ("Mault")

Exhibit A-4

Bank of America Corporation ("BofA" or "Defendant") provides this chart subject to all reservations, objections, statements, and disclaimers set forth herein and in BofA's Final Invalidity Contentions Cover Pleading, as well as any amendment, supplement, or modification thereof, which are incorporated herein by reference in their entirety.

On information and belief, Mault was filed on September 7, 2001, claiming the priority date of September 7, 2000. Mault published at least as of April 25, 2002, and is available as prior art at least under 35 U.S.C. § 102(e).

As shown in the chart below, Mault anticipates asserted claims 1, 6, and 18 of the '004 Patent. To the extent Mault is found not to expressly disclose certain limitations in the asserted claims, such limitations are inherent. To the extent Mault is found not to anticipate any asserted claims or claim elements of the '004 Patent, the reference nevertheless renders those claims or claim elements obvious under 35 U.S.C. § 103, either alone or in combination with other art identified herein or in the cover pleading of BofA's Final Invalidity Contentions, for the reasons identified in the cover pleading. These Final Invalidity Contentions are not an admission by BofA that the accused products, including any current or past versions of the accused products, are covered by, or infringe the asserted claims, but are based instead on the recognition that if the claims are interpreted to be broad enough to encompass the accused products, the claims must also be construed to have that same scope when considering whether they are invalid.

The following chart is based on the Court's claim construction rulings (Dkt. Nos. 145, 153, 236), and to the extent not contradictory, Plaintiffs' apparent interpretations of the claims in Plaintiff's "Preliminary" Final Infringement Contentions. BofA notes that in many instances, Plaintiff's "Preliminary" Final Infringement Contentions ignore and/or contradict the Court's claim construction rulings. BofA does not accept the assumptions concerning the scope and meaning implicit in Plaintiff's "Preliminary" Final Infringement Contentions, to the extent those assumptions are discernible, and reserves the right to challenge Plaintiff's proposed (or implied) constructions. Furthermore, Plaintiffs' "Preliminary" Final Contentions fail to comply with the level of specificity required by SPR 2.1.2 and 2.1.3, and include improper reservation of rights to supplement with "additional ways in which BoA Accused Products infringe." Any alleged deficient discovery is of Plaintiffs own making. Most if not all of the deficiencies in Plaintiffs' "Preliminary" Final Contentions have been documented by Defendants through multiple letters since at least May 18, 2021 and up to Bank of America's most recent letter of August 19, 2023. Plaintiffs' shifting infringement contentions, deficiencies in its infringement contentions, and changing claim scope have prejudiced BofA from being able to reasonably prepare its defenses, including its invalidity contentions. BofA also reserves the right to revise and supplement these charts if and when Plaintiff is permitted to provide revised or supplemental Infringement Contentions. Citations given in the chart below are merely representative of the respective elements to provide notice to Plaintiffs' of BofA's invalidity theories and are not meant to be exhaustive.

1. Claim 1

Claim Element	'004 Claim Recitation	Exemplary Citations in Reference
[1Pre]	A method for processing a video stream, comprising:	Mault discloses a method for processing a video stream. Mault at [0057] • "Embodiments of the invention can also be implemented using recorded video images. The person can view recorded videos at a convenient later time, so as to assist the person create a diet log." Mault at [0089] • "Referring to FIG. 16: IR lens **402** is used as part of an IR data communications system enabling watch **400** to communicate with other devices; camera lens **404** forms part of an image sensor; shutter button **406** is pressed to capture an image or video using the image sensor; mode button **408** changes the operating mode of the device (modes may include some or all of the following: image recording, image display, video recording, video display, television, wireless, timepiece, calculator, personal organizer, wireless phone, video phone, Internet access device, diet log, activity sensor, physiological sensor (e.g. blood glucose), and any other useful function); and set button **410**, reverse button **412**, change button **414**, forward button **416** are used in e.g. data entry and image review processes."
[1A]	analyzing, via a mobile device, a scene represented by the video stream for at least on object;	Mault discloses analyzing, via a mobile device, a scene represented by the video stream for at least on object Mault at [0006] • "A method of creating a record of foods consumed, or diet log, comprises recording images corresponding to the foods consumed. An image

Claim Element	**'004 Claim Recitation**	**Exemplary Citations in Reference**
		corresponding to a food item consumed can comprise an image of the food, food package, label, menu listing, barcode (such as a Universal Product Code or UPC), food serving on a plate, food serving in a restaurant package, package containing multiple food servings, alphanumeric codes, receipts, printed data (such as nutritional information conventionally found on food packages), notes written or typed about the food, and the like." Mault at [0009] • "After recording the image, at a convenient later time the person can retrieve a stored image and view the image on a display. The display can be part of the imaging device, or other device such as a personal digital assistant, personal computer, Internet appliance, WebTV, e-book, tablet computer, pager, cell phone, interactive TV, spectacle mounted display, visor, and the like. An electronic device having a display used to view an image may be in communication with a separate device used to create the image, for example using wireless or cable-based links." Mault at [0011] • "The imaging device, used to create an image of foods consumed, can be a digital camera, device having the functionality of a camera, such as a personal digital assistant (PDA), cell phone, electronic book, pager, calculator, or other consumer electronics device. The imaging device can also be an ornamental component, clothing item, or other functional or decorative item having image creation capabilities, such as spectacles, buttons, a pen, cutlery item, key ring, keys, belt, cigarette-like object, ring, body ornament, or jewelry." Mault at [0016] • "The food image can be recorded in a memory of the electronic device, or any other memory accessible to the electronic device, such as a memory module, remote database, and the like. The food item name can be selected

Claim Element	**'004 Claim Recitation**	**Exemplary Citations in Reference**
		from one or more lists of food item names presented on the display of the electronic device, for example from one or more menus presented on the display." Mault at [0017] • "A method of creating a diet log comprises: recording an image of a label, the label corresponding to a food item consumed; analyzing the image of the label (for example using optical character recognition, image analysis, image recognition, or neural network techniques) so as to determine the identity of the food item consumed; and storing the determined identity of the food item consumed in an electronic memory, so as to create a diet log. The label can be a printed label, a barcode, an alphanumeric code, a nutritional data label, a list of ingredients, a product code, graphical designs (Such as trademarks), product names, and the like." Mault at [0018]. • "A method by which a person can create a record of foods consumed comprises: recording a food image of a food consumed; transmitting the food image to a remote computer system, for example over a communications network such as the internet, receiving food identification data from the remote computer system (which may be generated by a software running on the remote computer system, a computer expert system, an image analysis system, a human operative in communication with the remote computer system, or some combination of techniques); and recording the food identification data So as to create a record of foods consumed." Mault at [0048] • "FIG. 5 is a schematic of a system comprising an electronic device 100, an imaging system 102, a physiological sensor 104, an activity sensor 106, and a communications link through communications network 108 to a remote

Claim Element	'004 Claim Recitation	Exemplary Citations in Reference
		computer system (such as a server system) 110." Mault at Fig. 5. 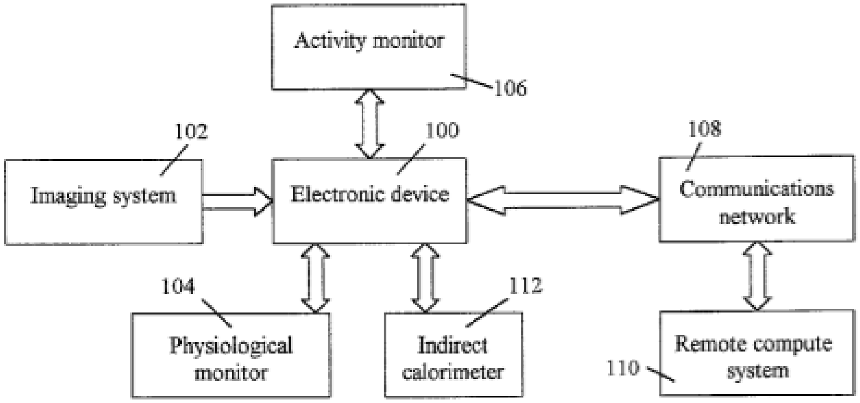 Figure 5 Mault at [0049] • "Images recorded with the imaging device are transmitted over the communications network to the remote computer. The remote computer system can be a server system, and may comprise a computer expert system for image analysis. Image analysis can be performed by one or more persons, a computer expert system, or other mechanism." Mault at [0051] • "FIG. 6 is a flow chart of a method for diet logging. Box 120 corresponds to recording an image of food. Box 122 corresponds to transmission of the image to a remote location, Such as a remote computer System. . . . Box 124

Claim Element	'004 Claim Recitation	Exemplary Citations in Reference
		corresponds to analysis of the image of the remote location." Mault at Fig. 6. 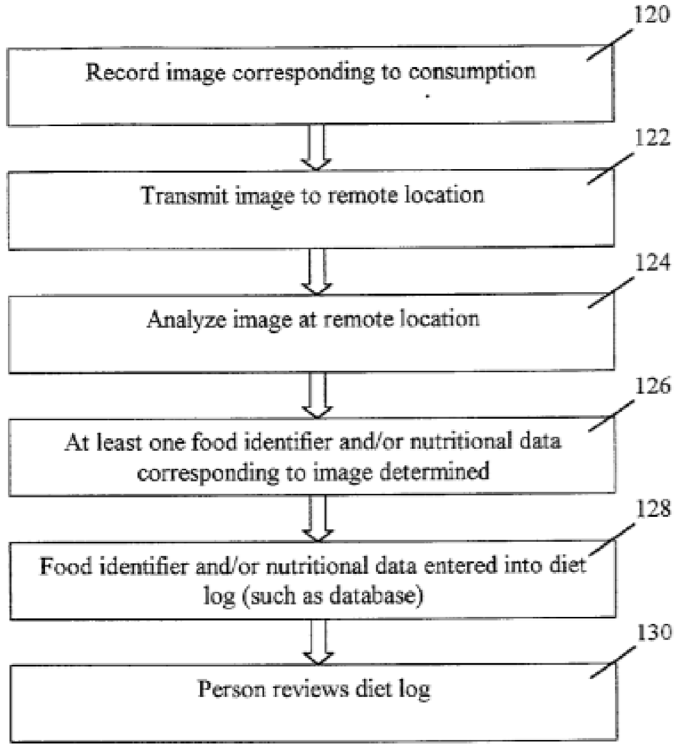 Figure 6 Mault at [0046] • "Other options which may be presented to the user include the entering of a food by name, for example into diet log software, or scanning the barcode of

Claim Element	'004 Claim Recitation	Exemplary Citations in Reference
		a food item. The electronic device may comprise or communicate with a barcode scanner, or the optical image Sensor can be used to record an image of the barcode. The barcode can then be interpreted in terms of a UPC (universal product code) and/or product name using a database accessible to the PDA. A UPC database may be accessible over a communications network Such as the Internet." Mault at [0047] • "The step corresponding to box **88** can be omitted, and corresponding nutritional data can be identified after the food identifier is stored in memory. The food item identifier can be a food name, product code, barcode, and the like." Mault at [0074] • "Fig. 11 shows a PDA **280** with display **284**, buttons **282**, and imaging device 286 directed at a can of pickles **288**. The display **284** can be used to indicate the image to be captured. A button such as 282 may be pressed to record the image. The display **184** can be used to view captured images of foods stored in the memory of the PDA, for assisting the person in creating a diet log." Mault at [0080] • "Barcodes can also be used in food item identification. The PDA may also incorporate a bar-code Scanner, or alternatively, barcodes may be identified from a recorded image of the bar-code using image analysis. The barcode (e.g. universal product code) can be used to retrieve nutritional data relating to the item from a database. The database may reside on a remote computer System (remote in this context meaning a computer System not carried by the person, for example a commercial Server System, home computer, etc.). The person may scan the barcode of a box of cereal and the UPC code used to retrieve nutrition information specific to that brand of cereal from the

Claim Element	**'004 Claim Recitation**	**Exemplary Citations in Reference**
		database." Mault at [0081] - "FIG. 12 shows an image of a barcode presented on the display 304 of PDA 300 (the PDA having data entry buttons 302, and an image sensor 306). Image analysis software can be used to analyze the image of the barcode, So as to determine UPC, product name, nutritional data, and the like. Alternatively, a barcode Scanner can be used to read the barcode, or an accessory device which in communication with the PDA may be used as a barcode scanner." Mault at Fig. 12

Claim Element	**'004 Claim Recitation**	**Exemplary Citations in Reference**
		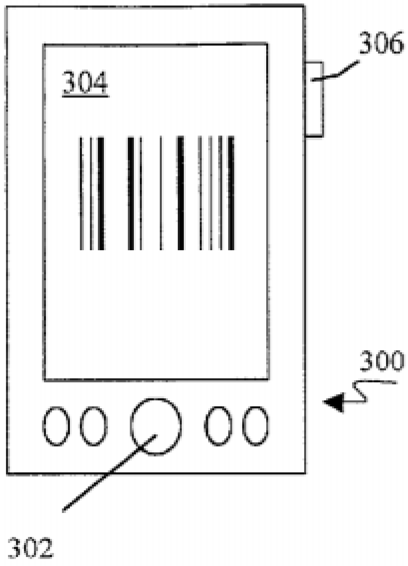 Figure 12 Mault at [0082] - "FIG. 13 shows an image of nutritional information on the display **324** of a PDA **320**, having data entry buttons **322** and optical image Sensor **326**. Optical character recognition Software running on a processor of the PDA can be used to extract nutritional information from the image, and place it into memory for possible later transfer to a database." Mault at Fig. 13

Claim Element	'004 Claim Recitation	Exemplary Citations in Reference
		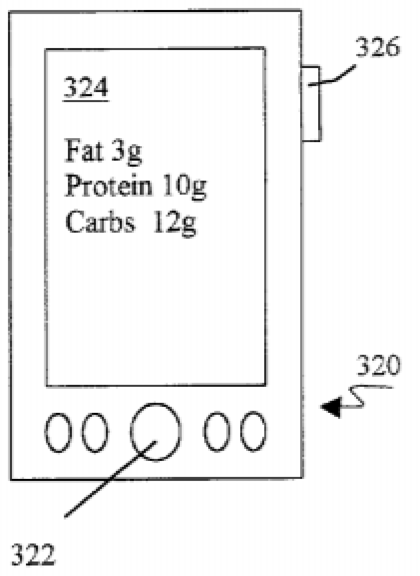 Figure 13 To the extent it is found that Mault does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Mault does not anticipate this claim, Mault renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Mault with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

Claim Element	**'004 Claim Recitation**	**Exemplary Citations in Reference**
[1B]	deriving at least one characteristic of the video stream;	Mault discloses deriving at least on characteristic of the video stream. Mault [0021] • "An imaging sensor having a two-dimensional array of sensor elements can capture an image by having the array exposed to the image, so as to create an electronic data set or file corresponding to the spatial distribution of intensity (and color, if required) of the image." Mault at Figure 1A: 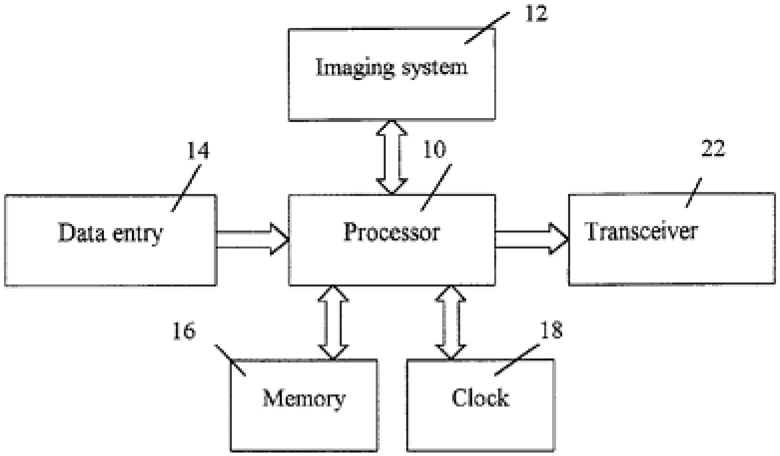 Mault at [0041] • "FIG. 1A, other components of the electronic device are omitted for clarity. On receiving a signal from the external control or from the processor, the controller **24** initiates the capture of an image by sensor **28**, and may store the

Claim Element	'004 Claim Recitation	Exemplary Citations in Reference
		image in memory **30** along with a time stamp provided by clock **32**." Mault at [0047] - "The time and location that the image was recorded can also be recorded in memory and displayed, so as to assist diet logging." Mault at [0076–77] - "For non-packaged foods, e.g. restaurant meals, image analysis may also be used in identifying the food items. Imaging at a number of wavelengths, followed by false color image generation, can be used to help identify image components. Spectroscopic imaging can be used in computer-assisted food recognition. Image processing, image recognition, and pattern recognition algorithms are useful in analyzing stored images so as to create a diet log. Algorithms may be applied to color images, or images recorded at a number of different wavelengths (e.g. in the IR, optical, and UV) which may assist identification. Stored images can be identified by comparison with previously stored images that the person has identified, using a software learning mechanism such as a neural network." Mault at [0090] - "A person can record images of medicines taken, such as tablets consumed, and these images can be used to create a medical log using embodiments of the present invention. Colors, shapes, printed codes, and fluorescent markers of tablets can be used to assist identification, either by a person or an software image analysis program." Mault at [0091]

Claim Element	'004 Claim Recitation	Exemplary Citations in Reference
		• "Images can be recorded at intervals by a device carried by the person, and the images used to create a lifestyle log for the person. This can include analysis of images of meals consumed, so as to create a diet log. Images can be recorded automatically, for example at time intervals, so as to spare the person the effort of remembering to record images of foods consumed. Image creation can be triggered by events, such as physiological parameters which change during eating, position (such as using a global positioning system), voice commands, wireless transmission from a remote device, such as a food vending machine, and the like." To the extent it is found that Mault does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Mault does not anticipate this claim, Mault renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Mault with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1C]	recognizing the at least one object in the scene as a target object based at least in part on the at least one characteristics of the video stream;	Mault discloses recognizing the at least one object in the scene as a target object based at least in part on the at least one characteristics of the video stream. Mault at [0017] • "A method of creating a diet log comprises: recording an image of a label, the label corresponding to a food item consumed; analyzing the image of the label (for example using optical character recognition, image analysis, image recognition, or neural network techniques) so as to determine the identity of

Claim Element	**’004 Claim Recitation**	**Exemplary Citations in Reference**
		the food item consumed; and storing the determined identity of the food item consumed in an electronic memory, so as to create a diet log. The label can be a printed label, a barcode, an alphanumeric code, a nutritional data label, a list of ingredients, a product code, graphical designs (Such as trademarks), product names, and the like." Mault at [0018] • "A method by which a person can create a record of foods consumed comprises: recording a food image of a food consumed; transmitting the food image to a remote computer system, for example over a communications network such as the internet, receiving food identification data from the remote computer System (which may be generated by a software running on the remote computer System, a computer expert system, an image analysis system, a human operative in communication with the remote computer system, or some combination of techniques); and recording the food identification data So as to create a record of foods consumed." Mault at [0046] • "Other options which may be presented to the user include the entering of a food by name, for example into diet log software, or scanning the barcode of a food item. The electronic device may comprise or communicate with a barcode scanner, or the optical image sensor can be used to record an image of the barcode. The barcode can then be interpreted in terms of a UPC (universal product code) and/or product name using a database accessible to the PDA. A UPC database may be accessible over a communications network such as the Internet. Other options include receiving transmissions such as wireless data transmissions from food vendors or food vending devices, viewing the recorded diet log, editing the diet log, or analyzing previously recorded images."

Claim Element	**'004 Claim Recitation**	**Exemplary Citations in Reference**
		Mault at [0051] • "Box 124 corresponds to analysis of the image of the remote location. Image recognition software may be used. A trained person, at a remote location, can select food names corresponding to the images, for example in conjunction with image recognition, providing names where image recognition has not been successful in identifying the food. A computer expert System, optical character recognition of a name on the package, barcode recognition, portion size estimate, and other food data can also be extracted from the image or images." Mault at [0075–76] • "Image processing techniques can be used to identify food items from images. For packaged foods, optical character recognition may be used to record and identify nutritional information and/or identify the item consumed. Computer analysis of images may be carried out on the PDA or using another computer system in communication with the PDA. For example, an image of food packaging may be used to identify the food contained. An image of the nutritional content information panel provided by the manufacturer on the package may also be recorded. Optical character recognition may be used to obtain information from the image for storage in a diet log. An image of a box of corn flakes may be recognized by a computer as such, and used to generate a corn flake serving record in a diet log. For non-packaged foods, e.g. restaurant meals, image analysis may also be used in identifying the food items. Imaging at a number of wavelengths, followed by false color image generation, can be used to help identify image components. Spectroscopic imaging can be used in computer-assisted food recognition." Mault at [0077]

Claim Element	'004 Claim Recitation	Exemplary Citations in Reference
		• "Image processing, image recognition, and pattern recognition algorithms are useful in analyzing stored images so as to create a diet log. Algorithms may be applied to color images, or images recorded at a number of different wavelengths (e.g. in the IR, optical, and UV) which may assist identification. Stored images can be identified by comparison with previously stored images that the person has identified, using a software learning mechanism such as a neural network." Mault at [0080] • "Barcodes can also be used in food item identification. The PDA may also incorporate a bar-code Scanner, or alternatively, barcodes may be identified from a recorded image of the bar-code using image analysis. The barcode (e.g. universal product code) can be used to retrieve nutritional data relating to the item from a database. The database may reside on a remote computer System (remote in this context meaning a computer System not carried by the person, for example a commercial Server System, home computer, etc.). The person may scan the barcode of a box of cereal and the UPC code used to retrieve nutrition information specific to that brand of cereal from the database." Mault at [0081] • "FIG. 12 shows an image of a barcode presented on the display **304** of PDA **300** (the PDA having data entry buttons **302**, and an image sensor **306**). Image analysis software can be used to analyze the image of the barcode, so as to determine UPC, product name, nutritional data, and the like." Mault at [0082] • "FIG. 13 shows an image of nutritional information on the display **324** of a PDA **320**, having data entry buttons **322** and optical image Sensor **326**.

Claim Element	**'004 Claim Recitation**	**Exemplary Citations in Reference**
		Optical character recognition Software running on a processor of the PDA can be used to extract nutritional information from the image, and place it into memory for possible later transfer to a database." Mault at [0090] • "Colors, shapes, printed codes, and fluorescent markers of tablets can be used to assist identification, either by a person or an [sic] software image analysis program." To the extent it is found that Mault does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Mault does not anticipate this claim, Mault renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Mault with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1D]	providing an indication upon recognizing the at least one object as the target object; and	Mault discloses providing an indication upon recognizing the at least one object as the target object. Mault at [0017] • "A method of creating a diet log comprises: recording an image of a label, the label corresponding to a food item consumed; analyzing the image of the label (for example using optical character recognition, image analysis, image recognition, or neural network techniques) so as to determine the identity of the food item consumed; and storing the determined identity of the food item

Claim Element	'004 Claim Recitation	Exemplary Citations in Reference
		consumed in an electronic memory, so as to create a diet log. The label can be a printed label, a barcode, an alphanumeric code, a nutritional data label, a list of ingredients, a product code, graphical designs (Such as trademarks), product names, and the like." Mault at [0051] • "Box **124** corresponds to analysis of the image of the remote location. Image recognition software may be used. A trained person, at a remote location, can select food names corresponding to the images, for example in conjunction with image recognition, providing names where image recognition has not been successful in identifying the food. A computer expert system, optical character recognition of a name on the package, barcode recognition, portion size estimate, and other food data can also be extracted from the image or images. Box **126** corresponds to determining food identity and/or nutritional content of the food within the image." Mault at [0052] • "A diet log channel, generated by remote server (or other remote computer systems, or other content provider) **176**, and received by the set top box over communications network **178**, can be selected for display on the interactive television **160**. A display of food item identifiers, such as food names in a menu format, can be displayed on the display **162** of the interactive TV **160**. An image stored in the local memory **172** can be added to the diet log channel signal, using techniques known in the art, so that the person views the diet log channel with an embedded image of the food consumed. The person then selects items corresponding to the food consumed from the displayed menu. Selection data is generated with user input **146** (for example a keyboard) and transmitted to the set top box and to the remote server. The display is changed appropriately on receipt of the selection data by the content provider and the selection is added to the diet log of the person, which is stored in a database

Claim Element	**'004 Claim Recitation**	**Exemplary Citations in Reference**
		associated with or otherwise accessible by the remote server **176**." Mault at [0059] - "For example, a person can sit down to a meal and record an image of the meal using the electronic device. The image is transmitted to a remote location where it is analyzed and converted into a caloric value. The caloric value can be transmitted back to the person, allowing the person to decide whether to eat the meal, choose another meal, or eat some portion of the meal consistent with reaching a calorie balance goal." To the extent it is found that Mault does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Mault does not anticipate this claim, Mault renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Mault with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1E]	presenting to a user, via an address, content information associated with the target object.	Mault discloses presenting to a user, via an address, content information associated with the target object. Mault at Fig. 6

Claim Element	’004 Claim Recitation	Exemplary Citations in Reference
		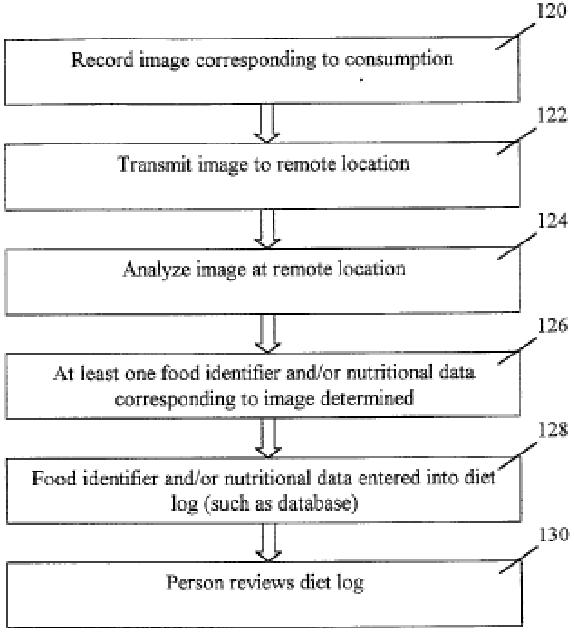 Figure 6 Mault at [0009] - “After recording the image, at a convenient later time the person can retrieve a stored image and view the image on a display. The display can be part of the imaging device, or other device such as a personal digital assistant, personal computer, Internet appliance, WebTV, e-book, tablet computer, pager, cell phone, interactive TV, spectacle mounted display, visor, and the like.” Mault at [0017] - “The identity of food items consumed can be correlated with nutritional data, for example using a nutritional database, and the nutritional data can stored

Claim Element	'004 Claim Recitation	Exemplary Citations in Reference
		in an electronic memory so as to create a nutritional record, for example as part of a diet log." Mault at [0052] • "A diet log channel, generated by remote server (or other remote computer systems, or other content provider) **176**, and received by the set top box over communications network **178**, can be selected for display on the interactive television **160**. A display of food item identifiers, such as food names in a menu format, can be displayed on the display **162** of the interactive TV **160**. An image stored in the local memory **172** can be added to the diet log channel signal, using techniques known in the art, so that the person views the diet log channel with an embedded image of the food consumed. The person then selects items corresponding to the food consumed from the displayed menu. Selection data is generated with user input **146** (for example a keyboard) and transmitted to the set top box and to the remote server. The display is changed appropriately on receipt of the selection data by the content provider and the selection is added to the diet log of the person, which is stored in a database associated with or otherwise accessible by the remote server **176**." Mault at [0059] • "For example, a person can sit down to a meal and record an image of the meal using the electronic device. The image is transmitted to a remote location where it is analyzed and converted into a caloric value. The caloric value can be transmitted back to the person, allowing the person to decide whether to eat the meal, choose another meal, or eat some portion of the meal consistent with reaching a calorie balance goal." Mault at [0074] • "FIG. 11 shows a PDA **280** with display **284**, buttons **282**, and imaging device **286** directed at a can of pickles **288**. The display **284** can be used to

Claim Element	**'004 Claim Recitation**	**Exemplary Citations in Reference**
		indicate the image to be captured. A button such as **282** may be pressed to record the image. The display **184** can be used to view captured images of foods stored in the memory of the PDA, for assisting the person in creating a diet log." Mault at [0082–83] • "FIG. 13 shows an image of nutritional information on the display **324** of a PDA **320**, having data entry buttons **322** and optical image sensor **326**. Optical character recognition software running on a processor of the PDA can be used to extract nutritional information from the image, and place it into memory for possible later transfer to a database. Food item identities and nutritional information can also be received from food vendors, e.g. grocery stores, on-line retailers, restaurants, vending machines, diet food retailers, etc. The PDA may prompt the person to identify the source of the food being imaged, and this information can be used to help in identification." To the extent it is found that Mault does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Mault does not anticipate this claim, Mault renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Mault with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

2. Claim 6

Claim	'004 Claim Recitation	Exemplary Citations in Reference
6	The method of claim 1, wherein the mobile device comprises a camera capable of at least one of image and video capture.	Mault discloses the mobile device comprises a camera capable of at least one of image and video capture. Mault at [0052] • "FIG. 7 shows an electronic device **140** comprising a processor **142**, an image sensor **144**, data entry mechanism **146**, memory **148**, clock **150**, and wireless transmission unit **152**." • "A diet log channel, generated by remote server (or other remote computer systems, or other content provider) **176**, and received by the set top box over communications network **178**, can be selected for display on the interactive television **160**. A display of food item identifiers, such as food names in a menu format, can be displayed on the display **162** of the interactive TV **160**. An image stored in the local memory **172** can be added to the diet log channel signal, using techniques known in the art, so that the person views the diet log channel with an embedded image of the food consumed. The person then selects items corresponding to the food consumed from the displayed menu. Selection data is generated with user input **146** (for example a keyboard) and transmitted to the set top box and to the remote server. The display is changed appropriately on receipt of the selection data by the content provider and the selection is added to the diet log of the person, which is stored in a database associated with or otherwise accessible by the remote server **176**."

		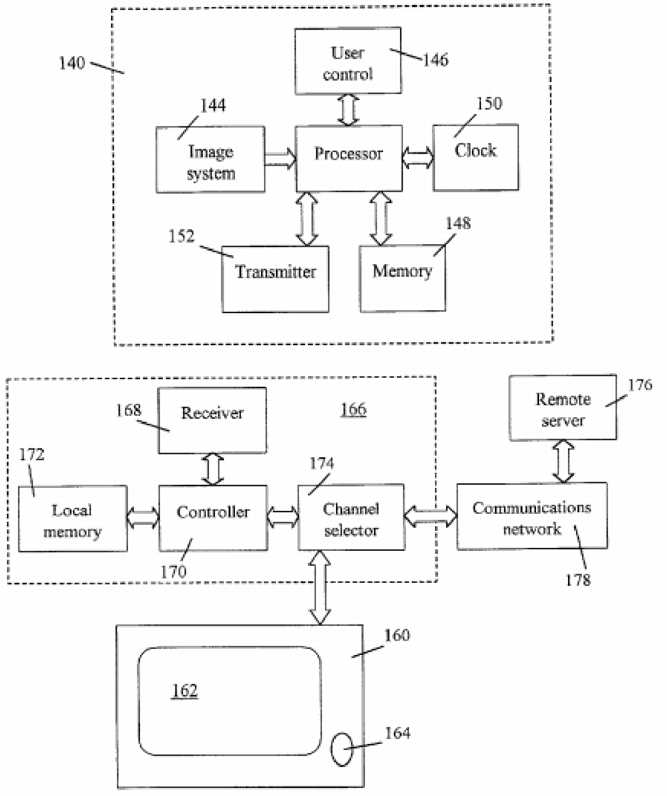 Figure 7 To the extent it is found that Mault does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Mault does not anticipate this claim, Mault renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Mault with any of the other prior art references identified in

		the cover pleading or herein with a reasonable expectation of success because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

3. Claim 18

Claim	'004 Claim Recitation	Exemplary Citations in Reference
18	The method of claim 1, further comprising communicating at least one of banking account information, user data and terminal specific data to a banking company.	To the extent it is found that Mault does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Mault does not anticipate this claim, Mault renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Mault with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

Exhibit A-5

Bank of America Corporation ("BofA" or "Defendant") provides this chart subject to all reservations, objections, statements, and disclaimers set forth herein and in BofA's Final Invalidity Contentions Cover Pleading, as well as any amendment, supplement, or modification thereof, which are incorporated herein by reference in their entirety.

On information and belief, the application that matured into Sizer was filed on March 28, 1997. Sizer published at least as of March 14, 2000, and is available as prior art at least under 35 U.S.C. § 102(a), (b) and (e).

As shown in the chart below, Sizer anticipates asserted claims 1, 6, and 18 of the '004 Patent. To the extent Sizer is found not to expressly disclose certain limitations in the asserted claims, such limitations are inherent. To the extent Sizer is found not to anticipate any asserted claims or claim elements of the '004 Patent, the reference nevertheless renders those claims or claim elements obvious under 35 U.S.C. § 103, either alone or in combination with other art identified herein or in the cover pleading of BofA's Final Invalidity Contentions, for the reasons identified in the cover pleading. These Final Invalidity Contentions are not an admission by BofA that the accused products, including any current or past versions of the accused products, are covered by, or infringe the asserted claims, but are based instead on the recognition that if the claims are interpreted to be broad enough to encompass the accused products, the claims must also be construed to have that same scope when considering whether they are invalid.

The following chart is based on the Court's claim construction rulings (Dkt. Nos. 145, 153, 236), and to the extent not contradictory, Plaintiffs' apparent interpretations of the claims in Plaintiffs' "Preliminary" Final Infringement Contentions. BofA notes that in many instances, Plaintiffs' "Preliminary" Final Infringement Contentions ignore and/or contradict the Court's claim construction rulings. BofA does not accept the assumptions concerning the scope and meaning implicit in Plaintiffs' Preliminary Final Infringement Contentions, to the extent those assumptions are discernible, and reserves the right to challenge Plaintiffs' proposed (or implied) constructions. Furthermore, Plaintiffs' "Preliminary" Final Contentions fail to comply with the level of specificity required by SPR 2.1.2 and 2.1.3, and include improper reservation of rights to supplement with "additional ways in which BoA Accused Products infringe." Any alleged deficient discovery is of Plaintiffs own making. Most if not all of the deficiencies in Plaintiffs' Preliminary Final Contentions have been documented by Defendants through multiple letters since at least May 18, 2021 and up to Bank of America's most recent letter of August 19, 2023. Plaintiffs' shifting infringement contentions, deficiencies in its infringement contentions, and changing claim scope have prejudiced BofA from being able to reasonably prepare its defenses, including its invalidity contentions. BofA also reserves the right to revise and supplement these charts if and when Plaintiffs are permitted to provide revised or supplemental Infringement Contentions. Citations given in the chart below are merely representative of the respective elements to provide notice to Plaintiffs of BofA's invalidity theories and are not meant to be exhaustive.

1. Claim 1

Element	'004 Claim Recitation	Exemplary Citations in Reference
[1Pre]	A method for processing a video stream, comprising:	Sizer discloses a method for processing a video stream. Sizer at FIG. 1 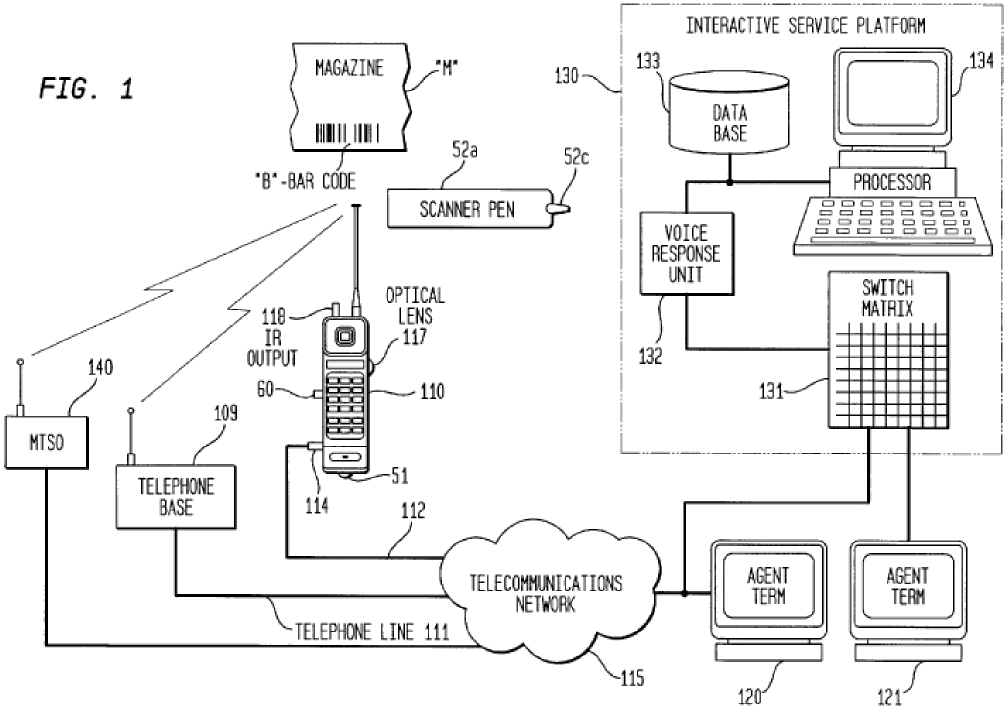 FIG. 1
[1A]	analyzing, via a mobile device, a scene represented by the video stream for at least one object;	Sizer discloses analyzing, via a mobile device, a scene represented by the video stream for at least one object.

Element	'004 Claim Recitation	Exemplary Citations in Reference
		Sizer at FIG. 1 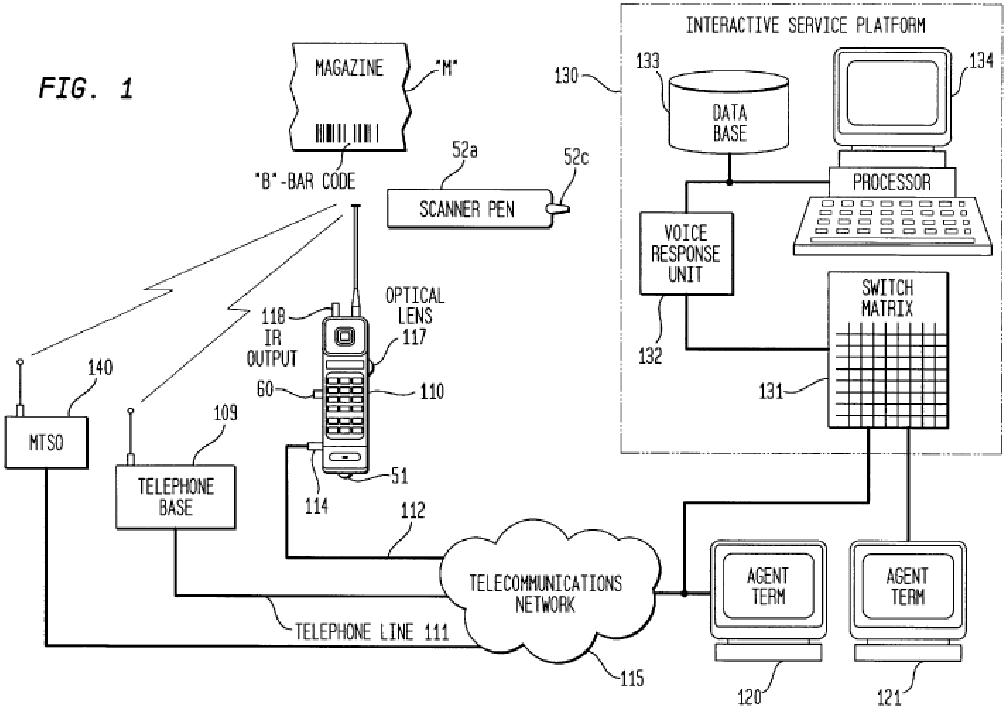 FIG. 1 Sizer at FIG. 2

Element	**'004 Claim Recitation**	**Exemplary Citations in Reference**
		FIG. 2 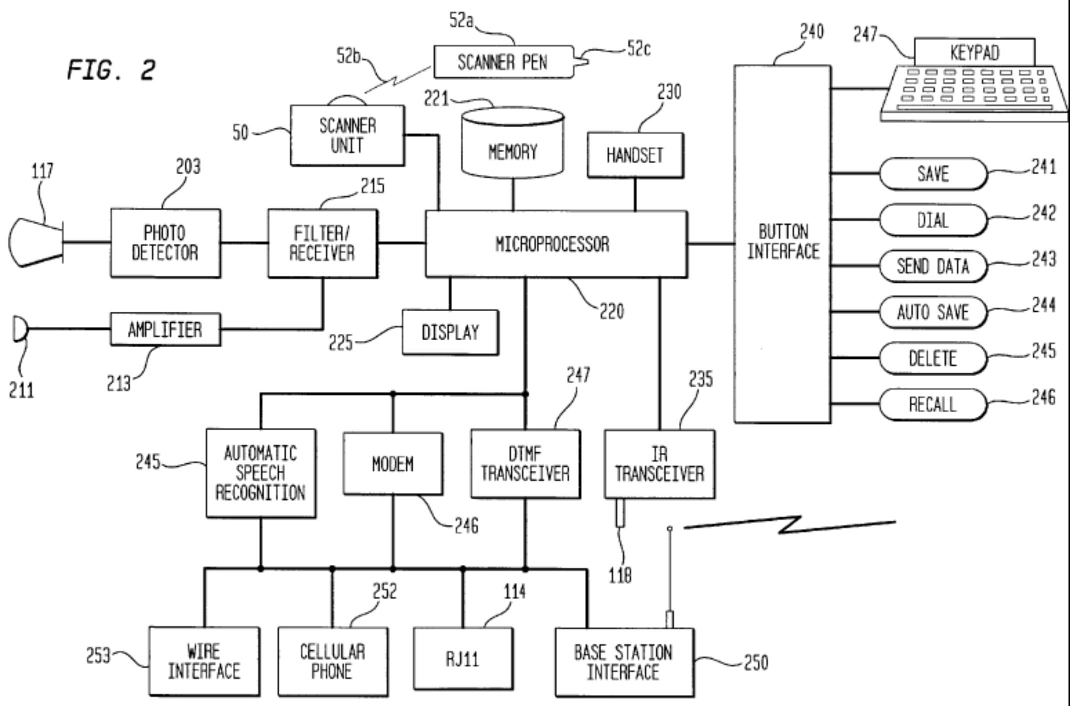 Sizer at 5:3-20 - "Briefly, it was found using a spectrum analyzer on a typical video signal, that there are comparatively large frequency components at the line rate and at the frame rate and its harmonics, but that between these frequencies, there are other frequency bands in which little information is carried. One such frequency band is between 15 and 30 kHz. By adding a low level carrier signal or tone at a frequency in this band, say 25 kHz, the video image is not degraded, but a properly tuned decoder can receive and decode the encoded information. In this way, digital information can be subliminally inserted in a video signal by adding to the video signal an amplitude shift

Element	'004 Claim Recitation	Exemplary Citations in Reference
		keyed (ASK) or frequency shift keyed (FSK) carrier signal, and the digital information can later be recovered using ASK or FSK decoding. Encoded data can also be inserted in a television signal and recovered by a received that responds to the picture displayed on the television, in the manner described in U.S. Pat. No. 4,807,031 cited above." Sizer at 5:23-26 • "If any non-discernible encoded data is part of a video signal, such sensing takes the form of light from display area being collected by an optical lens 117." Sizer at 6:42-47 • "If the encoded non-perceptible data is inserted in the video portion of a television signal, the images displayed on a visual display area a monitor or television are captured by optical lens 117 and an associated photo-detector 203, which is arranged to supply an electrical signal to a filter/receiver 215 representing the image." Sizer at 9:33-53 • "For background purposes, a televised advertisement for a product could include additional ordering or "coupon" information embedded in a subliminal visual pattern within the video image displayed on a television. The coupon information describes the item(s) offered for sale, the regular price, the coupon value, and contains additional information needed for ordering. This embedded information could be transmitted between a television and capture device at a rate of about 60 bits per second, that is sufficient such that the necessary information can be transmitted in a very brief period of time. In the application just described, the embedded product information, dialing number, coupon information and necessary checkbits together comprise approximately 2,000 bits of information, so that between three and four seconds would be required for transmission. Once the coupon

Element	'004 Claim Recitation	Exemplary Citations in Reference
		information is stored in capture device 110, the user of the capture device can supplement the coupon information with user entered data which can be entered using keypad 247. This supplemental information may include credit card or other billing information pertaining to the user, as an example." Sizer at 10:27-37 • In the active mode, a user operates capture device 110 in real time, typically while scanning a product or bar code, or in more complicated devices viewing a video display or listening to an audio presentation. Scanning a product or bar code, or in more complicated devices the user activates the SAVE button 241. When button 241 is activated, a positive result occurs in step 503, thereby capturing the encoded data, displaying all or a portion of the data on display 225, and storing same in memory 221, all in step 507. Until the SAVE button is activated, capture device 110 may be in a "wait state" 505." To the extent it is found that Sizer does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Sizer does not anticipate this claim, Sizer renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Sizer with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1B]	deriving at least one characteristic of the video stream;	Sizer discloses deriving at least one characteristic of the video stream. Sizer at 6:42-47

Element	'004 Claim Recitation	Exemplary Citations in Reference
		• "If the encoded non-perceptible data is inserted in the video portion of a television signal, the images displayed on a visual display area a monitor or television are captured by optical lens 117 and an associated photo-detector 203, which is arranged to supply an electrical signal to a filter/receiver 215 representing the image." Sizer at 9:33-53 • "For background purposes, a televised advertisement for a product could include additional ordering or "coupon" information embedded in a subliminal visual pattern within the video image displayed on a television. The coupon information describes the item(s) offered for sale, the regular price, the coupon value, and contains additional information needed for ordering. This embedded information could be transmitted between a television and capture device at a rate of about 60 bits per second, that is sufficient such that the necessary information can be transmitted in a very brief period of time. In the application just described, the embedded product information, dialing number, coupon information and necessary checkbits together comprise approximately 2,000 bits of information, so that between three and four seconds would be required for transmission. Once the coupon information is stored in capture device 110, the user of the capture device can supplement the coupon information with user entered data which can be entered using keypad 247. This supplemental information may include credit card or other billing information pertaining to the user, as an example." Sizer at 10:27-37 • In the active mode, a user operates capture device 110 in real time, typically while scanning a product or bar code, or in more complicated devices viewing a video display or listening to an audio presentation. Scanning a product or bar code, or in more complicated devices the user activates the SAVE button 241. When button 241 is activated, a positive result occurs in

Element	'004 Claim Recitation	Exemplary Citations in Reference
		step 503, thereby capturing the encoded data, displaying all or a portion of the data on display 225, and storing same in memory 221, all in step 507. Until the SAVE button is activated, capture device 110 may be in a "wait state" 505." To the extent it is found that Sizer does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Sizer does not anticipate this claim, Sizer renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Sizer with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1C]	recognizing the at least one object in the scene as a target object based at least in part on the at least one characteristic of the video stream;	Sizer discloses recognizing the at least one object in the scene as a target object based at least in part on the at least one characteristic of the video stream. Sizer at 6:42-47 • "If the encoded non-perceptible data is inserted in the video portion of a television signal, the images displayed on a visual display area a monitor or television are captured by optical lens 117 and an associated photo-detector 203, which is arranged to supply an electrical signal to a filter/receiver 215 representing the image." Sizer at 9:33-53 • "For background purposes, a televised advertisement for a product could include additional ordering or "coupon" information embedded in a subliminal visual pattern within the video image displayed on a television. The coupon information describes the item(s) offered for sale, the regular

Element	'004 Claim Recitation	Exemplary Citations in Reference
		price, the coupon value, and contains additional information needed for ordering. This embedded information could be transmitted between a television and capture device at a rate of about 60 bits per second, that is sufficient such that the necessary information can be transmitted in a very brief period of time. In the application just described, the embedded product information, dialing number, coupon information and necessary checkbits together comprise approximately 2,000 bits of information, so that between three and four seconds would be required for transmission. Once the coupon information is stored in capture device 110, the user of the capture device can supplement the coupon information with user entered data which can be entered using keypad 247. This supplemental information may include credit card or other billing information pertaining to the user, as an example." Sizer at 10:27-37 • In the active mode, a user operates capture device 110 in real time, typically while scanning a product or bar code, or in more complicated devices viewing a video display or listening to an audio presentation. Scanning a product or bar code, or in more complicated devices the user activates the SAVE button 241. When button 241 is activated, a positive result occurs in step 503, thereby capturing the encoded data, displaying all or a portion of the data on display 225, and storing same in memory 221, all in step 507. Until the SAVE button is activated, capture device 110 may be in a "wait state" 505." Sizer at 11:11-15 • In accordance with one optional arrangement of the present invention, when data is captured and stored in step 507, they newly collected data is compared with previously stored data. When a match is found, any duplicate data is dropped. In addition, the information in various fields is examined, to assure that the data captured is in the appropriate format. If an

Element	'004 Claim Recitation	Exemplary Citations in Reference
		error is detected in any data, the information may be deleted, and "new" data collected. This is accomplished by repeating capture step 507 several times. To the extent it is found that Sizer does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Sizer does not anticipate this claim, Sizer renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Sizer with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1D]	providing an indication upon recognizing the at least one object as the target object; and	Sizer discloses providing an indication upon recognizing the at least one object as the target object. Sizer at 6:30-38 • For certain transactions, information connected with a transaction may also be transmitted from service platform 130 to capture device 110, either for display to a user, storage for later use, or output to a point of sale system or device. This information may, for example, be confirmation information, indicating that a transaction has been effected, or details regarding a just completed transaction, such as ticket information relating to a travel reservation. Sizer at 7:12-15 • "The captured data may also be displayed on a display 225, such as a liquid crystal display, so that a user will be aware of exactly what

Element	**'004 Claim Recitation**	**Exemplary Citations in Reference**
		information was captured." Sizer at 7:42-46 • "When the data has been captured, the microprocessor 220 may control the display 225 to actually display some or all of the data, so that the user can be aware of information that had been subliminal and indiscernible. The display can also indicate that the device is 'aimed' correctly." To the extent it is found that Sizer does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Sizer does not anticipate this claim, Sizer renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Sizer with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1E]	presenting to a user, via an address, content information associated with the target object.	Sizer discloses presenting to a user, via an address, content information associated with the target object. Sizer at FIG. 1

Element	'004 Claim Recitation	Exemplary Citations in Reference
		FIG. 1 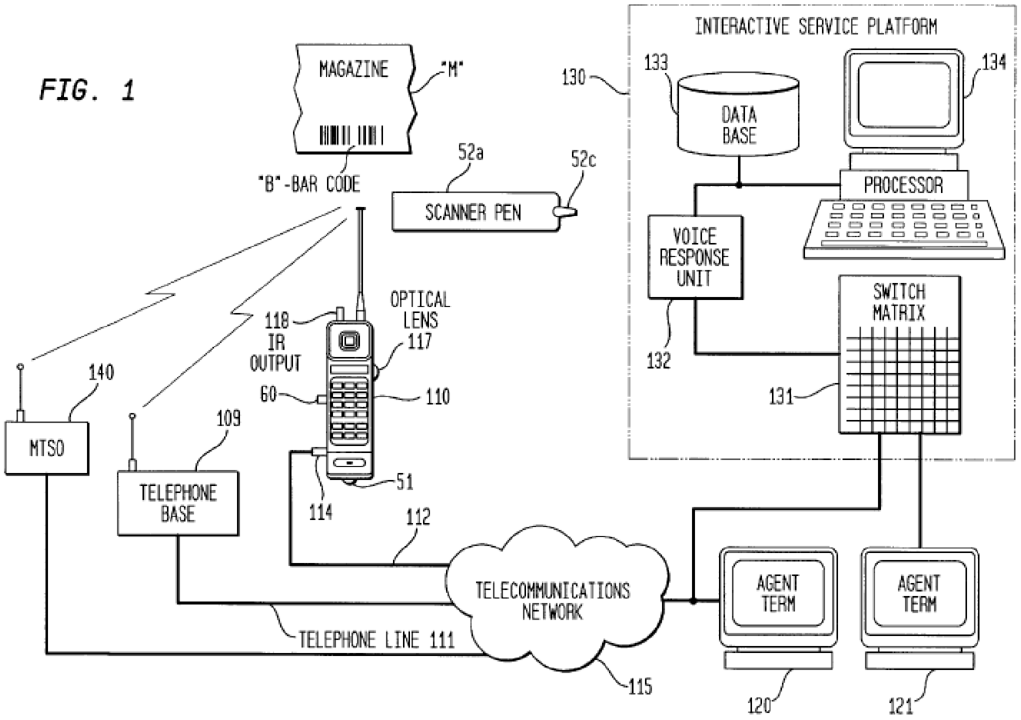 Sizer at 6:18-38 - "If an outgoing telephone call in initiated by capture device 110, the call can illustratively be routed through telecommunications network 115 directly to an agent terminal 120, or to an interactive service platform 130, which is connected to another agent terminal 121. When a connection is completed to platform 130, the call may be routed by a switch matrix 131 to a voice response unit 132 that operates under the control of a processor 134. Information included in the captured encoded can be provided to the

Element	'004 Claim Recitation	Exemplary Citations in Reference
		attendants, or processor 134 can perform a look-up and retrieval operation in database 131, using encoded data as a key, in order to effectuate a desired transaction. For certain transactions, information connected with a transaction may also be transmitted from service platform 130 to capture device 110, either for display to a user, storage for later use, or output to a point of sale system or device. This information may, for example, be confirmation information, indicating that a transaction has been effected, or details regarding a just completed transaction, such as ticket information relating to a travel reservation." Sizer at 10:45-49 • "A call is then originated launch, such by transmitting the dialing information from memory 221 through base station interface 250 to base station 109 and thence to telephone line 111 and telecommunications network 115 to interactive service platform 130." Sizer at 10:59-63 • After the call has been originated, a determination is made in step 515 that the call has been answered and that the SEND DATA button 243 has been pushed. (Until the SEND DATA button is activated, capture device 110 may be in a "wait state" 517.). When a positive result occurs in step 515, information contained in fields 403-411 is transmitted to the called destination in step 519. As stated previously, from the point of view of the called party, the information received includes details concerning the advertised product, the identity of the caller, and other information associated with the item that was initially displayed or mentioned on the television program. Advantageously, the advertiser will know more about their customer that ever before. This is because the advertiser will be able to determine from the transmitted information which advertisement draw the business, how long it took the caller to respond to the ad, and the

Element	'004 Claim Recitation	Exemplary Citations in Reference
		geographic location of the callers, based upon ANI or called ID. It will be easy to tell which ads in a campaign are most successful. Sizer at 12:41-43 • "The call may be placed to an interactive service platform, and additional captured data, such as coupon or similar information, can then be used to effectuate a transaction." To the extent it is found that Sizer does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Sizer does not anticipate this claim, Sizer renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Sizer with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

2. Claim 6

Claim	'004 Claim Recitation	Exemplary Citations in Reference
6	The method of claim 1, wherein the at least one object is identified as the target object based at least in part on at least one user input.	Sizer discloses the at least one object is identified as the target object based at least in part on at least one user input. Sizer at 9:47-53 • Once the coupon information is stored in capture device 110, the user of the capture device can supplement the coupon information with user entered data which can be entered using keypad 247. This supplemental

		information may include credit card or other billing information pertaining to the user, as an example. Sizer at 10:27-37 • "In the active mode, a user operates capture device 110 in real time, typically while scanning a product or bar code, or in more complicated devices viewing a video display or listening to an audio presentation. Scanning a product or bar code, or in more complicated devices the user activates the SAVE button 241. When button 241 is activated, a positive result occurs in step 503, thereby capturing the encoded data, displaying all or a portion of the data on display 225, and storing same in memory 221, all in step 507. Until the SAVE button is activated, capture device 110 may be in a "wait state" 505." Sizer at 10:38-49 • "The user of capture device 110 determines that the appropriate data has been captured by reviewing the information seen in display 225. After the captured data has been reviewed and determined to be correct, activation of the DIAL button 242 on the portable capture device 110 in step 509 initiates step 513, by which stored dialing and/or routing information, illustratively in field 401 of FIG. 4, is recalled from memory 221. A call is then originated launch, such by transmitting the dialing information from memory 221 through base station interface 250 to base station 109 and thence to telephone line 111 and telecommunications network 115 to interactive service platform 130. The call may be a local or long distance call made using a "plain old telephone service" (POTS) line; alternatively, a toll-free call may be launched using an 800 number. Desirably, automatic number identification (ANI), called ID, and other features

		available in the telecommunications network may be used to provide additional information to the called party, in addition to the captured data that is transmitted in succeeding steps in the process. Until the DIAL button is activated, capture device 110 may be in a "wait state" 511." To the extent it is found that Sizer does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Sizer does not anticipate this claim, Sizer renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Sizer with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

3. Claim 18

Claim	'004 Claim Recitation	Exemplary Citations in Reference
18	The method of claim 1, further comprising communicating at least one of banking account information, user data and terminal specific data to a banking company.	Sizer discloses communicating at least one of banking account information, user data and terminal specific data to a banking company. Sizer at 9:33-53 • "For background purposes, a televised advertisement for a product could include additional ordering or "coupon" information embedded in a subliminal visual pattern within the video image displayed on a television. The coupon information describes the item(s) offered for sale, the regular price, the coupon value, and contains additional information

		needed for ordering. This embedded information could be transmitted between a television and capture device at a rate of about 60 bits per second, that is sufficient such that the necessary information can be transmitted in a very brief period of time. In the application just described, the embedded product information, dialing number, coupon information and necessary checkbits together comprise approximately 2,000 bits of information, so that between three and four seconds would be required for transmission. Once the coupon information is stored in capture device 110, the user of the capture device can supplement the coupon information with user entered data which can be entered using keypad 247. This supplemental information may include credit card or other billing information pertaining to the user, as an example." Sizer at 12:62-13:12 • In accordance with the present invention, a capture device now can initiate a transaction by capturing transaction data from marks contained on an object, such as a magazine advertisement, where the marks have embedded therein a code corresponding to transaction data for initiating a transaction. The capture device includes a scanner, operable by a user, for reading marks contained on the object and a controller for interpreting the marks and retrieving transaction data embedded in the marks. Through associated apparatus as described above, a telephone call can be originated by information contained in the transaction data in at least a portion of the transaction data transferred to a desired destination for initiating a transaction. As a result, it is now possible now for magazines, newspapers, printed media and many other objects to contain marks relating to an advertisement or solicitation which can be easily scanned and a transaction completed to the vendor of the advertisement or solicitation. This saves much time and allows easier access to services and products.

		To the extent it is found that Sizer does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Sizer does not anticipate this claim, Sizer renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Sizer with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

Nantworks, LLC v. Bank of America Corporation
Central District of California, No. 2:20-cv-7872
Invalidity Chart for U.S. Patent 9,324,004 based on
U.S. Patent No. 6,666,377 to Harris ("Harris")

Exhibit A-6

Bank of America Corporation ("BofA" or "Defendant") provides this chart subject to all reservations, objections, statements, and disclaimers set forth herein and in BofA's Final Invalidity Contentions Cover Pleading, as well as any amendment, supplement, or modification thereof, which are incorporated herein by reference in their entirety.

On information and belief, the application that matured into Harris was filed on July 18, 2000. Harris published at least as of December 23, 2003, and is available as prior art at least under 35 U.S.C. § 102(e).

As shown in the chart below, Harris anticipates asserted claims 1, 2, 3, 6, 11, 15, 18 of the '004 Patent. To the extent Harris is found not to expressly disclose certain limitations in the asserted claims, such limitations are inherent. To the extent Harris is found not to anticipate any asserted claims or claim elements of the '004 Patent, the reference nevertheless renders those claims or claim elements obvious under 35 U.S.C. § 103, either alone or in combination with other art identified herein or in the cover pleading of BofA's Final Invalidity Contentions, for the reasons identified in the cover pleading. These Final Invalidity Contentions are not an admission by BofA that the accused products, including any current or past versions of the accused products, are covered by, or infringe the asserted claims, but are based instead on the recognition that if the claims are interpreted to be broad enough to encompass the accused products, the claims must also be construed to have that same scope when considering whether they are invalid.

The following chart is based on the Court's claim construction rulings (Dkt. Nos. 145, 153, 236), and to the extent not contradictory, Plaintiffs' apparent interpretations of the claims in Plaintiffs' "Preliminary" Final Infringement Contentions. BofA notes that in many instances, Plaintiffs' "Preliminary" Final Infringement Contentions ignore and/or contradict the Court's claim construction rulings. BofA does not accept the assumptions concerning the scope and meaning implicit in Plaintiffs' Preliminary Final Infringement Contentions, to the extent those assumptions are discernible, and reserves the right to challenge Plaintiffs' proposed (or implied) constructions. Furthermore, Plaintiffs' "Preliminary" Final Contentions fail to comply with the level of specificity required by SPR 2.1.2 and 2.1.3, and include improper reservation of rights to supplement with "additional ways in which BoA Accused Products infringe." Any alleged deficient discovery is of Plaintiffs own making. Most if not all of the deficiencies in Plaintiffs' Preliminary Final Contentions have been documented by Defendants through multiple letters since at least May 18, 2021 and up to Bank of America's most recent letter of August 19, 2023. Plaintiffs' shifting infringement contentions, deficiencies in its infringement contentions, and changing claim scope have prejudiced BofA from being able to reasonably prepare its defenses, including its invalidity contentions. BofA also reserves the right to revise and supplement these charts if and when Plaintiffs are permitted to provide revised or supplemental Infringement Contentions. Citations given in the chart below are merely representative of the respective elements to provide notice to Plaintiffs of BofA's invalidity theories and are not meant to be exhaustive.

1. Claim 1

Element	'004 Claim Recitation	Exemplary Citations in Reference
[1Pre]	A method for processing a video stream, comprising:	Harris discloses a method: Harris at 2: 3-14 "The bar codes can be imaged/scanned in a number of different ways. One embodiment discloses using a camera to input and decode these bar codes. The embodiments are shown in FIGS. 1A and 1B. FIG. 1A describes using a personal digital assistant 100 as the input device. One preferred input device of this type is the Palm V(TM) type hand held computer. A bar code scanner can be used, such as the commercially available Symbol Technology SPT 1700. Alternatively, a camera add-on unit can be added to the Palm V and used as described herein. The PDA includes a screen 145 and user interface 146 all within the same housing 99."
[1A]	analyzing, via a mobile device, a scene represented by the video stream for at least one object;	Harris discloses analyzing, via a mobile device, a scene represented by the video stream for at least one object. Harris at FIG. 1A and FIG. 1B

Element	'004 Claim Recitation	Exemplary Citations in Reference
		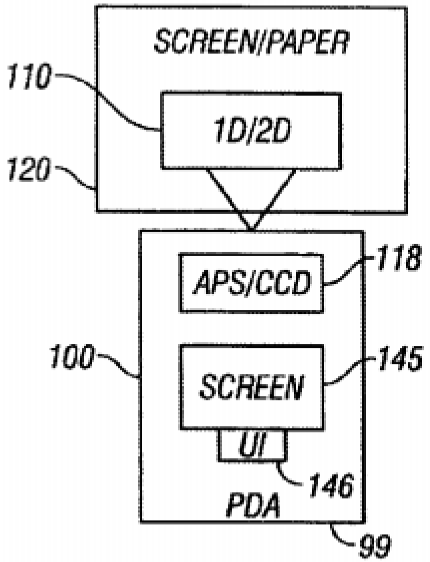 FIG. 1A 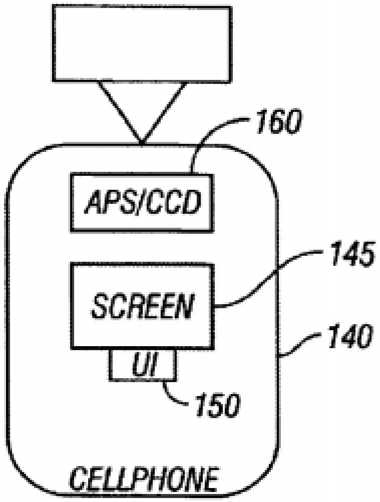 FIG. 1B Harris at 2: 3-14 • "The bar codes can be imaged/scanned in a number of different ways. One embodiment discloses using a camera to input and decode these bar codes. The embodiments are shown in FIGS. 1A and 1B. FIG. 1A describes using a personal digital assistant 100 as the input device. One preferred input device of this type is the Palm V(TM) type hand held computer. A bar code scanner can be used, such as the commercially available Symbol Technology SPT 1700. Alternatively, a camera add-on unit can be added to the Palm V and used as described herein. The PDA includes a screen 145 and user interface 146 all within the same housing 99." Harris at FIG. 4

Element	'004 Claim Recitation	Exemplary Citations in Reference
		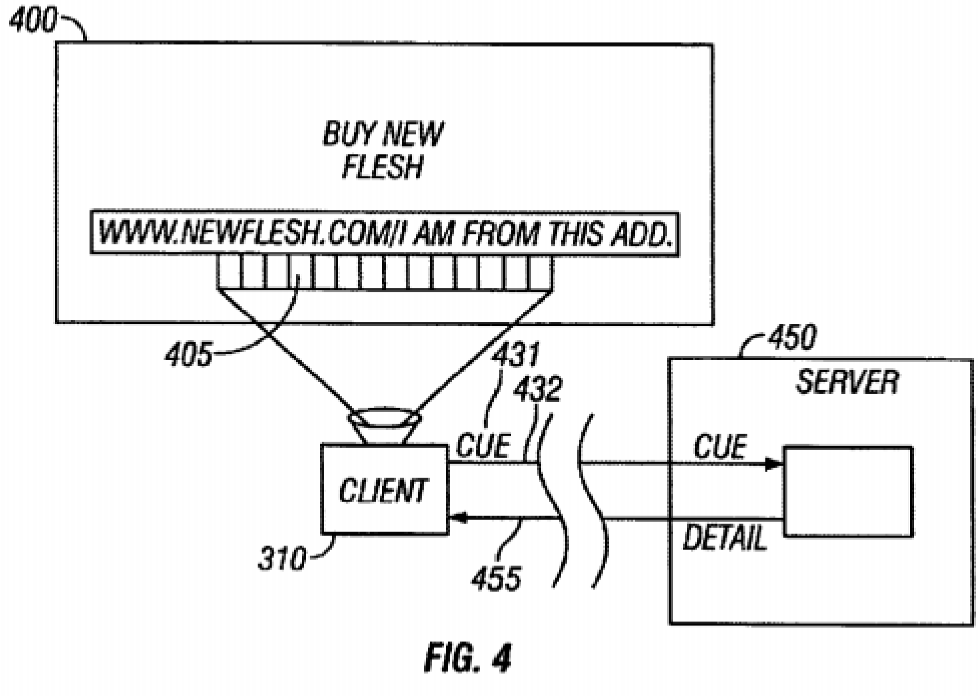 FIG. 4 Harris at 4: 9-15 • "Another embodiment shown in FIG. 4 uses bar codes to enter information into the computer. The FIG. 4 embodiment stores scannable non-alphanumeric information, e.g. bar code information, as some part of a communication—here an advertisement. The advertisement can be a print advertisement, a television advertisement, or an Internet advertisement for example." To the extent it is found that Harris does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Harris does not anticipate

Element	'004 Claim Recitation	Exemplary Citations in Reference
		this claim, Harris renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Harris with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1B]	deriving at least one characteristic of the video stream;	Harris discloses deriving at least one characteristic of the video stream. Harris at FIG. 4

Element	'004 Claim Recitation	Exemplary Citations in Reference
		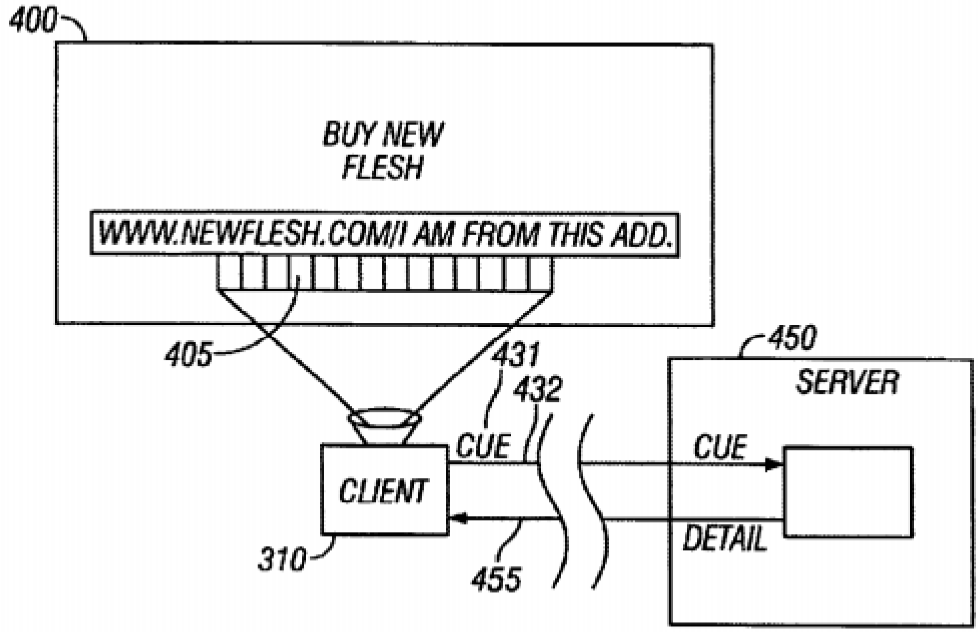 FIG. 4 Harris at 4:15-30 • "The advertisement 400 includes a bar code 405 therein. The bar code 405 is associated with the advertisement, and includes some additional information about the advertisement. For example, the bar code may include the web site address of the company preparing the advertisement, or appointment information about the advertisement, or a "vcf file" or the like which is an importable file with address information about the company or author sponsoring the advertisement. In operation, the user brings one of the clients, either the cellular telephone 310 or PDA 315, into range of the bar code 405. The client reads the bar code and decodes it as noted above. The decoded information can represent

Element	'004 Claim Recitation	Exemplary Citations in Reference
		ASCII information, a compressed file such as a zip file, G code information, or any other compressed or non-compressed information." To the extent it is found that Harris does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Harris does not anticipate this claim, Harris renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Harris with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1C]	recognizing the at least one object in the scene as a target object based at least in part on the at least one characteristic of the video stream;	Harris discloses recognizing at least one object in the scene as a target object based at least in part on the at least on characteristic of the video stream. Harris at FIG. 4

Element	'004 Claim Recitation	Exemplary Citations in Reference
		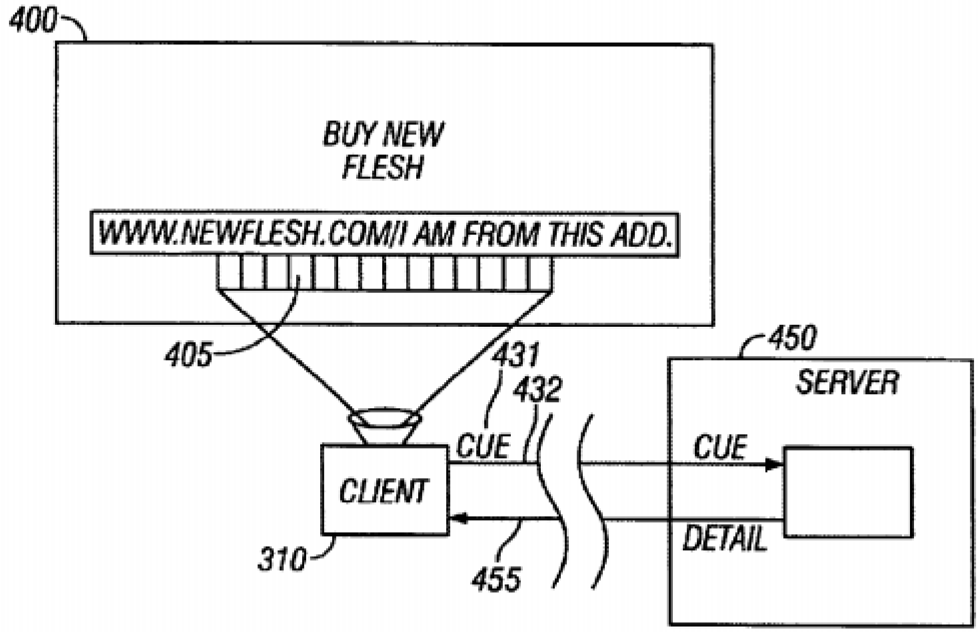 FIG. 4 Harris at 4:15-30 - "The advertisement 400 includes a bar code 405 therein. The bar code 405 is associated with the advertisement, and includes some additional information about the advertisement. For example, the bar code may include the web site address of the company preparing the advertisement, or appointment information about the advertisement, or a "vcf file" or the like which is an importable file with address information about the company or author sponsoring the advertisement. In operation, the user brings one of the clients, either the cellular telephone 310 or PDA 315, into range of the bar code 405. The client reads the bar code and decodes it as noted above. The decoded information can represent

Element	**'004 Claim Recitation**	**Exemplary Citations in Reference**
		ASCII information, a compressed file such as a zip file, G code information, or any other compressed or non-compressed information." To the extent it is found that Harris does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Harris does not anticipate this claim, Harris renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Harris with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1D]	providing an indication upon recognizing the at least one object as the target object; and	Harris discloses providing an indication upon recognizing the at least one object as the target object. Harris at FIG. 4

Element	'004 Claim Recitation	Exemplary Citations in Reference
		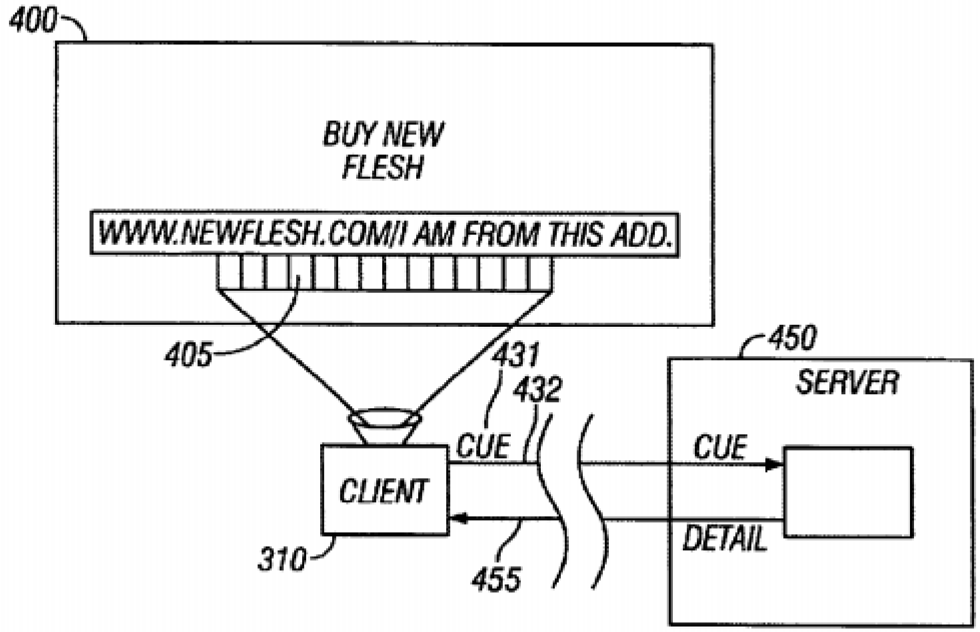 FIG. 4 Harris at 4:30-40 • "The contents are automatically input into the client. The contents can directly represent that information, in which case the information is input into the client. For example, the information can directly represent the ASCII information indicating the website. Alternatively, the information can represent a pointer to a database, e.g. a publicly available database. This database can later be accessed as part of an information transfer. For example, Palm systems enable a hot sync where the portable computer is synced with another computer that is running hot sync software."

Element	'004 Claim Recitation	Exemplary Citations in Reference
		To the extent it is found that Harris does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Harris does not anticipate this claim, Harris renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Harris with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1E]	presenting to a user, via an address, content information associated with the target object.	Harris discloses presenting to a user, via an address, content information associated with the target object. Harris at FIG. 4

Element	'004 Claim Recitation	Exemplary Citations in Reference
		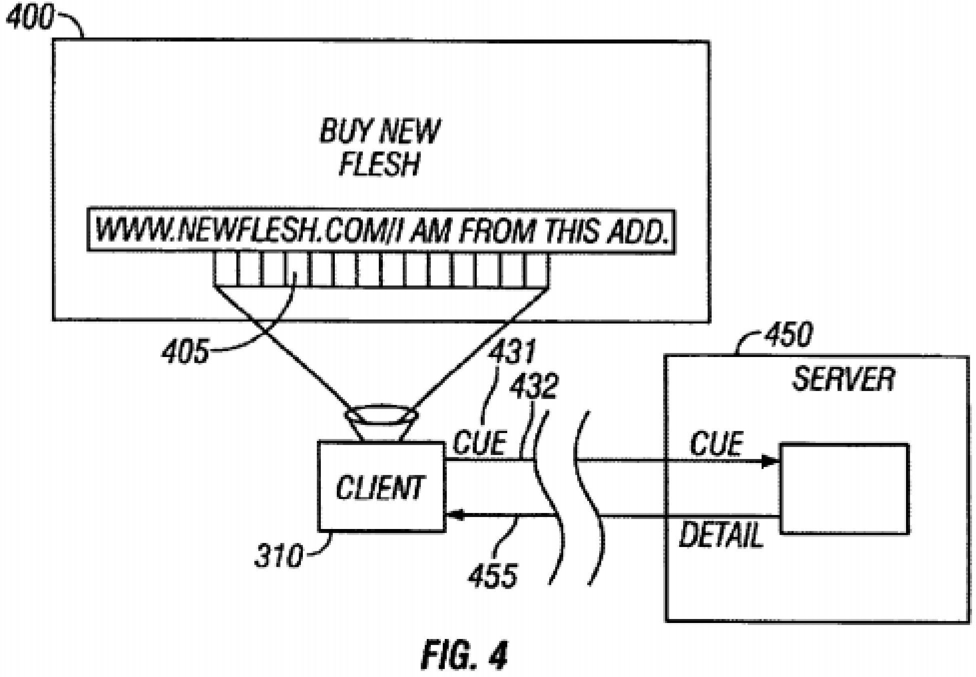 FIG. 4 Harris at 4:41-48 • "In this system, the information can represent a cue, e.g. a number, that is associated with a special function. The cue is used during a hot sync to access a publicly available information network, e.g. by direct connection or by the Internet. The cue 431 is sent over the Internet 432 to the server 450, and addresses more detailed information in a memory of the server. The server returns that information as 455, and the client receives the more detailed information." To the extent it is found that Harris does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Harris does not anticipate

Element	'004 Claim Recitation	Exemplary Citations in Reference
		this claim, Harris renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Harris with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

2. Claim 6

Claim	'004 Claim Recitation	Exemplary Citations in Reference
6	The method of claim 1, wherein the at least one object is identified as the target object based at least in part on at least one user input.	Harris discloses the method of claim 1, wherein the at least one object is identified as the target object based at least in part on at least one user input. Harris at FIG. 4

		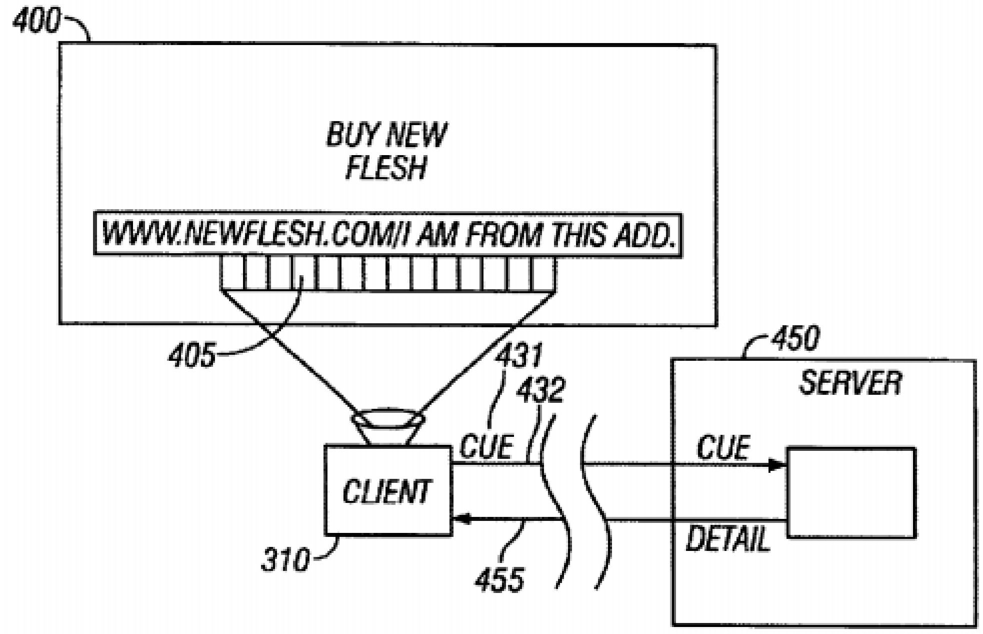 FIG. 4 Harris at 4:15-30 - "The advertisement 400 includes a bar code 405 therein. The bar code 405 is associated with the advertisement, and includes some additional information about the advertisement. For example, the bar code may include the web site address of the company preparing the advertisement, or appointment information about the advertisement, or a "vcf file" or the like which is an importable file with address information about the company or author sponsoring the advertisement. In operation, the user brings one of the clients, either the cellular telephone 310 or PDA 315, into range of the bar code 405. The client reads the bar code and decodes it as noted above. The decoded information can represent ASCII information, a compressed file such as a

		zip file, G code information, or any other compressed or non-compressed information." To the extent it is found that Harris does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Harris does not anticipate this claim, Harris renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Harris with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

3. Claim 18

Claim	'004 Claim Recitation	Exemplary Citations in Reference
18	The method of claim 1, further comprising communicating at least one of banking account information, user data and terminal specific data to a banking company.	To the extent it is found that Harris does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Harris does not anticipate this claim, Harris renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Harris with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

Nantworks, LLC v. Bank of America Corporation
Central District of California, No. 2:20-cv-7872
Invalidity Chart for U.S. Patent 9,324,004 based on
U.S. Patent No. 6,722,569 to Ehrhart et al. (“Ehrhart”)

Exhibit A-11

Bank of America Corporation (“BofA” or “Defendant”) provides this chart subject to all reservations, objections, statements, and disclaimers set forth herein and in BofA’s Final Invalidity Contentions Cover Pleading, as well as any amendment, supplement, or modification thereof, which are incorporated herein by reference in their entirety.

On information and belief, the application that matured into Ehrhart was filed on July 13, 2001, claiming the priority date of July 13, 2001. Ehrhart published at least as of April 20, 2004, and is available as prior art at least under 35 U.S.C. § 102(e).

As shown in the chart below, Ehrhart anticipates asserted claims 1, 6, 18 of the ’004 Patent. To the extent Ehrhart is found not to expressly disclose certain limitations in the asserted claims, such limitations are inherent. To the extent Ehrhart is found not to anticipate any asserted claims or claim elements of the ’004 Patent, the reference nevertheless renders those claims or claim elements obvious under 35 U.S.C. § 103, either alone or in combination with other art identified in the cover pleading of BofA’s Invalidity Contentions, for the reasons identified in the cover pleading. These Invalidity Contentions are not an admission by BofA that the accused products, including any current or past versions of the accused products, are covered by, or infringe the asserted claims, but are based instead on the recognition that if the claims are interpreted to be broad enough to encompass the accused products, the claims must also be construed to have that same scope when considering whether they are invalid.

The following chart is based on the Court’s claim construction rulings (Dkt. Nos. 145, 153, 236), and to the extent not contradictory, Plaintiffs’ apparent interpretations of the claims in Plaintiffs’ “Preliminary” Final Infringement Contentions. BofA notes that in many instances, Plaintiffs’ “Preliminary” Final Infringement Contentions ignore and/or contradict the Court’s claim construction rulings. BofA does not accept the assumptions concerning the scope and meaning implicit in Plaintiffs’ Preliminary Final Infringement Contentions, to the extent those assumptions are discernible, and reserves the right to challenge Plaintiffs’ proposed (or implied) constructions. Furthermore, Plaintiffs’ “Preliminary” Final Contentions fail to comply with the level of specificity required by SPR 2.1.2 and 2.1.3, and include improper reservation of rights to supplement with “additional ways in which BoA Accused Products infringe.” Any alleged deficient discovery is of Plaintiffs own making. Most if not all of the deficiencies in Plaintiffs’ Preliminary Final Contentions have been documented by Defendants through multiple letters since at least May 18, 2021 and up to Bank of America’s most recent letter of August 19, 2023. Plaintiffs’ shifting infringement contentions, deficiencies in its infringement contentions, and changing claim scope have prejudiced BofA from being able to reasonably prepare its defenses, including its invalidity contentions. BofA also reserves the right to revise and supplement these charts if and when Plaintiffs are permitted to provide revised or supplemental Infringement Contentions. Citations given in the chart below are merely representative of the respective elements to provide notice to Plaintiffs of BofA’s invalidity theories and are not meant to be exhaustive.

1. Claim 1

Element	'004 Claim Recitation	Exemplary Citations in Reference
[1Pre]	A method for processing a video stream, comprising:	Ehrhart discloses a method for processing a video stream. Ehrhart at Abstract. • "The present invention relates to an optical reader that includes a color imaging assembly that generates color imaging data. An image analysis circuit determines if the acquired image should be characterized as a color photograph or as including a graphical symbol. A processing circuit processes the imaging data based on the image analysis circuit's determination of whether the image is a graphical symbol or a color photograph. The present invention allows a user to acquire and process both color images and graphical symbols, such as bar codes, text, OCR symbols or signatures. The optical reader of the present invention is also configured to associate an acquired image with at least one other acquired image." Ehrhart at 2:62-64. • "In another aspect, the present invention includes a method for acquiring an image of an object with an optical reader."
[1A]	analyzing, via a mobile device, a scene represented by the video stream for at least one object;	Ehrhart discloses analyzing, via a mobile device, a scene represented by the video stream for at least one object. *Id.* at [1Pre]. Ehrhart at 1:57-2:16.

Element	'004 Claim Recitation	Exemplary Citations in Reference
		- "The optical reader of the present invention automatically, or through manual Selection, determines whether a captured image is a color photographic image or, a color image that includes a graphical Symbol. Subsequently, the optical reader of the present invention processes the acquired imaging data in accordance with that determination. The optical reader of the present invention is operative to acquire and associate a plurality of acquired images. One aspect of the present invention is an optical reader. The optical reader includes a color imaging assembly for acquiring an image of an object, the color imaging assembly generating imaging data corresponding to the image. An image analysis circuit is coupled to the color imaging assembly. The image analysis circuit being configured to determine if the color imaging data includes at least one graphical Symbol. The image is classified as a graphical Symbol, or the image is classified as a color photograph if the color imaging data does not include at least one graphical Symbol. A processing circuit is coupled to the image analysis circuit. The processing circuit is operative to process the imaging data based on the determination. In another aspect, the present invention includes an optical reader for capturing an image of an object. The optical reader includes a color imaging assembly for converting the image of the object into color digital data corresponding to the image." Ehrhart at 11:46-12:18 - "As embodied herein and depicted in FIG. 8, a flow chart showing a method for performing the 1D autodiscrimination of step 702 in FIG. 7 is disclosed. In step 800 processor 40 calculates the activities of selected image data elements. The activity is defined as a measure of the rate of change of the image data over a small two-dimensional portion of the region surrounding the selected data element. In one embodiment, the activity is calculated along any two arbitrarily selected directions which are orthogonal one to the other. Two

Element	'004 Claim Recitation	Exemplary Citations in Reference
		mutually perpendicular directions are used because the orientation of the symbol is unknown. In step 802, processor 40 looks for "high activity" regions. These high activity regions are referred to as candidate symbol regions(CSRs). A high activity region indicates a transition from a black region to a white region, or vice-versa. If there is more than one CSR, it may indicate the presence of more than one bar code symbol. In step 804, processor 40 selects the largest CSR. In step 806, processor 40 calculates the centroid of the largest CSR. Subsequently, processor 40 finds the direction of the highest activity in the largest CSR. In a 1D bar code, this will be the direction perpendicular to the direction of the bars. In steps 810 and 812, processor defines the initial scan line(SC=0), as being the scan line bisecting the centroid of the bar code. Processor calculates the brightness values of sampling points along the initial scan line. These brightness values are converted to digital data in step 816. In decoding step 818, processor 40 applies one 1D decoding program after another. If decoding is unsuccessful, processor 40 checks if the entire CSR has been scanned. If not, it establishes a new scan line, and repeats the decoding process. If in step 822, the entire CSR has been scanned, and there are no CSRs remaining to be decoded, processor 40 exits the routine. If in step 820, 1D decoding is successful, processor 40 determines if the symbol is a 1D stacked symbol. If it is a 1D stacked symbol, processor 40 scans and decodes the remaining CSRs in the stacked symbol. If it is not a stacked symbol, the decoded 1D data is stored or output to display 60 in step 830. In step 838, processor 40 determines if there any unexamined regions. If there are unexamined regions, the decoding process is repeated. Otherwise, processor 40 exits the routine." To the extent it is found that Ehrhart does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Ehrhart does not anticipate this claim, Ehrhart renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified in the cover pleading of BofA's Invalidity

Element	'004 Claim Recitation	Exemplary Citations in Reference
		Contentions. A person of ordinary skill in the art would have been motivated to combine Ehrhart with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1B]	deriving at least one characteristic of the video stream;	Ehrhart discloses deriving at least one characteristic of the video stream. *Id.* at [1Pre-A]. Ehrhart at 2:26-42: • "In another aspect, the present invention includes an optical reader for capturing an image of an object. The optical reader includes a color imaging assembly for capturing the image as color imaging data. A classification circuit is coupled to the color imaging assembly, the classification circuit being configured to process at least a portion of the color imaging data to thereby Select one of a plurality of classifications, whereby the image is classified as a color photographic image, or as an image that includes at least one graphical Symbol. An automatic mode Selector is coupled to the classification circuit, the automatic mode Selector being configured to Select an optical reader mode in accordance with the Selected classification. A processor is coupled to the classification circuit, the processor being programmed to process the color imaging data in accordance with the optical reader mode Selected by the automatic mode Selector." Ehrhart at 5:36-50: • "In accordance with the invention, the present invention for an optical reader includes a color imaging assembly for acquiring color imaging data. An image analysis circuit determines if the acquired image includes at least one graphical symbol. A processing circuit processes the imaging data based on the determination of whether the image

Element	'004 Claim Recitation	Exemplary Citations in Reference
		includes at least one graphical symbol. The present invention allows a user to read graphical symbols, such as bar codes, text, OCR characters or signatures using a color imager. The color optical reader of the present invention is configured to automatically determine whether a color image includes a graphical symbol, or is merely a color photographic image. The optical reader of the present invention also is operative to associate one acquired image with at least one subsequently acquired image." Ehrhart at 11:46-12:18. • "As embodied herein and depicted in FIG. 8, a flow chart showing a method for performing the 1D autodiscrimination of step 702 in FIG. 7 is disclosed. In step 800 processor 40 calculates the activities of selected image data elements. The activity is defined as a measure of the rate of change of the image data over a small two-dimensional portion of the region surrounding the selected data element. In one embodiment, the activity is calculated along any two arbitrarily selected directions which are orthogonal one to the other. Two mutually perpendicular directions are used because the orientation of the symbol is unknown. In step 802, processor 40 looks for "high activity" regions. These high activity regions are referred to as candidate symbol regions(CSRs). A high activity region indicates a transition from a black region to a white region, or vice-versa. If there is more than one CSR, it may indicate the presence of more than one bar code symbol. In step 804, processor 40 selects the largest CSR. In step 806, processor 40 calculates the centroid of the largest CSR. Subsequently, processor 40 finds the direction of the highest activity in the largest CSR. In a 1D bar code, this will be the direction perpendicular to the direction of the bars. In steps 810 and 812, processor defines the initial scan line(SC=0), as being the scan line bisecting the centroid of the bar code. Processor calculates the brightness values of sampling points along the initial scan line. These brightness values are converted to digital data in step 816. In decoding step 818, processor 40 applies one 1D decoding program after another. If decoding is unsuccessful, processor 40 checks if the entire CSR has

Element	'004 Claim Recitation	Exemplary Citations in Reference
		been scanned. If not, it establishes a new scan line, and repeats the decoding process. If in step 822, the entire CSR has been scanned, and there are no CSRs remaining to be decoded, processor 40 exits the routine. If in step 820, 1D decoding is successful, processor 40 determines if the symbol is a 1D stacked symbol. If it is a 1D stacked symbol, processor 40 scans and decodes the remaining CSRs in the stacked symbol. If it is not a stacked symbol, the decoded 1D data is stored or output to display 60 in step 830. In step 838, processor 40 determines if there any unexamined regions. If there are unexamined regions, the decoding process is repeated. Otherwise, processor 40 exits the routine." To the extent it is found that Ehrhart does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Ehrhart does not anticipate this claim, Ehrhart renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Ehrhart with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1C]	recognizing the at least one object in the scene as a target object based at least in part on the at least one characteristic of the video stream;	Ehrhart discloses recognizing the at least one object in the scene as a target object based at least in part on the at least one characteristics of the video stream. *Id.* at [1Pre-B]. Ehrhart at 8:36-49. • "In the Automatic imaging mode, processor 40 is programmed to analyze the color imaging data to determine if an acquired image includes a graphical

Element	'004 Claim Recitation	Exemplary Citations in Reference
		symbol or is merely a color photographic image. If it makes the determination that the color image includes a graphical symbol, it further analyzes the acquired image and classifies it as a bar code, OCR symbol, text, or a signature. Based on the classification, optical reader 10 jumps to the appropriate routine in EROM 48. The semi-automatic mode is similar. Thus, in the automatic or semi-automatic modes, the bar code scanning mode, the OCR/text mode, the signature capture mode, the color photography mode, and the association mode are controlled by the application program, not by the user." Ehrhart at 10:5-20. • "If processor 40 does not locate a bar code symbol it moves onto step 420 and looks for OCR-A or OCR-B characters. If it finds these characters it performs optical character recognition in step 422. If it does not, processor evaluates the image for the presence of text. If text is located, the image is cropped, and the text is compressed and stored in steps 428 and 430. If the image does not include text, processor 40 evaluates the image for the presence of a signature. If one is present, the image is cropped, and the data is compressed and stored in steps 428 and 430. In another embodiment, optical reader 10 is networked, and processor 40 communicates with remote network resources to provide signature verification services. If processor 40 cannot detect a bar code symbol, OCR symbols, text, or a signature, the user is asked in step 432 if he desires to store the image. If he does, the color photographic image is stored in memory in step 434." Ehrhart at 11:46-12:18. • "As embodied herein and depicted in FIG. 8, a flow chart showing a method for performing the 1D autodiscrimination of step 702 in FIG. 7 is disclosed. In step 800 processor 40 calculates the activities of selected image data elements. The activity is defined as a measure of the rate of change of the image data over a small two-dimensional portion of the region surrounding the selected

Element	'004 Claim Recitation	Exemplary Citations in Reference
		data element. In one embodiment, the activity is calculated along any two arbitrarily selected directions which are orthogonal one to the other. Two mutually perpendicular directions are used because the orientation of the symbol is unknown. In step 802, processor 40 looks for "high activity" regions. These high activity regions are referred to as candidate symbol regions(CSRs). A high activity region indicates a transition from a black region to a white region, or vice-versa. If there is more than one CSR, it may indicate the presence of more than one bar code symbol. In step 804, processor 40 selects the largest CSR. In step 806, processor 40 calculates the centroid of the largest CSR. Subsequently, processor 40 finds the direction of the highest activity in the largest CSR. In a 1D bar code, this will be the direction perpendicular to the direction of the bars. In steps 810 and 812, processor defines the initial scan line(SC=0), as being the scan line bisecting the centroid of the bar code. Processor calculates the brightness values of sampling points along the initial scan line. These brightness values are converted to digital data in step 816. In decoding step 818, processor 40 applies one 1D decoding program after another. If decoding is unsuccessful, processor 40 checks if the entire CSR has been scanned. If not, it establishes a new scan line, and repeats the decoding process. If in step 822, the entire CSR has been scanned, and there are no CSRs remaining to be decoded, processor 40 exits the routine. If in step 820, 1D decoding is successful, processor 40 determines if the symbol is a 1D stacked symbol. If it is a 1D stacked symbol, processor 40 scans and decodes the remaining CSRs in the stacked symbol. If it is not a stacked symbol, the decoded 1D data is stored or output to display 60 in step 830. In step 838, processor 40 determines if there any unexamined regions. If there are unexamined regions, the decoding process is repeated. Otherwise, processor 40 exits the routine." Ehrhart at 13:1-9. • "As embodied herein and depicted in FIG. 11, a flow chart showing a method for performing OCR in accordance with yet another embodiment of the present

Element	'004 Claim Recitation	Exemplary Citations in Reference
		invention is disclosed. In step 1100, reader 10 produces a bit-mapped image of the page. Subsequently, processor 40 finds lines of text in the image, locates the white spaces in each line, and isolates the characters. In step 1108, processor 40 performs character recognition, either OCR-A or OCR-B, as desired. The decoded characters are stored in memory." To the extent it is found that Ehrhart does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Ehrhart does not anticipate this claim, Ehrhart renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Ehrhart with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1D]	providing an indication upon recognizing the at least one object as the target object; and	Ehrhart discloses providing an indication upon recognizing the at least on object as the target object. *Id.* at [1Pre-C]. Ehrhart at 8:58-62: • "The user may also click on OCR/Text icon 660. Clicking icon 660 provides the user with a check validation mode, a text scanning mode, or a bi-tonal image capture mode. The check validation mode is performed in conjunction with network services." Ehrhart at 8:64-9:10:

Element	'004 Claim Recitation	Exemplary Citations in Reference
		• "In one embodiment, this mode includes a Signature verification program wherein the user may select between a Static verification or a dynamic verification. In the Static mode, the user captures the image of a Signature. The captured image is compared with a reference image Stored in a remote database. In the dynamic mode, optical reader 10 uses the Stylus and Signature block to capture the Signature. In this mode, Signature block 62 measures unique dynamic parameters, Such as applied pressure, direction and timing of movements, or a combination of these parameters. One of ordinary skill in the art will recognize that this list is not meant to be all-inclusive, but rather, is a representative example. The captured dynamic parameters are compared with a reference data Stored in a remote database." Ehrhart at Fig. 17.

		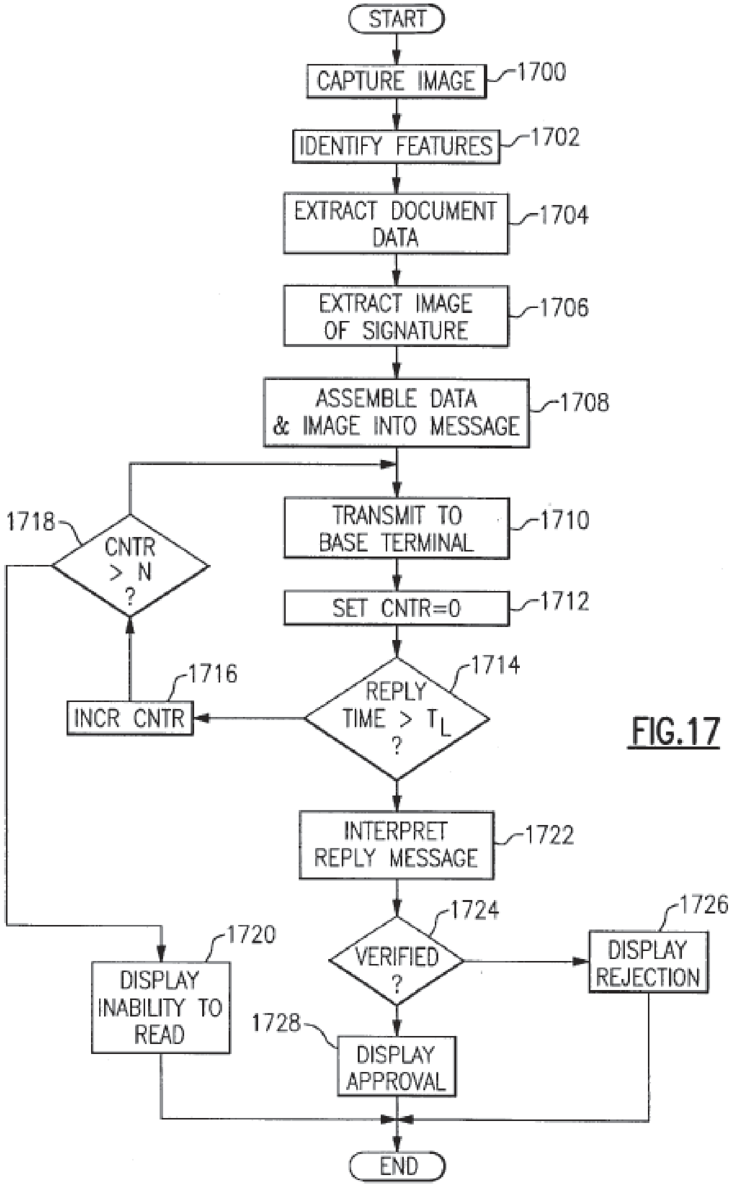START CAPTURE IMAGE 1700 IDENTIFY FEATURES 1702 EXTRACT DOCUMENT DATA 1704 EXTRACT IMAGE OF SIGNATURE 1706 ASSEMBLE DATA & IMAGE INTO MESSAGE 1708 TRANSMIT TO BASE TERMINAL 1710 SET CNTR=0 1712 1718 CNTR > N ? 1716 INCR CNTR REPLY TIME > T_L ? 1714 INTERPRET REPLY MESSAGE 1722 1724 VERIFIED ? 1726 DISPLAY REJECTION 1720 DISPLAY INABILITY TO READ 1728 DISPLAY APPROVAL END FIG.17

Element	'004 Claim Recitation	Exemplary Citations in Reference
		To the extent it is found that Ehrhart does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Ehrhart does not anticipate this claim, Ehrhart renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Ehrhart with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1E]	presenting to a user, via an address, content information associated with the target object.	Ehrhart discloses presenting to a user, via an address, content information associated with the target object. *Id.* at [1Pre-D]. Ehrhart at 13:27-43. • "As embodied herein and depicted in FIG. 13, an example of image association in accordance with the present invention is disclosed. One or ordinary skill in the art will recognize that associated images 1300 can be disposed on paper, displayed electronically on display 60, or displayed electronically sing other electronic means, such as a computer monitor. In this example, the first image captured is color photograph 1302 which shows a damaged parcel. The second image captured is bar code 1304 affixed to the side of the damaged parcel. Processor 40 decodes bar code 1304 and associates decoded bar code data 1306 with color photograph 1302. In this example, the user elected to associate a third image, signature 1308. Thus, personnel viewing record 1300 may reasonably conclude that a damaged parcel was delivered to Company XYZ,

Element	**'004 Claim Recitation**	**Exemplary Citations in Reference**
		and that the person signing for the parcel delivery was someone named John W. Smith." To the extent it is found that Ehrhart does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Ehrhart does not anticipate this claim, Ehrhart renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Ehrhart with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

2. Claim 6

Claim	'004 Claim Recitation	Exemplary Citations in Reference
6	The method of claim 1, wherein the at least one object is identified as the target object based at least in part on at least one user input.	Ehrhart discloses the method of claim 1, wherein the at least one object is identified as the target object based at least in part on at least one user input. *Id.* at [1]. Ehrhart at 10:40-50. • "In steps 520, 522, 526, and 532, the user is given the opportunity to select the type of graphical imaging that is to be performed. The method for performing OCR, text capture, and signature capture and/or verification are discussed above in the automatic mode description with one difference. In the semi-automatic mode, the user is asked in step 538 if he desires to associate the processed image with a subsequent captured image. If so, process flow is directed back to step 508 and another image is captured and displayed. The association feature can be used several times to associate multiple images." To the extent it is found that Ehrhart does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Ehrhart does not anticipate this claim, Ehrhart renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Ehrhart with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

3. Claim 18

Claim	'004 Claim Recitation	Exemplary Citations in Reference
18	The method of claim 1, further comprising communicating at least one of banking account information, user data and terminal specific data to a banking company.	Ehrhart discloses the method of claim 1, further comprising communicating at least one of banking account information, user data and terminal specific data to a banking company. *Id.* at [1]. Ehrhart at 15:1-21. • "As embodied herein and depicted in FIG. 17, a flow chart showing a method for performing signature verification is disclosed. In step 1700, optical reader 10 captures the image of the document to thereby generate a bit-map of the image. One of ordinary skill in the art will recognize that in the automatic mode or semi-automatic mode, processor 40 determines that the image object is a graphical symbol in a subsequent step. Step 1202 is similar to steps 1002 and 1004 of FIG. 10. The image is sampled by analyzing every Nth scan line of the bit mapped image. As discussed above, the image must be scanned in such a way so as to provide sufficient resolution to locate and classify the various regions on the document. In the case of a check, the location of the various fields on the instrument are relatively standard. Check sizes may differ somewhat, but the check number, bank code, account number, date, signature block, and etc. are in the same relative locations from check to check. In step 1704, document data such as the name, check number, bank code, account number, and date, are extracted from the document using any OCR program and stored in memory. In step 1706, the image of the hand writing in the signature block is captured." Ehrhart at 16:19-31. • "PAN 1850 includes at least one color optical reader 10 coupled to point-of-sale (POS) terminal 1854. POS terminal 1854 is coupled to network 1810 via interface 182. POS terminal 1854 includes a credit card reader and a

		signature capture block. In the scenario depicted in FIG. 18, a merchant user of POS terminal 1854 transmits an associated customer credit card number, signature, and in one embodiment, a color image of the customer, to Center 1830. Authentication module 1840 is used to authenticate the credit card and signature verification module is used to authenticate the signature. In another embodiment, database 1836 is used to store the customer's image, credit card number, and signature for verification purposes." To the extent it is found that Ehrhart does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Ehrhart does not anticipate this claim, Ehrhart renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Ehrhart with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

Nantworks, LLC v. Bank of America Corporation
Central District of California, No. 2:20-cv-7872
Invalidity Chart for U.S. Patent 9,324,004 based on
U.S. Patent No. 6,097,834 to Krouse ("Krouse")

Exhibit A-18

Bank of America Corporation ("BofA" or "Defendant") provides this chart subject to all reservations, objections, statements, and disclaimers set forth herein and in BofA's Final Invalidity Contentions Cover Pleading, as well as any amendment, supplement, or modification thereof, which are incorporated herein by reference in their entirety.

On information and belief, the application that matured into Krouse was filed on June 13, 1997, claiming the priority date of June 13, 1997. Krouse published at least as of August 1, 2000, and is available as prior art at least under 35 U.S.C. § 102(a).

As shown in the chart below, Krouse anticipates asserted claims 1, 2, 3, 6, 11, 15, 18 of the '004 Patent. To the extent Krouse is found not to expressly disclose certain limitations in the asserted claims, such limitations are inherent. To the extent Krouse is found not to anticipate any asserted claims or claim elements of the '004 Patent, the reference nevertheless renders those claims or claim elements obvious under 35 U.S.C. § 103, either alone or in combination with other art identified herein or in the cover pleading of BofA's Invalidity Contentions, for the reasons identified in the cover pleading. These Invalidity Contentions are not an admission by BofA that the accused products, including any current or past versions of the accused products, are covered by, or infringe the asserted claims, but are based instead on the recognition that if the claims are interpreted to be broad enough to encompass the accused products, the claims must also be construed to have that same scope when considering whether they are invalid.

The following chart is based on the Court's claim construction rulings (Dkt. Nos. 145, 153, 236), and to the extent not contradictory, Plaintiffs' apparent interpretations of the claims in Plaintiffs' "Preliminary" Final Infringement Contentions. BofA notes that in many instances, Plaintiffs' "Preliminary" Final Infringement Contentions ignore and/or contradict the Court's claim construction rulings. BofA does not accept the assumptions concerning the scope and meaning implicit in Plaintiffs' Preliminary Final Infringement Contentions, to the extent those assumptions are discernible, and reserves the right to challenge Plaintiffs' proposed (or implied) constructions. Furthermore, Plaintiffs' "Preliminary" Final Contentions fail to comply with the level of specificity required by SPR 2.1.2 and 2.1.3, and include improper reservation of rights to supplement with "additional ways in which BoA Accused Products infringe." Any alleged deficient discovery is of Plaintiffs own making. Most if not all of the deficiencies in Plaintiffs' Preliminary Final Contentions have been documented by Defendants through multiple letters since at least May 18, 2021 and up to Bank of America's most recent letter of August 19, 2023. Plaintiffs' shifting infringement contentions, deficiencies in its infringement contentions, and changing claim scope have prejudiced BofA from being able to reasonably prepare its defenses, including its invalidity contentions. BofA also reserves the right to revise and supplement these charts if and when Plaintiffs are permitted to provide revised or supplemental Infringement Contentions. Citations given in the chart below are merely representative of the respective elements to provide notice to Plaintiffs of BofA's invalidity theories and are not meant to be exhaustive.

1. Claim 1

Element	'004 Claim Recitation	Exemplary Citations in Reference
[1Pre]	A method for processing a video stream, comprising:	Krouse discloses a method for processing a video stream. Krouse at 1:4-12: • "The present invention relates generally to financial transaction processing Systems and methods, and more. Specifically, to Such Systems and methods wherein at least one document containing financial transaction-related information is optically Scanned to generate at least one computer-readable image from which the information may be extracted for use in processing the transaction. Krouse at 6:4-34: • "Each of the terminals 18, 20, 22 includes an optical Scanner 34 for Scanning negotiable instrument documents, (exemplified by check 110 shown in FIG. 6) tendered at point of Sale in payment for goods or Services provided by the merchant or retailer associated with that point of Sale, to generate respective computer-readable Scanned images of the documents. Preferably, Scanner 34 comprises a conventional monotonal image Scanner having a resolution of at least 300 dots-per-inch and being adapted for being controlled by local controller/processor 40 to scan one or both sides of a document 110 placed in the scanner 34. To the extent it is found that Krouse does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Krouse does not anticipate this claim, Krouse renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been

Element	**'004 Claim Recitation**	**Exemplary Citations in Reference**
		motivated to combine Krouse with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1A]	analyzing, via a mobile device, a scene represented by the video stream for at least one object;	Krouse discloses analyzing, via a mobile device, a scene represented by the video stream for at least one object. *Id.* at [1Pre]. Krouse at 1:4-12: • "The present invention relates generally to financial transaction processing Systems and methods, and more. Specifically, to Such Systems and methods wherein at least one document containing financial transaction-related information is optically Scanned to generate at least one computer-readable image from which the information may be extracted for use in processing the transaction. Krouse at 6:4-34: • "Each of the terminals 18, 20, 22 includes an optical Scanner 34 for Scanning negotiable instrument documents, (exemplified by check 110 shown in FIG. 6) tendered at point of Sale in payment for goods or Services provided by the merchant or retailer associated with that point of Sale, to generate respective computer-readable Scanned images of the documents. Preferably, Scanner 34 comprises a conventional monotonal image Scanner having a resolution of at least 300 dots-per-inch and being adapted for being controlled by local controller/processor 40 to scan one or both sides of a document 110 placed in the scanner 34.

Element	'004 Claim Recitation	Exemplary Citations in Reference
		To the extent it is found that Krouse does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Krouse does not anticipate this claim, Krouse renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Krouse with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1B]	deriving at least one characteristic of the video stream;	Krouse discloses deriving at least one characteristic of the video stream. Krouse at 1:4-12: • "The present invention relates generally to financial transaction processing Systems and methods, and more. Specifically, to Such Systems and methods wherein at least one document containing financial transaction-related information is optically Scanned to generate at least one computer-readable image from which the information may be extracted for use in processing the transaction. Krouse at 4:27-39: • "One preferred embodiment of the system of the second aspect of the present invention includes an optical Scanner for generating a Scanned image of at least a portion of the document containing the Visual representation. An image characterization generator is provided for generating recognition characteristics from the Scanned image, and a recognition characteristic comparitor is provided for comparing the recognition characteristics to respective Sets of reference recognition characteristics generated from respective other transaction documents having different respective formats and

Element	'004 Claim Recitation	Exemplary Citations in Reference
		for determining therefrom whether the particular format of the Visual data on the Scanned document matches one of the respective formats of the other documents. A field location generator is also provided for determining, based upon the one respective format when the particular format is determined to match the one respective format, location of a field in the Scanned image to which optical character recognition (OCR) may be applied to generate therefrom the transaction-related information represented by the Visual data. OCR is then applied by an image processor to this location to generate the Visual data for use in processing the transaction." Krouse at 4:40-57: • "One preferred embodiment of the method according to the Second aspect of the present invention includes generating an optically Scanned image of at least a portion of document containing visual data, in a particular format, representing information related to the financial transaction. Recognition characteristics are generated from the Scanned image and are compared to respective Sets of reference recognition characteristics generated from respective other transaction documents having different respective formats to determine therefrom whether the particular format of the visual data matches one of the respective formats of the other documents. When Such a match is found to exist, location is determined of a field in the Scanned image to which optical character recognition may be applied to generate therefrom the information, based upon the respective format found to match the particular format of the Visual data. Optical character recognition is then utilized to generate Said Visual data from Said location." Krouse at 6:4-34: • "Each of the terminals 18, 20, 22 includes an optical Scanner 34 for Scanning negotiable instrument documents, (exemplified by check 110 shown in FIG. 6) tendered at point of Sale in payment for goods or Services provided by the merchant or retailer associated with that point of Sale, to generate respective

Element	'004 Claim Recitation	Exemplary Citations in Reference
		computer-readable Scanned images of the documents. Preferably, Scanner 34 comprises a conventional monotonal image Scanner having a resolution of at least 300 dots-per-inch and being adapted for being controlled by local controller/processor 40 to scan one or both sides of a document 110 placed in the scanner 34. To the extent it is found that Krouse does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Krouse does not anticipate this claim, Krouse renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Krouse with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1C]	recognizing the at least one object in the scene as a target object based at least in part on the at least one characteristic of the video stream;	Krouse discloses recognizing the at least one object in the scene as a target object based at least in part on the at least one characteristic of the video stream. *Id.* at [1Pre-B]. Krouse at 1:4-12: • "The present invention relates generally to financial transaction processing Systems and methods, and more. Specifically, to Such Systems and methods wherein at least one document containing financial transaction-related information is optically Scanned to generate at least one computer-readable image from which the information may be extracted for use in processing the transaction.

Element	'004 Claim Recitation	Exemplary Citations in Reference
		Krouse at 4:27-39: • "One preferred embodiment of the system of the second aspect of the present invention includes an optical Scanner for generating a Scanned image of at least a portion of the document containing the Visual representation. An image characterization generator is provided for generating recognition characteristics from the Scanned image, and a recognition characteristic comparitor is provided for comparing the recognition characteristics to respective Sets of reference recognition characteristics generated from respective other transaction documents having different respective formats and for determining therefrom whether the particular format of the Visual data on the Scanned document matches one of the respective formats of the other documents. A field location generator is also provided for determining, based upon the one respective format when the particular format is determined to match the one respective format, location of a field in the Scanned image to which optical character recognition (OCR) may be applied to generate therefrom the transaction-related information represented by the Visual data. OCR is then applied by an image processor to this location to generate the Visual data for use in processing the transaction." Krouse at 4:40-57: • "One preferred embodiment of the method according to the Second aspect of the present invention includes generating an optically Scanned image of at least a portion of document containing visual data, in a particular format, representing information related to the financial transaction. Recognition characteristics are generated from the Scanned image and are compared to respective Sets of reference recognition characteristics generated from respective other transaction documents having different respective formats to determine therefrom whether the particular format of the visual data matches one of the respective formats of the other documents. When Such a match is found to exist, location is determined of a field in the Scanned image to which optical character recognition may be applied to generate therefrom the

Element	'004 Claim Recitation	Exemplary Citations in Reference
		information, based upon the respective format found to match the particular format of the Visual data. Optical character recognition is then utilized to generate Said Visual data from Said location." Krouse at 6:4-34: "Each of the terminals 18, 20, 22 includes an optical Scanner 34 for Scanning negotiable instrument documents, (exemplified by check 110 shown in FIG. 6) tendered at point of Sale in payment for goods or Services provided by the merchant or retailer associated with that point of Sale, to generate respective computer-readable Scanned images of the documents. Preferably, Scanner 34 comprises a conventional monotonal image Scanner having a resolution of at least 300 dots-per-inch and being adapted for being controlled by local controller/processor 40 to scan one or both sides of a document 110 placed in the scanner 34. To the extent it is found that Krouse does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Krouse does not anticipate this claim, Krouse renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Krouse with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1D]	providing an indication upon recognizing the at least one object as the target object; and	Krouse discloses providing an indication upon recognizing the at least one object as the target object. Krouse at 3:43–4:12

Element	'004 Claim Recitation	Exemplary Citations in Reference
		• "One preferred embodiment of the system of this aspect of the present invention includes an optical scanner for generating a computer-readable scanned image of the document, and an optical character recognizer for generating, from the scanned image, transaction data indicative of at least a portion of the visual information. This embodiment of the system of this aspect of the present invention also includes an authorization record generator for generating, based upon the transaction data, a record of the transaction for being assented to by the party tendering the document prior to completion of the transaction, and for generating from the record once it is assented to by the party, a computer-readable record of assent of the party to the transaction. A collection system is also included in this embodiment of the system of this aspect of the present invention for generating in response to the transaction data and the computer-readable record, an electronic financial transaction request for being used by an electronic transaction system to execute the electronic financial transaction. One embodiment of the method of the first aspect of the present invention includes generating a computer-readable scanned image of the document, and generating, from the scanned image, transaction data indicative of at least a portion of the visual information contained on the document. A record is generated, from the transaction data, of the electronic financial transaction being processed, which record is for being assented to by the party tendering the document, prior to completion of said transaction. Once the record has been assented to by the party tendering the document, a computer-readable record of the party's assent to the transaction is generated from the assented-to record. An electronic financial transaction request is then generated in response to the transaction data and the computer-readable assented-to record, which request is for being used by an electronic financial transaction system to execute the electronic financial transaction."

Element	'004 Claim Recitation	Exemplary Citations in Reference
		Krouse at 7:7–31 - "Once the aforesaid data (i.e., payment amount being tendered and payee) has been input via the device **30**, document **110** has been placed in the scanner **34**, and presence of the document **110** in the scanner **34** has been signalled to the terminal operator, the terminal operator continues processing the document **110** by commanding the processor **40** via the device **30** to initiate scanning of the document **110**. In response, the processor **40** issues commands to the scanner **34** to scan the document **110** placed therein, whereupon, the scanner **34** scans at least one entire side **112** of the draft **110** so as to generate a computer-readable scanned image of at least that entire side **112** of the draft document **110** (i.e., computer-readable data from which an appropriately provisioned computer may generate a display image of the at least one scanned side **112** of the document **110**). The computer-readable scanned image data is then transmitted by the scanner **34** to the processor **40**; the processor **40** then transmits to the optical character recognizer **32** a copy of the scanned image data generated by the scanner **34** in order to permit the recognizer **32** to process same in the manner described below. A separate copy of the scanned image data is also transmitted to the storage system **42** where it is stored in association with the previously stored payment amount and payee data of the negotiable instrument being processed." Krouse at 7:38–67 - "One such standardized portion in check **110** is Magnetic Ink Character Recognizable (MICR) line **94**, which is positioned at the bottom of the check **110** and in check **110**, includes three fields **114, 116, 118** of printed digits separated by special control characters. Fields **114, 116** contain the financial institution transaction routing number and account number on which the check **110** is drawn, and field **118** contains the check number of the check **110**. It should be noted, however, that other

Element	'004 Claim Recitation	Exemplary Citations in Reference
		than the left most nine digits (and associated separator characters) of the MICR line **94**, the contents of said line may vary, depending upon the financial institution upon which the negotiable instrument is drawn. System **10** may be adapted to recognize and parse information contained in different types of negotiable instrument MICR lines (not shown). Based upon the standardized location, size, style, etc. of the MICR line **94** in document **110** and the fields **114, 116, 118** contained therein, recognizer **32** utilizes optical character recognition and image processing techniques to extract from the scanned image data the data generated by the scanner **34** from scanning the MICR line portion **94** of the document **110**, and once thus extracted, to generate numeric transaction data indicative of the visual information (i.e., the numbers printed in fields **114, 116, 118**) in MICR line **94**. This numeric data is then transmitted to the processor **40**, which transmits a copy of the numeric data to the storage system **42** where it is stored in association with the scanned image data and the payee and payment amount data previously stored therein." Krouse at 16:66-17:19: • "The generator 224 then causes display 206 to display this information and to prompt via the display 206 for the user/customer at the display 206 to indicate via the user interface 206 whether the displayed information is correct. If the user/customer indicates that the displayed information is correct, the generator 224 then causes the display 206 to prompt the user/customer to provide means to pay the bill whose billing information is displayed on the display 206, otherwise the user/customer is prompted to input to the generator 224 the corrected billing information, which corrected information, once input, is displayed on the display with a request that the user verify correctness of Same. Once the user/customer indicates that the billing information displayed on the display 206 is

Element	'004 Claim Recitation	Exemplary Citations in Reference
		correct, the generator 224 causes the display 206 to request that the user/customer indicate to the system 200 via the user interface 206 how the user/customer wishes to make payment based upon the displayed billing information. Interface 206 may comprise a conventional point of Sale debit card reading/transaction mechanism (not shown) to permit the user/customer to accomplish this." To the extent it is found that Krouse does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Krouse does not anticipate this claim, Krouse renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Krouse with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1E]	presenting to a user, via an address, content information associated with the target object.	Krouse discloses presenting to a user, via an address, content information associated with the target object. Krouse at 7:7–31 • "Once the aforesaid data (i.e., payment amount being tendered and payee) has been input via the device **30**, document **110** has been placed in the scanner **34**, and presence of the document **110** in the scanner **34** has been signalled to the terminal operator, the terminal operator continues processing the document **110** by commanding the processor **40** via the device **30** to initiate scanning of the document **110**. In response, the processor **40** issues commands to the scanner **34** to scan the document **110** placed therein, whereupon, the scanner **34** scans at least one entire

Element	'004 Claim Recitation	Exemplary Citations in Reference
		side **112** of the draft **110** so as to generate a computer-readable scanned image of at least that entire side **112** of the draft document **110** (i.e., computer-readable data from which an appropriately provisioned computer may generate a display image of the at least one scanned side **112** of the document **110**). The computer-readable scanned image data is then transmitted by the scanner **34** to the processor **40**; the processor **40** then transmits to the optical character recognizer **32** a copy of the scanned image data generated by the scanner **34** in order to permit the recognizer **32** to process same in the manner described below. A separate copy of the scanned image data is also transmitted to the storage system **42** where it is stored in association with the previously stored payment amount and payee data of the negotiable instrument being processed." Krouse at 8:44–64 • "Each transaction record **70** also includes a portion **78** which specifies the date and time of processing of the transaction, a portion **76** specifying the terminal identification number of the terminal processing the transaction, and a transaction data section **82** specifying the numerical transaction data generated from the scanned image. Each transaction record **70** also contains a unique transaction identification number generated by the generator **38** by concatenating a 5-digit form of the terminal identification number to the time and date of the transaction in numerically expressed in the form YYYYMMDDHHMM (wherein YYYY represents the year, the left most MM represents the month, DD represents the day of the month, HH represents the hour in said day, and the right most MM represents the minutes in said hour of the transaction) and to a system-wide, unique 5-digit daily transaction sequence number (which is reset system-wide, periodically, e.g., once per day), and a tendering party assent portion **84** for receiving a signature **72** from the tendering party or other authorized representative thereof for signifying assent by the tendering party to the terms and conditions specified in the record **70**."

Element	'004 Claim Recitation	Exemplary Citations in Reference
		To the extent it is found that Krouse does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Krouse does not anticipate this claim, Krouse renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Krouse with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

2. Claim 6

Claim	'004 Claim Recitation	Exemplary Citations in Reference
6	The method of claim 1, wherein the at least one object is identified as the target object based at least in part on at least one user input.	Krouse discloses the method of claim 1, wherein the at least one object is identified as the target object based at least in part on at least one user input. *Id.* at [1]. Krouse at 3:43–4:12 • "One preferred embodiment of the system of this aspect of the present invention includes an optical scanner for generating a computer-readable scanned image of the document, and an optical character recognizer for generating, from the scanned image, transaction data indicative of at least a portion of the visual information. This embodiment of the system of this aspect of the present invention also includes an authorization record generator for generating, based upon

		the transaction data, a record of the transaction for being assented to by the party tendering the document prior to completion of the transaction, and for generating from the record once it is assented to by the party, a computer-readable record of assent of the party to the transaction. A collection system is also included in this embodiment of the system of this aspect of the present invention for generating in response to the transaction data and the computer-readable record, an electronic financial transaction request for being used by an electronic transaction system to execute the electronic financial transaction. One embodiment of the method of the first aspect of the present invention includes generating a computer-readable scanned image of the document, and generating, from the scanned image, transaction data indicative of at least a portion of the visual information contained on the document. A record is generated, from the transaction data, of the electronic financial transaction being processed, which record is for being assented to by the party tendering the document, prior to completion of said transaction. Once the record has been assented to by the party tendering the document, a computer-readable record of the party's assent to the transaction is generated from the assented-to record. An electronic financial transaction request is then generated in response to the transaction data and the computer-readable assented-to record, which request is for being used by an electronic financial transaction system to execute the electronic financial transaction." Krouse at 7:7–31 • "Once the aforesaid data (i.e., payment amount being tendered and payee) has been input via the device **30**, document **110** has been placed in the scanner **34**, and presence of the document **110** in the scanner **34** has been signalled to the terminal operator, the terminal operator continues processing the document **110** by commanding the processor **40** via the device **30** to initiate scanning of the document **110**. In

		response, the processor **40** issues commands to the scanner **34** to scan the document **110** placed therein, whereupon, the scanner **34** scans at least one entire side **112** of the draft **110** so as to generate a computer-readable scanned image of at least that entire side **112** of the draft document **110** (i.e., computer-readable data from which an appropriately provisioned computer may generate a display image of the at least one scanned side **112** of the document **110**). The computer-readable scanned image data is then transmitted by the scanner **34** to the processor **40**; the processor **40** then transmits to the optical character recognizer **32** a copy of the scanned image data generated by the scanner **34** in order to permit the recognizer **32** to process same in the manner described below. A separate copy of the scanned image data is also transmitted to the storage system **42** where it is stored in association with the previously stored payment amount and payee data of the negotiable instrument being processed." Krouse at 7:38–67 • "One such standardized portion in check **110** is Magnetic Ink Character Recognizable (MICR) line **94**, which is positioned at the bottom of the check **110** and in check **110**, includes three fields **114, 116, 118** of printed digits separated by special control characters. Fields **114, 116** contain the financial institution transaction routing number and account number on which the check **110** is drawn, and field **118** contains the check number of the check **110**. It should be noted, however, that other than the left most nine digits (and associated separator characters) of the MICR line **94**, the contents of said line may vary, depending upon the financial institution upon which the negotiable instrument is drawn. System **10** may be adapted to recognize and parse information contained in different types of negotiable instrument MICR lines (not shown). Based upon the standardized location, size, style, etc. of the MICR line

		94 in document **110** and the fields **114, 116, 118** contained therein, recognizer **32** utilizes optical character recognition and image processing techniques to extract from the scanned image data the data generated by the scanner **34** from scanning the MICR line portion **94** of the document **110**, and once thus extracted, to generate numeric transaction data indicative of the visual information (i.e., the numbers printed in fields **114, 116, 118**) in MICR line **94**. This numeric data is then transmitted to the processor **40**, which transmits a copy of the numeric data to the storage system **42** where it is stored in association with the scanned image data and the payee and payment amount data previously stored therein." To the extent it is found that Krouse does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Krouse does not anticipate this claim, Krouse renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Krouse with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

3. Claim 18

Claim	'004 Claim Recitation	Exemplary Citations in Reference

18	The method of claim 1, further comprising communicating at least one of banking account information, user data and terminal specific data to a banking company.	Krouse discloses the method of claim 1, further comprising communicating at least one of banking account information, user data and terminal specific data to a banking company. *Id.* at [1]. Krouse at Fig. 6 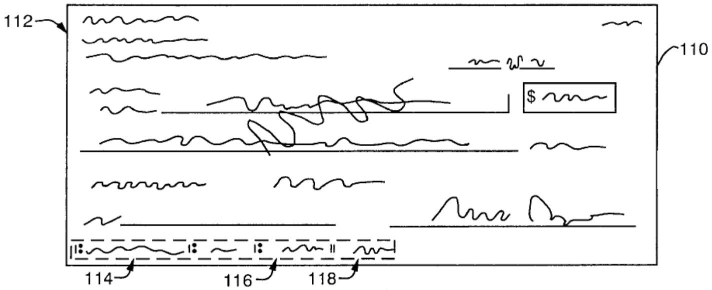 FIG. 6 Krouse at 7:43–53 • “Fields **114, 116** contain the financial institution transaction routing number and account number on which the check **110** is drawn, and field **118** contains the check number of the check **110**. It should be noted, however, that other than the left most nine digits (and associated separator characters) of the MICR line **94**, the contents of said line may vary, depending upon the financial institution upon which the negotiable instrument is drawn. System **10** may be adapted to recognize and parse information contained in different types of negotiable instrument MICR lines (not shown).” Krouse at 7:38–67 • “One such standardized portion in check **110** is Magnetic Ink Character Recognizable (MICR) line **94**, which is positioned at the bottom of the

		check **110** and in check **110**, includes three fields **114, 116, 118** of printed digits separated by special control characters. Fields **114, 116** contain the financial institution transaction routing number and account number on which the check **110** is drawn, and field **118** contains the check number of the check **110**. It should be noted, however, that other than the left most nine digits (and associated separator characters) of the MICR line **94**, the contents of said line may vary, depending upon the financial institution upon which the negotiable instrument is drawn. System **10** may be adapted to recognize and parse information contained in different types of negotiable instrument MICR lines (not shown). Based upon the standardized location, size, style, etc. of the MICR line **94** in document **110** and the fields **114, 116, 118** contained therein, recognizer **32** utilizes optical character recognition and image processing techniques to extract from the scanned image data the data generated by the scanner **34** from scanning the MICR line portion **94** of the document **110**, and once thus extracted, to generate numeric transaction data indicative of the visual information (i.e., the numbers printed in fields **114, 116, 118**) in MICR line **94**. This numeric data is then transmitted to the processor **40**, which transmits a copy of the numeric data to the storage system **42** where it is stored in association with the scanned image data and the payee and payment amount data previously stored therein." To the extent it is found that Krouse does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Krouse does not anticipate this claim, Krouse renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Krouse with any of the other prior art references identified

		in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

Nantworks, LLC v. Bank of America Corporation
Central District of California, No. 2:20-cv-7872
Invalidity Chart for U.S. Patent 9,324,004 based on the Query by Image and Video Content ("QBIC") System

Exhibit A-23

Bank of America Corporation ("BofA" or "Defendant") provides this chart subject to all reservations, objections, statements, and disclaimers set forth herein and in BofA's Final Invalidity Contentions Cover Pleading, as well as any amendment, supplement, or modification thereof, which are incorporated herein by reference in their entirety. A statement of reasons and good cause for each supplement in this chart is identified in the cover pleading.

QBIC was known or used by others in the United States and publicly used or on sale in the United States by November 5, 2000. *See generally* W. Niblack et al., The QBIC Project: Querying Images By Content Using Color, Texture, and Shape, SPIE Vol. 1908 (1993), pp. 173–187 (IBM0002390–404); M. Flickner et al., Query by Image and Video Content: The QBIC System ("QBIC"), IEEE Sept. 1995, pp. 23–32 (IBM0000893–902); W. Niblack et al, Updates to the QBIC System, SPIE Vol. 3312 (1997), pp. 150–161 (IBM 0002419–430); and Query By Image Content QBIC Demonstration Program, IBM Research, Sept. 1998 (IBM000747). QBIC is available as prior art at least under 35 U.S.C. § 102(a), (b), and (g), having been known or used by others in the United States and publicly used or on sale in the United States by at least 1995 and certainly by November 5, 2000, including, for example, at http://libra.uc.davis.edu through the Art History Department at U.C. Davis; at www..thinker.org/imagebase/indox-2 .htm3 through the Fine Arts Museum of San Francisco; through demos at http://wwwqbic.almaden.ibm.com; and through IBM's commercial products, including in IBM's Ultimedia Manager, the DB2 series of products such as the Image Extender in DB2 Universal Database, and the Enterprise Information Portal. *See generally* Deposition of IBM Corporate Representative Myron Flickner, *Nantworks, LLC and Nant Holdings IP, LLC v. Bank of America Corporation and Bank of America, N.A.*, 2:20-cv-07872-GW-PVC (C.D.C.A. 2020) ("Flickner 8/29/23 Depo") at e.g., 32–33, 39–40, 44, 62–63, 66–67,72–74, 103–107, 110–112, 127, 163–174, 183–189, 191–195, 215, 219–220, 241, 307–309, 311–317, Exs. 3–13; *see also* Deposition of IBM Corporate Representative Myron Flickner, *Network-1 Technologies, Inc., v. Google LLC and Youtube LLC*, 1:14-cv-09558-PGG (S.D.N.Y. 2019) ("Flickner Depo"), including, e.g., at 13, 14, 15–18, 29–30, 34, 36, 37, 43–46, 48, 58, 60, 61–62, 67–69, 72, 74–76, 84–86, 88–89, 98–100, 102, 106, 113, 125, 135, 152–154, 160–161, 170–172, 180–181, 194, 198–201, 203–205, 215–216, 220–221, and Exs. 1–16.

As shown in the chart below, QBIC anticipates asserted claims 1, 6, and 18 of the '004 Patent. The disclosures/exemplary citations for each element incorporate the disclosures citations of the proceeding limitations. To the extent QBIC is found not to expressly disclose certain limitations in the asserted claims, such limitations are inherent. To the extent QBIC is found not to anticipate any asserted claims or claim elements of the '004 Patent, the reference nevertheless renders those claims or claim elements obvious under 35 U.S.C. § 103, either alone or in combination with other art identified herein or in the cover pleading of BofA's Final Invalidity Contentions, for the reasons identified in the cited references and/or in the cover pleading. These Invalidity Contentions are not an admission by BofA that the accused products, including any current or past versions of the accused products, are covered by, or infringe the asserted claims, but are based instead on the recognition that if the claims are interpreted to be broad enough to encompass the accused products, the claims must also be construed to have that same scope when considering whether they are invalid. The following chart incorporates by reference the analysis from Exhibit A-21 for each limitation.

The following chart is based on the Court's claim construction rulings (Dkt. Nos. 145, 153, 236), and to the extent not contradictory, Plaintiffs' apparent interpretations of the claims in Plaintiffs' "Preliminary" Final Infringement Contentions. BofA notes that in many instances, Plaintiffs' "Preliminary" Final Infringement Contentions ignore and/or contradict the Court's claim construction rulings. BofA does not accept the assumptions concerning the scope and meaning implicit in Plaintiffs' Preliminary Final Infringement Contentions, to the extent those assumptions are discernible, and reserves the right to challenge Plaintiffs' proposed (or implied) constructions. Furthermore, Plaintiffs' "Preliminary" Final Contentions fail to comply with the level of specificity required by SPR 2.1.2 and 2.1.3, and include improper reservation of rights to supplement with "additional ways in which BoA Accused Products infringe." Any alleged deficient discovery is of Plaintiffs own making. Most if not all of the deficiencies in Plaintiffs' Preliminary Final Contentions have been documented by Defendants through multiple letters since at least May 18, 2021 and up to Bank of America's most recent letter of August 19, 2023. Plaintiffs' shifting infringement contentions, deficiencies in its infringement contentions, and changing claim scope have prejudiced BofA from being able to reasonably prepare its defenses, including its invalidity contentions. BofA also reserves the right to revise and supplement these charts if and when Plaintiffs are permitted to provide revised or supplemental Infringement Contentions. Citations given in the chart below are merely representative of the respective elements to provide notice to Plaintiffs of BofA's invalidity theories and are not meant to be exhaustive.

1. Claim 1

Element	’004 Claim Recitation	Exemplary Citations in Reference
[1Pre]	A method for processing a video stream, comprising:	The QBIC Project: Querying Images by Content Using Color, Texture, and Shape (IBM 0002390–404) at 2390 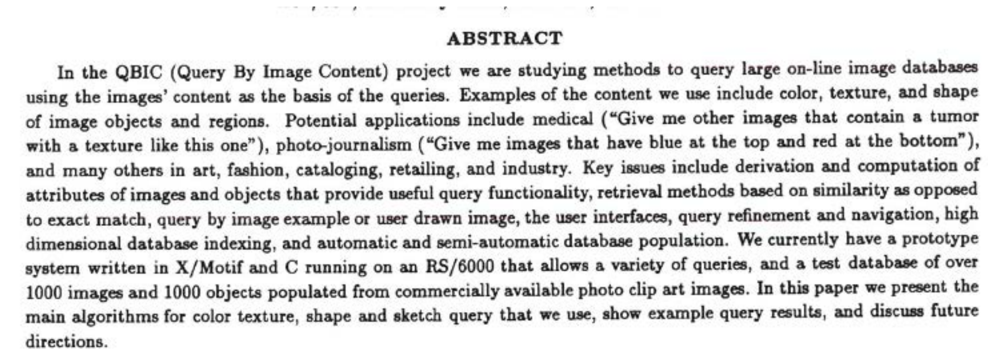 ABSTRACT In the QBIC (Query By Image Content) project we are studying methods to query large on-line image databases using the images’ content as the basis of the queries. Examples of the content we use include color, texture, and shape of image objects and regions. Potential applications include medical (“Give me other images that contain a tumor with a texture like this one”), photo-journalism (“Give me images that have blue at the top and red at the bottom”), and many others in art, fashion, cataloging, retailing, and industry. Key issues include derivation and computation of attributes of images and objects that provide useful query functionality, retrieval methods based on similarity as opposed to exact match, query by image example or user drawn image, the user interfaces, query refinement and navigation, high dimensional database indexing, and automatic and semi-automatic database population. We currently have a prototype system written in X/Motif and C running on an RS/6000 that allows a variety of queries, and a test database of over 1000 images and 1000 objects populated from commercially available photo clip art images. In this paper we present the main algorithms for color texture, shape and sketch query that we use, show example query results, and discuss future directions. The QBIC Project: Querying Images by Content Using Color, Texture, and Shape (IBM 0002390–404) at 3 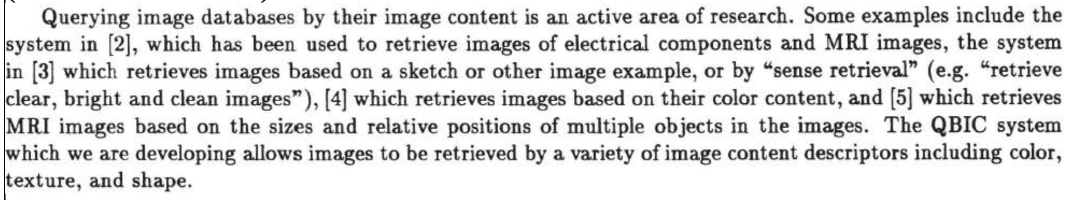 Querying image databases by their image content is an active area of research. Some examples include the system in [2], which has been used to retrieve images of electrical components and MRI images, the system in [3] which retrieves images based on a sketch or other image example, or by “sense retrieval” (e.g. “retrieve clear, bright and clean images”), [4] which retrieves images based on their color content, and [5] which retrieves MRI images based on the sizes and relative positions of multiple objects in the images. The QBIC system which we are developing allows images to be retrieved by a variety of image content descriptors including color, texture, and shape. Query by Image and Video Content: The QBIC System (IBM 0000893–902): “We have developed the QBIC (Query by Image Content) system to explore content-based retrieval methods. QBIC allows queries on large image and video databases based on • example images, • user-constructed sketches and drawings, • selected color and texture patterns, • camera and object motion, and • other graphical information.” *Id.* at pp. 1–3.

Element	'004 Claim Recitation	Exemplary Citations in Reference
		"To achieve this functionality, QBIC has two main components: database population (the process of creating an image database) and database query. During the population, images and videos are processed to extract features describing their content-colors, textures, shapes, and camera and object motion-and the features are stored in a database. During the query, the user composes a query graphically. Features are generated from the graphical query and then input to a matching engine that finds images or videos from the database with similar features. Figure 2 shows the system architecture." *Id*. at p. 3.

Element	'004 Claim Recitation	Exemplary Citations in Reference
		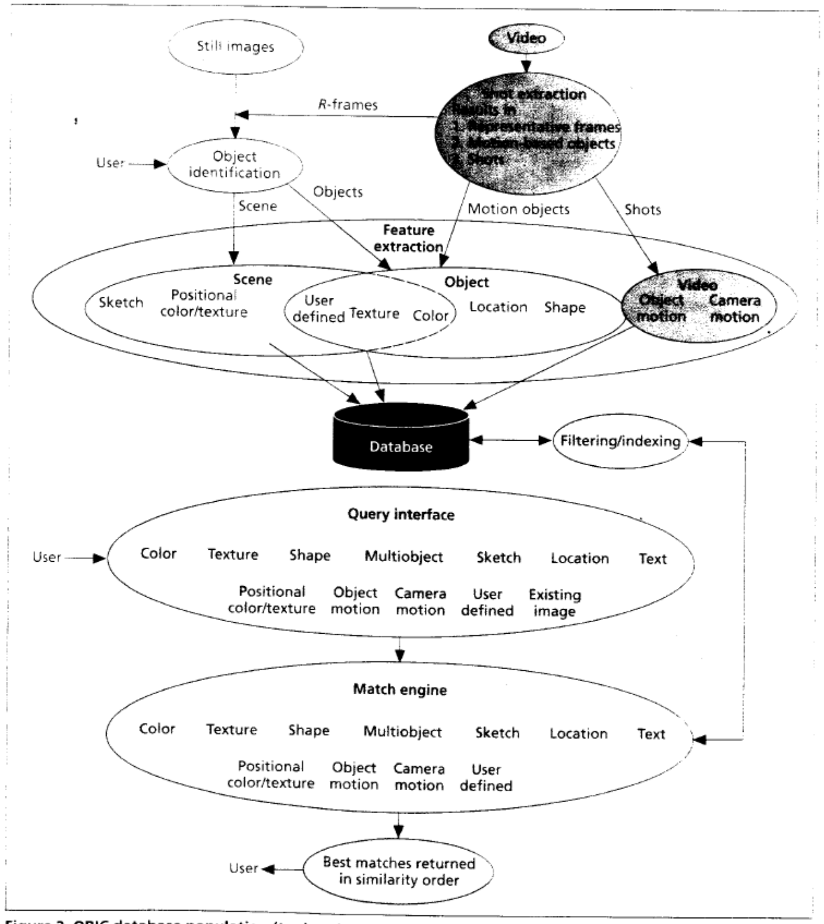 Figure 2. QBIC database population (top) and query (bottom) architecture. *Id*. at Fig. 2.

Element	'004 Claim Recitation	Exemplary Citations in Reference
		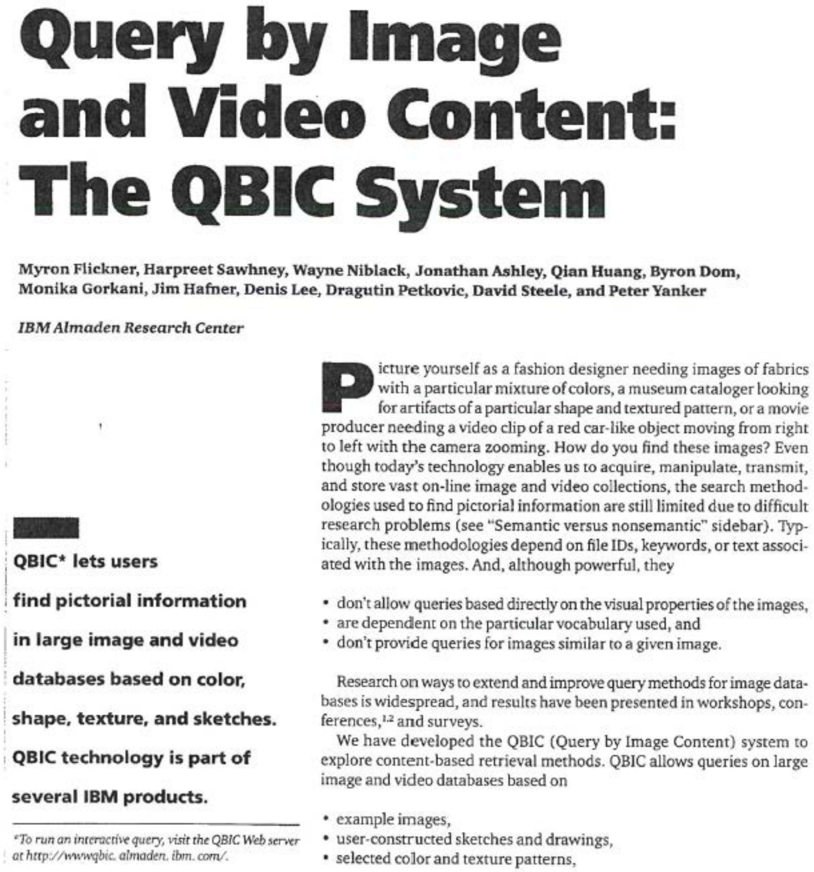 **Query by Image and Video Content: The QBIC System** Myron Flickner, Harpreet Sawhney, Wayne Niblack, Jonathan Ashley, Qian Huang, Byron Dom, Monika Gorkani, Jim Hafner, Denis Lee, Dragutin Petkovic, David Steele, and Peter Yanker *IBM Almaden Research Center* **QBIC* lets users find pictorial information in large image and video databases based on color, shape, texture, and sketches. QBIC technology is part of several IBM products.** Picture yourself as a fashion designer needing images of fabrics with a particular mixture of colors, a museum cataloger looking for artifacts of a particular shape and textured pattern, or a movie producer needing a video clip of a red car-like object moving from right to left with the camera zooming. How do you find these images? Even though today's technology enables us to acquire, manipulate, transmit, and store vast on-line image and video collections, the search methodologies used to find pictorial information are still limited due to difficult research problems (see "Semantic versus nonsemantic" sidebar). Typically, these methodologies depend on file IDs, keywords, or text associated with the images. And, although powerful, they • don't allow queries based directly on the visual properties of the images, • are dependent on the particular vocabulary used, and • don't provide queries for images similar to a given image. Research on ways to extend and improve query methods for image databases is widespread, and results have been presented in workshops, conferences,[1,2] and surveys. We have developed the QBIC (Query by Image Content) system to explore content-based retrieval methods. QBIC allows queries on large image and video databases based on • example images, • user-constructed sketches and drawings, • selected color and texture patterns, **To run an interactive query, visit the QBIC Web server at http://wwwqbic.almaden.ibm.com/.* IBM0000893 at 893

Element	'004 Claim Recitation	Exemplary Citations in Reference
		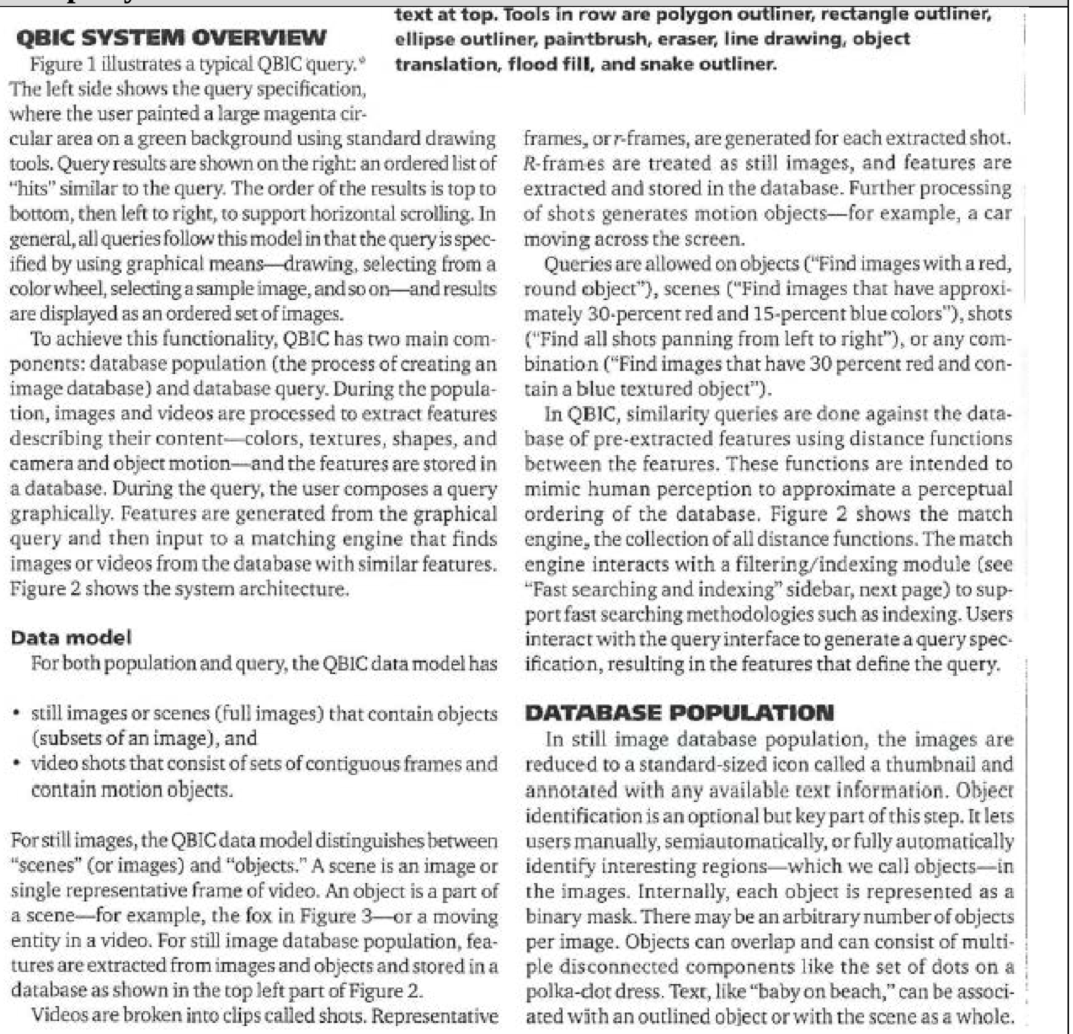 **text at top. Tools in row are polygon outliner, rectangle outliner, ellipse outliner, paintbrush, eraser, line drawing, object translation, flood fill, and snake outliner.** **QBIC SYSTEM OVERVIEW** Figure 1 illustrates a typical QBIC query.* The left side shows the query specification, where the user painted a large magenta circular area on a green background using standard drawing tools. Query results are shown on the right: an ordered list of "hits" similar to the query. The order of the results is top to bottom, then left to right, to support horizontal scrolling. In general, all queries follow this model in that the query is specified by using graphical means—drawing, selecting from a color wheel, selecting a sample image, and so on—and results are displayed as an ordered set of images. To achieve this functionality, QBIC has two main components: database population (the process of creating an image database) and database query. During the population, images and videos are processed to extract features describing their content—colors, textures, shapes, and camera and object motion—and the features are stored in a database. During the query, the user composes a query graphically. Features are generated from the graphical query and then input to a matching engine that finds images or videos from the database with similar features. Figure 2 shows the system architecture. **Data model** For both population and query, the QBIC data model has • still images or scenes (full images) that contain objects (subsets of an image), and • video shots that consist of sets of contiguous frames and contain motion objects. For still images, the QBIC data model distinguishes between "scenes" (or images) and "objects." A scene is an image or single representative frame of video. An object is a part of a scene—for example, the fox in Figure 3—or a moving entity in a video. For still image database population, features are extracted from images and objects and stored in a database as shown in the top left part of Figure 2. Videos are broken into clips called shots. Representative frames, or *r*-frames, are generated for each extracted shot. *R*-frames are treated as still images, and features are extracted and stored in the database. Further processing of shots generates motion objects—for example, a car moving across the screen. Queries are allowed on objects ("Find images with a red, round object"), scenes ("Find images that have approximately 30-percent red and 15-percent blue colors"), shots ("Find all shots panning from left to right"), or any combination ("Find images that have 30 percent red and contain a blue textured object"). In QBIC, similarity queries are done against the database of pre-extracted features using distance functions between the features. These functions are intended to mimic human perception to approximate a perceptual ordering of the database. Figure 2 shows the match engine, the collection of all distance functions. The match engine interacts with a filtering/indexing module (see "Fast searching and indexing" sidebar, next page) to support fast searching methodologies such as indexing. Users interact with the query interface to generate a query specification, resulting in the features that define the query. **DATABASE POPULATION** In still image database population, the images are reduced to a standard-sized icon called a thumbnail and annotated with any available text information. Object identification is an optional but key part of this step. It lets users manually, semiautomatically, or fully automatically identify interesting regions—which we call objects—in the images. Internally, each object is represented as a binary mask. There may be an arbitrary number of objects per image. Objects can overlap and can consist of multiple disconnected components like the set of dots on a polka-dot dress. Text, like "baby on beach," can be associated with an outlined object or with the scene as a whole. IBM0000893 at 895

Element	'004 Claim Recitation	Exemplary Citations in Reference
		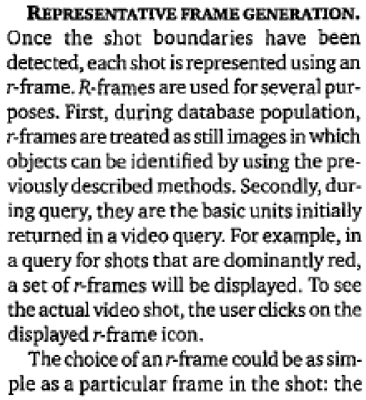 REPRESENTATIVE FRAME GENERATION. Once the shot boundaries have been detected, each shot is represented using an *r*-frame. *R*-frames are used for several purposes. First, during database population, *r*-frames are treated as still images in which objects can be identified by using the previously described methods. Secondly, during query, they are the basic units initially returned in a video query. For example, in a query for shots that are dominantly red, a set of *r*-frames will be displayed. To see the actual video shot, the user clicks on the displayed *r*-frame icon. The choice of an *r*-frame could be as simple as a particular frame in the shot: the IBM0000893 at 898 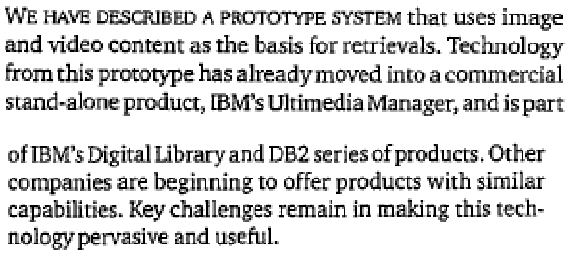 WE HAVE DESCRIBED A PROTOTYPE SYSTEM that uses image and video content as the basis for retrievals. Technology from this prototype has already moved into a commercial stand-alone product, IBM's Ultimedia Manager, and is part of IBM's Digital Library and DB2 series of products. Other companies are beginning to offer products with similar capabilities. Key challenges remain in making this technology pervasive and useful. IBM0000893 at 900 QBIC Demo (IBM000747)

Element	'004 Claim Recitation	Exemplary Citations in Reference
		 QBIC(TM) -- IBM's Query By Image Content Goto: How to Use QBIC IBM Image Search Demo Stop QBIC Http Server QBIC allows you to make queries of image databases large and small based on visual characteristics -- properties such as color percentages, color layout and textures occurring in the images. The queries match image colors, textures and their positions without having to describe them with in words. QBIC's content-based queries can also be combined with keywords to provide a powerful multimedia search and retrieval tool. As image collections grow larger and become more prolific, QBIC is the best tool to more efficiently search and navigate through them. In this demo , we use QBIC to index about 470 color images based on their color and texture information. You can evaluate the QBIC Search capability by following the link located at left called IBM Image Search Demo. For license and business related questins, please contact Ted Loewenberg by phone 408-927-1202, or by email Flickner Depo at 38-41

Element	'004 Claim Recitation	Exemplary Citations in Reference
		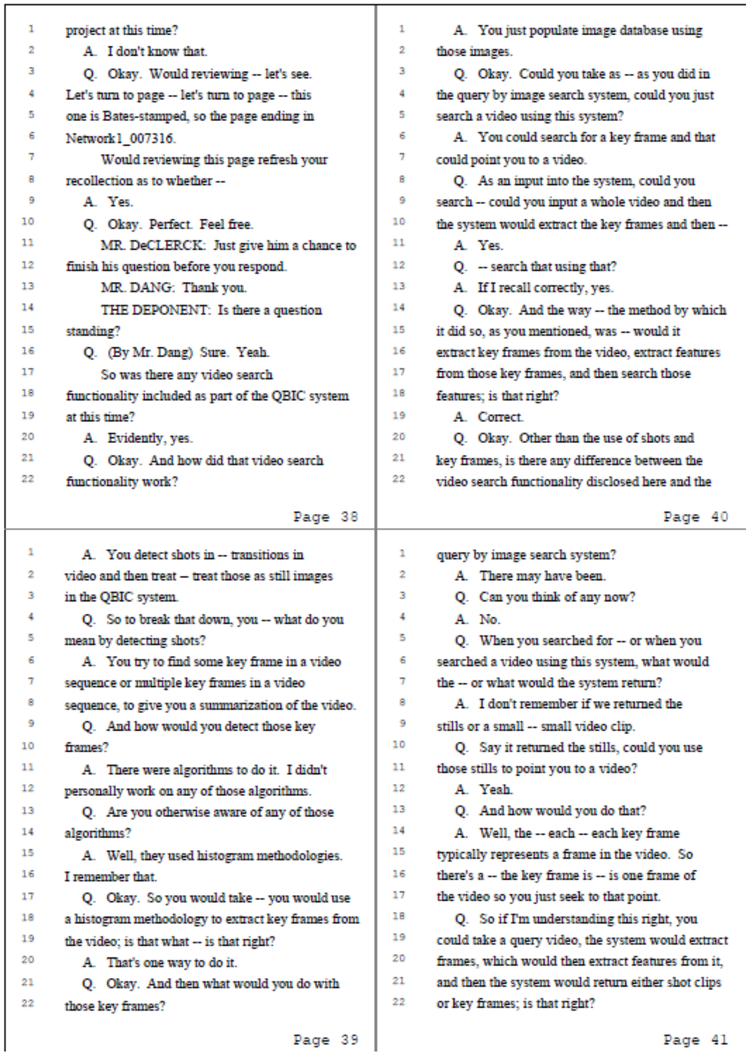 project at this time? A. I don't know that. Q. Okay. Would reviewing -- let's see. Let's turn to page -- let's turn to page -- this one is Bates-stamped, so the page ending in Network1_007316. Would reviewing this page refresh your recollection as to whether -- A. Yes. Q. Okay. Perfect. Feel free. MR. DeCLERCK: Just give him a chance to finish his question before you respond. MR. DANG: Thank you. THE DEPONENT: Is there a question standing? Q. (By Mr. Dang) Sure. Yeah. So was there any video search functionality included as part of the QBIC system at this time? A. Evidently, yes. Q. Okay. And how did that video search functionality work? Page 38 A. You detect shots in -- transitions in video and then treat -- treat those as still images in the QBIC system. Q. So to break that down, you -- what do you mean by detecting shots? A. You try to find some key frame in a video sequence or multiple key frames in a video sequence, to give you a summarization of the video. Q. And how would you detect those key frames? A. There were algorithms to do it. I didn't personally work on any of those algorithms. Q. Are you otherwise aware of any of those algorithms? A. Well, they used histogram methodologies. I remember that. Q. Okay. So you would take -- you would use a histogram methodology to extract key frames from the video; is that what -- is that right? A. That's one way to do it. Q. Okay. And then what would you do with those key frames? Page 39 A. You just populate image database using those images. Q. Okay. Could you take as -- as you did in the query by image search system, could you just search a video using this system? A. You could search for a key frame and that could point you to a video. Q. As an input into the system, could you search -- could you input a whole video and then the system would extract the key frames and then -- A. Yes. Q. -- search that using that? A. If I recall correctly, yes. Q. Okay. And the way -- the method by which it did so, as you mentioned, was -- would it extract key frames from the video, extract features from those key frames, and then search those features; is that right? A. Correct. Q. Okay. Other than the use of shots and key frames, is there any difference between the video search functionality disclosed here and the Page 40 query by image search system? A. There may have been. Q. Can you think of any now? A. No. Q. When you searched for -- or when you searched a video using this system, what would the -- or what would the system return? A. I don't remember if we returned the stills or a small -- small video clip. Q. Say it returned the stills, could you use those stills to point you to a video? A. Yeah. Q. And how would you do that? A. Well, the -- each -- each key frame typically represents a frame in the video. So there's a -- the key frame is -- is one frame of the video so you just seek to that point. Q. So if I'm understanding this right, you could take a query video, the system would extract frames, which would then extract features from it, and then the system would return either shot clips or key frames; is that right? Page 41 Flickner Depo at 209-210

Element	'004 Claim Recitation	Exemplary Citations in Reference
		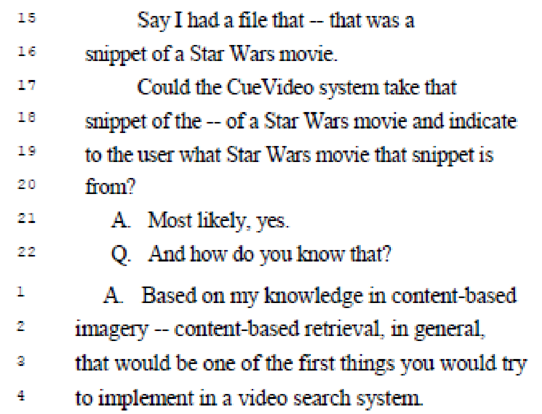 Q. Okay. Let me ask -- let me ask that question using -- using an example. Say I had a file that -- that was a snippet of a Star Wars movie. Could the CueVideo system take that snippet of the -- of a Star Wars movie and indicate to the user what Star Wars movie that snippet is from? A. Most likely, yes. Q. And how do you know that? A. Based on my knowledge in content-based imagery -- content-based retrieval, in general, that would be one of the first things you would try to implement in a video search system. Flickner Depo at 201 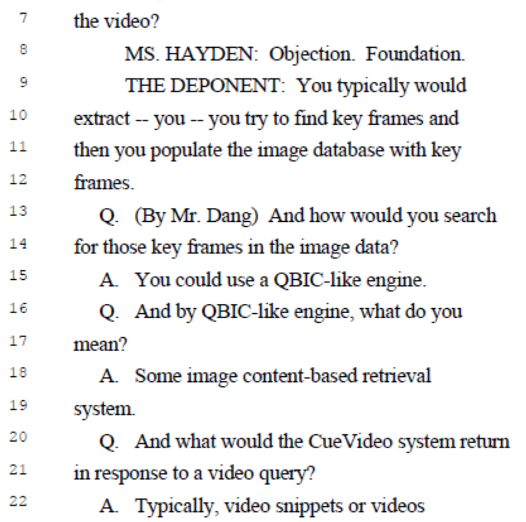 Q. Now, to search a video by content using CueVideo, how would the features be extracted from the video? MS. HAYDEN: Objection. Foundation. THE DEPONENT: You typically would extract -- you -- you try to find key frames and then you populate the image database with key frames. Q. (By Mr. Dang) And how would you search for those key frames in the image data? A. You could use a QBIC-like engine. Q. And by QBIC-like engine, what do you mean? A. Some image content-based retrieval system. Q. And what would the CueVideo system return in response to a video query? A. Typically, video snippets or videos Flickner Depo at 214-25

Element	'004 Claim Recitation	Exemplary Citations in Reference
		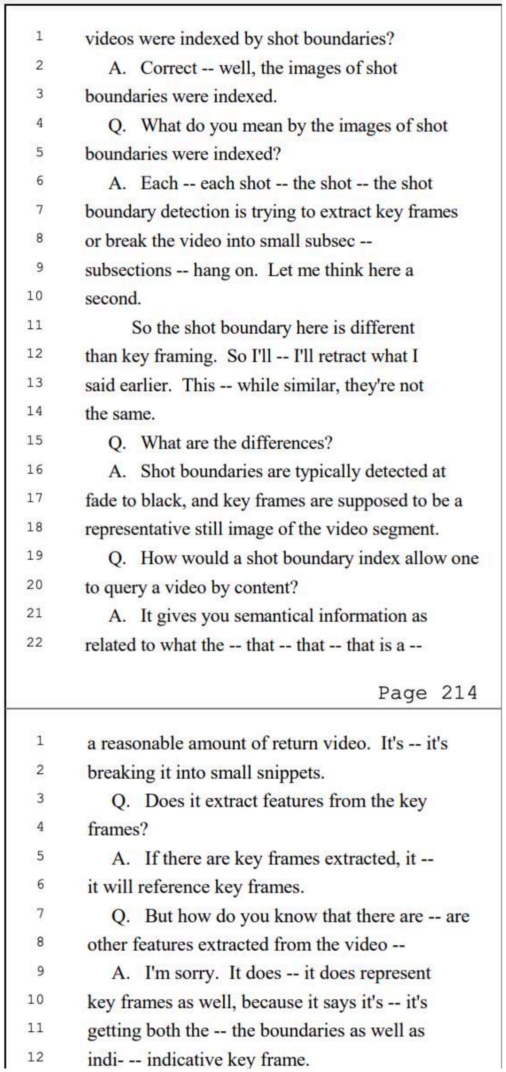 videos were indexed by shot boundaries? A. Correct -- well, the images of shot boundaries were indexed. Q. What do you mean by the images of shot boundaries were indexed? A. Each -- each shot -- the shot -- the shot boundary detection is trying to extract key frames or break the video into small subsec -- subsections -- hang on. Let me think here a second. So the shot boundary here is different than key framing. So I'll -- I'll retract what I said earlier. This -- while similar, they're not the same. Q. What are the differences? A. Shot boundaries are typically detected at fade to black, and key frames are supposed to be a representative still image of the video segment. Q. How would a shot boundary index allow one to query a video by content? A. It gives you semantical information as related to what the -- that -- that -- that is a -- Page 214 a reasonable amount of return video. It's -- it's breaking it into small snippets. Q. Does it extract features from the key frames? A. If there are key frames extracted, it -- it will reference key frames. Q. But how do you know that there are -- are other features extracted from the video -- A. I'm sorry. It does -- it does represent key frames as well, because it says it's -- it's getting both the -- the boundaries as well as indi- -- indicative key frame. Flickner 8/29/23 Depo at 32–33

Element	'004 Claim Recitation	Exemplary Citations in Reference
		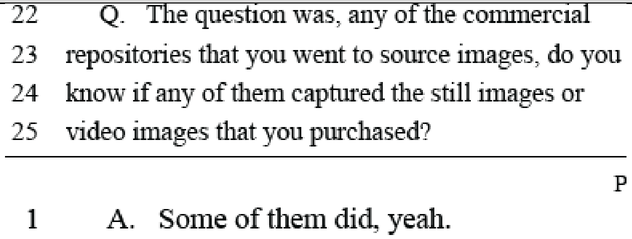 Q. The question was, any of the commercial repositories that you went to source images, do you know if any of them captured the still images or video images that you purchased? P A. Some of them did, yeah. *Id.* at 34 – 35

Element	’004 Claim Recitation	Exemplary Citations in Reference
		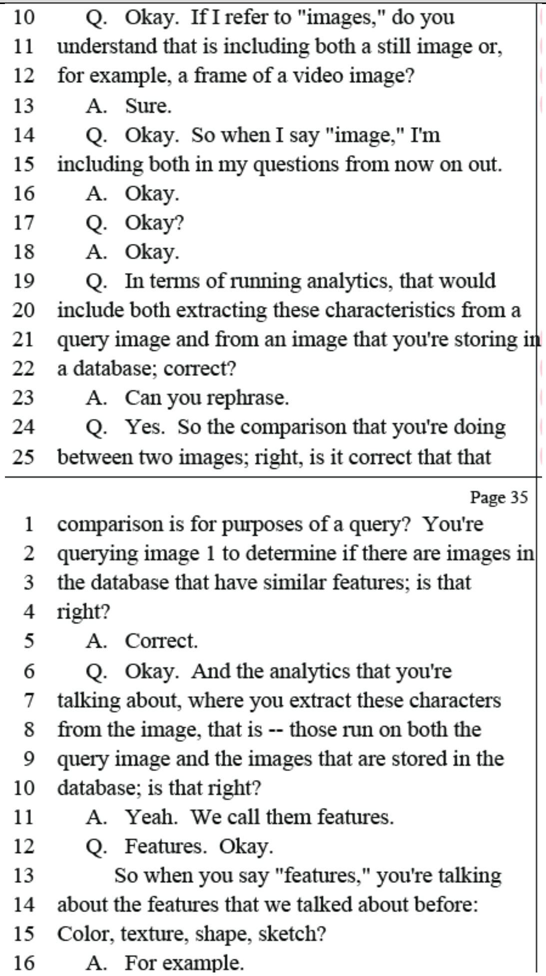 Q. Okay. If I refer to "images," do you understand that is including both a still image or, for example, a frame of a video image? A. Sure. Q. Okay. So when I say "image," I'm including both in my questions from now on out. A. Okay. Q. Okay? A. Okay. Q. In terms of running analytics, that would include both extracting these characteristics from a query image and from an image that you're storing in a database; correct? A. Can you rephrase. Q. Yes. So the comparison that you're doing between two images; right, is it correct that that Page 35 comparison is for purposes of a query? You're querying image 1 to determine if there are images in the database that have similar features; is that right? A. Correct. Q. Okay. And the analytics that you're talking about, where you extract these characters from the image, that is -- those run on both the query image and the images that are stored in the database; is that right? A. Yeah. We call them features. Q. Features. Okay. So when you say "features," you're talking about the features that we talked about before: Color, texture, shape, sketch? A. For example. *Id.* at 42

Element	'004 Claim Recitation	Exemplary Citations in Reference
		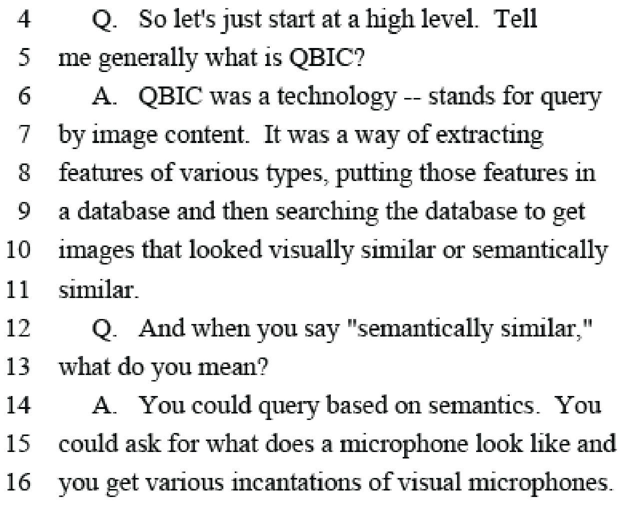 Q. So let's just start at a high level. Tell me generally what is QBIC? A. QBIC was a technology -- stands for query by image content. It was a way of extracting features of various types, putting those features in a database and then searching the database to get images that looked visually similar or semantically similar. Q. And when you say "semantically similar," what do you mean? A. You could query based on semantics. You could ask for what does a microphone look like and you get various incantations of visual microphones. *Id.* at 205–206

Element	'004 Claim Recitation	Exemplary Citations in Reference
		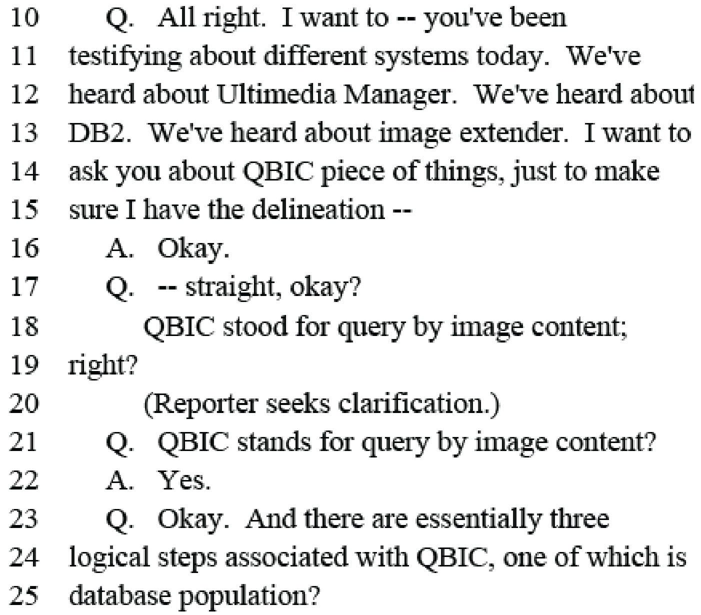 Q. All right. I want to -- you've been testifying about different systems today. We've heard about Ultimedia Manager. We've heard about DB2. We've heard about image extender. I want to ask you about QBIC piece of things, just to make sure I have the delineation -- A. Okay. Q. -- straight, okay? QBIC stood for query by image content; right? (Reporter seeks clarification.) Q. QBIC stands for query by image content? A. Yes. Q. Okay. And there are essentially three logical steps associated with QBIC, one of which is database population? 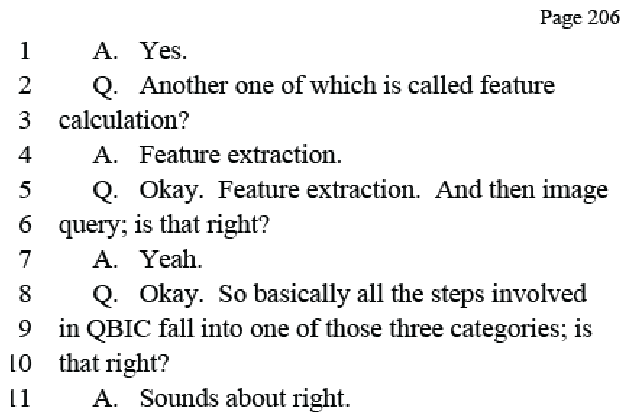 Page 206 A. Yes. Q. Another one of which is called feature calculation? A. Feature extraction. Q. Okay. Feature extraction. And then image query; is that right? A. Yeah. Q. Okay. So basically all the steps involved in QBIC fall into one of those three categories; is that right? A. Sounds about right. *Id.* at 65–66

Element	'004 Claim Recitation	Exemplary Citations in Reference
		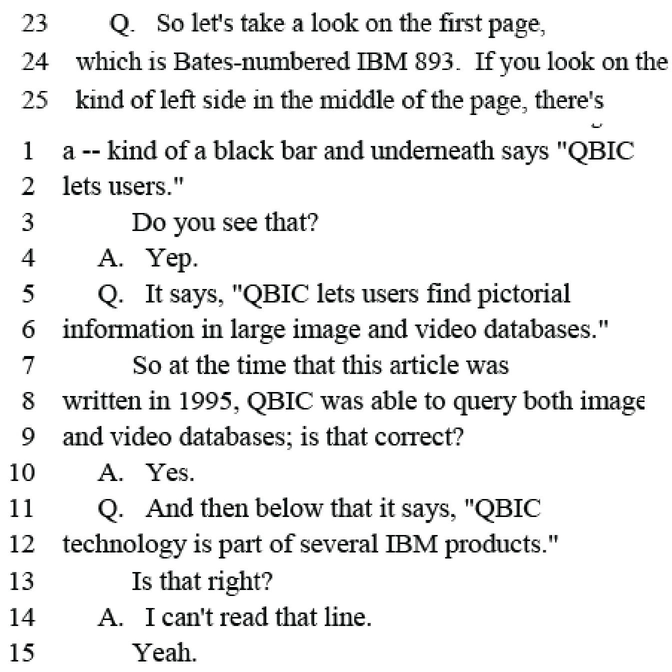 Q. So let's take a look on the first page, which is Bates-numbered IBM 893. If you look on the kind of left side in the middle of the page, there's a -- kind of a black bar and underneath says "QBIC lets users." Do you see that? A. Yep. Q. It says, "QBIC lets users find pictorial information in large image and video databases." So at the time that this article was written in 1995, QBIC was able to query both image and video databases; is that correct? A. Yes. Q. And then below that it says, "QBIC technology is part of several IBM products." Is that right? A. I can't read that line. Yeah. *Id.* at 68 – 72

Element	'004 Claim Recitation	Exemplary Citations in Reference
		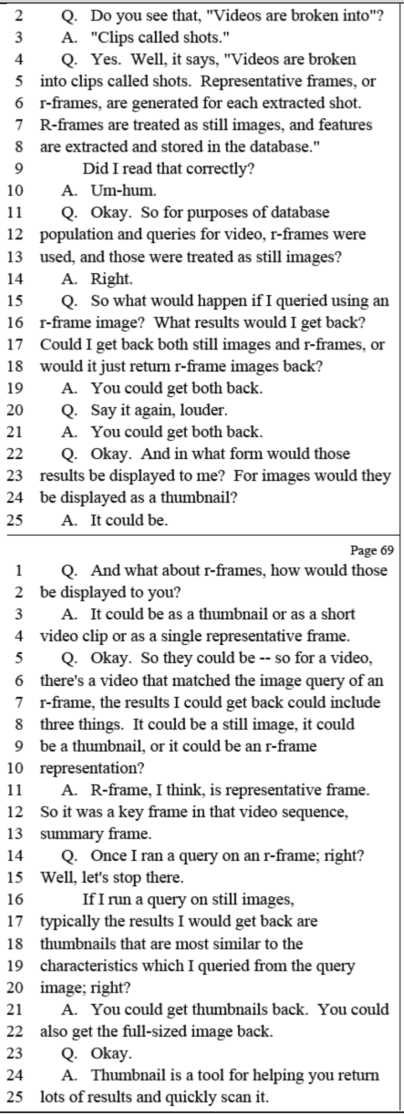 Q. Do you see that, "Videos are broken into"? A. "Clips called shots." Q. Yes. Well, it says, "Videos are broken into clips called shots. Representative frames, or r-frames, are generated for each extracted shot. R-frames are treated as still images, and features are extracted and stored in the database." Did I read that correctly? A. Um-hum. Q. Okay. So for purposes of database population and queries for video, r-frames were used, and those were treated as still images? A. Right. Q. So what would happen if I queried using an r-frame image? What results would I get back? Could I get back both still images and r-frames, or would it just return r-frame images back? A. You could get both back. Q. Say it again, louder. A. You could get both back. Q. Okay. And in what form would those results be displayed to me? For images would they be displayed as a thumbnail? A. It could be. Page 69 Q. And what about r-frames, how would those be displayed to you? A. It could be as a thumbnail or as a short video clip or as a single representative frame. Q. Okay. So they could be -- so for a video, there's a video that matched the image query of an r-frame, the results I could get back could include three things. It could be a still image, it could be a thumbnail, or it could be an r-frame representation? A. R-frame, I think, is representative frame. So it was a key frame in that video sequence, summary frame. Q. Once I ran a query on an r-frame; right? Well, let's stop there. If I run a query on still images, typically the results I would get back are thumbnails that are most similar to the characteristics which I queried from the query image; right? A. You could get thumbnails back. You could also get the full-sized image back. Q. Okay. A. Thumbnail is a tool for helping you return lots of results and quickly scan it.

Element	'004 Claim Recitation	Exemplary Citations in Reference
		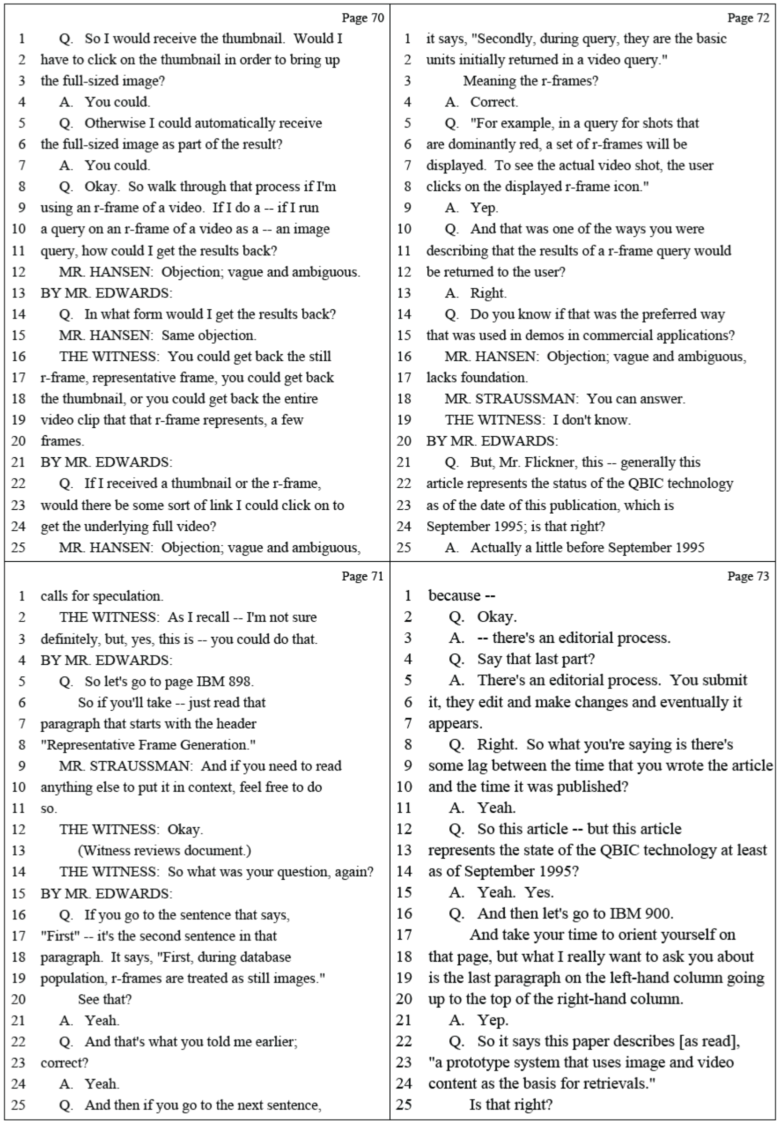

Page 70

Q. So I would receive the thumbnail. Would I have to click on the thumbnail in order to bring up the full-sized image?
A. You could.
Q. Otherwise I could automatically receive the full-sized image as part of the result?
A. You could.
Q. Okay. So walk through that process if I'm using an r-frame of a video. If I do a -- if I run a query on an r-frame of a video as a -- an image query, how could I get the results back?
MR. HANSEN: Objection; vague and ambiguous.
BY MR. EDWARDS:
Q. In what form would I get the results back?
MR. HANSEN: Same objection.
THE WITNESS: You could get back the still r-frame, representative frame, you could get back the thumbnail, or you could get back the entire video clip that that r-frame represents, a few frames.
BY MR. EDWARDS:
Q. If I received a thumbnail or the r-frame, would there be some sort of link I could click on to get the underlying full video?
MR. HANSEN: Objection; vague and ambiguous,

Page 71

calls for speculation.
THE WITNESS: As I recall -- I'm not sure definitely, but, yes, this is -- you could do that.
BY MR. EDWARDS:
Q. So let's go to page IBM 898.
So if you'll take -- just read that paragraph that starts with the header "Representative Frame Generation."
MR. STRAUSSMAN: And if you need to read anything else to put it in context, feel free to do so.
THE WITNESS: Okay.
(Witness reviews document.)
THE WITNESS: So what was your question, again?
BY MR. EDWARDS:
Q. If you go to the sentence that says, "First" -- it's the second sentence in that paragraph. It says, "First, during database population, r-frames are treated as still images."
See that?
A. Yeah.
Q. And that's what you told me earlier; correct?
A. Yeah.
Q. And then if you go to the next sentence,

Page 72

it says, "Secondly, during query, they are the basic units initially returned in a video query."
Meaning the r-frames?
A. Correct.
Q. "For example, in a query for shots that are dominantly red, a set of r-frames will be displayed. To see the actual video shot, the user clicks on the displayed r-frame icon."
A. Yep.
Q. And that was one of the ways you were describing that the results of a r-frame query would be returned to the user?
A. Right.
Q. Do you know if that was the preferred way that was used in demos in commercial applications?
MR. HANSEN: Objection; vague and ambiguous, lacks foundation.
MR. STRAUSSMAN: You can answer.
THE WITNESS: I don't know.
BY MR. EDWARDS:
Q. But, Mr. Flickner, this -- generally this article represents the status of the QBIC technology as of the date of this publication, which is September 1995; is that right?
A. Actually a little before September 1995

Page 73

because --
Q. Okay.
A. -- there's an editorial process.
Q. Say that last part?
A. There's an editorial process. You submit it, they edit and make changes and eventually it appears.
Q. Right. So what you're saying is there's some lag between the time that you wrote the article and the time it was published?
A. Yeah.
Q. So this article -- but this article represents the state of the QBIC technology at least as of September 1995?
A. Yeah. Yes.
Q. And then let's go to IBM 900.
And take your time to orient yourself on that page, but what I really want to ask you about is the last paragraph on the left-hand column going up to the top of the right-hand column.
A. Yep.
Q. So it says this paper describes [as read], "a prototype system that uses image and video content as the basis for retrievals."
Is that right?

Element	'004 Claim Recitation	Exemplary Citations in Reference
		Id. at 191–192 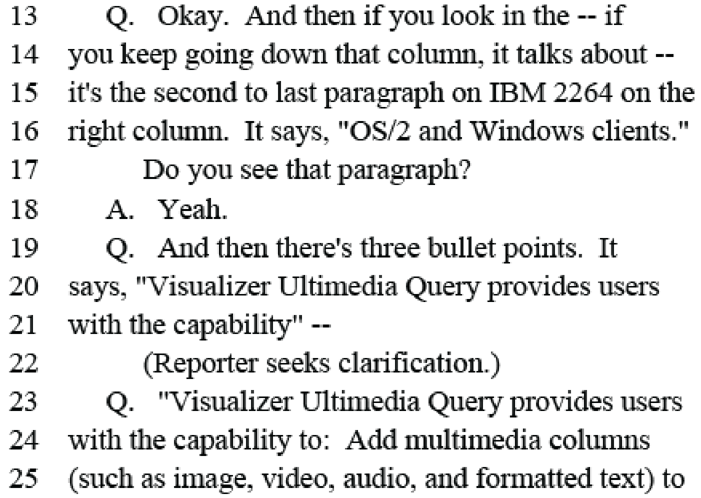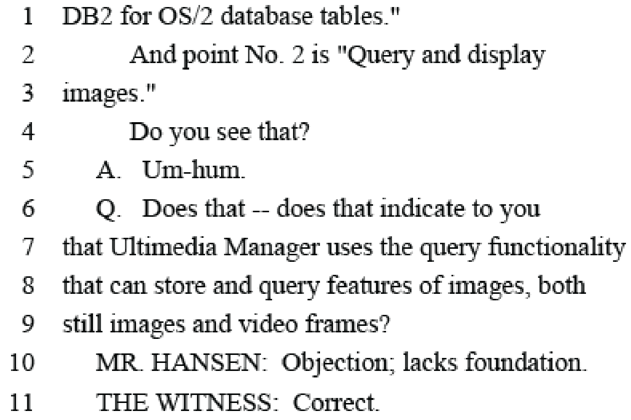 Q. Okay. And then if you look in the -- if you keep going down that column, it talks about -- it's the second to last paragraph on IBM 2264 on the right column. It says, "OS/2 and Windows clients." Do you see that paragraph? A. Yeah. Q. And then there's three bullet points. It says, "Visualizer Ultimedia Query provides users with the capability" -- (Reporter seeks clarification.) Q. "Visualizer Ultimedia Query provides users with the capability to: Add multimedia columns (such as image, video, audio, and formatted text) to DB2 for OS/2 database tables." And point No. 2 is "Query and display images." Do you see that? A. Um-hum. Q. Does that -- does that indicate to you that Ultimedia Manager uses the query functionality that can store and query features of images, both still images and video frames? MR. HANSEN: Objection; lacks foundation. THE WITNESS: Correct. *Id.* at 201

Element	'004 Claim Recitation	Exemplary Citations in Reference
		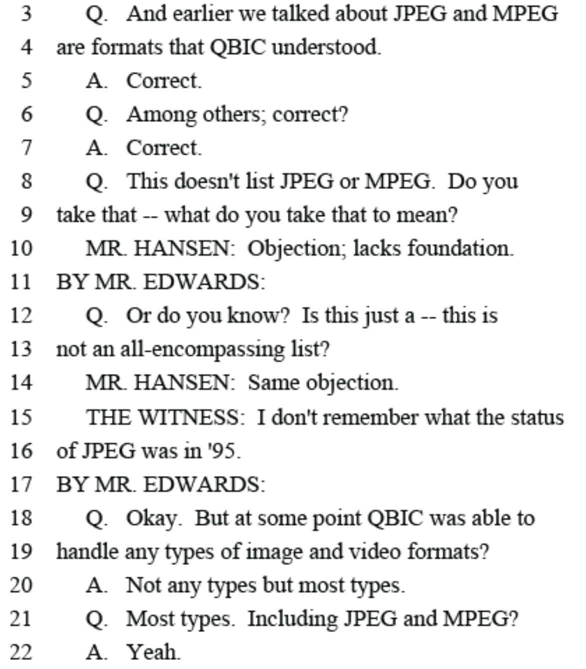 Q. And earlier we talked about JPEG and MPEG are formats that QBIC understood. A. Correct. Q. Among others; correct? A. Correct. Q. This doesn't list JPEG or MPEG. Do you take that -- what do you take that to mean? MR. HANSEN: Objection; lacks foundation. BY MR. EDWARDS: Q. Or do you know? Is this just a -- this is not an all-encompassing list? MR. HANSEN: Same objection. THE WITNESS: I don't remember what the status of JPEG was in '95. BY MR. EDWARDS: Q. Okay. But at some point QBIC was able to handle any types of image and video formats? A. Not any types but most types. Q. Most types. Including JPEG and MPEG? A. Yeah. *Id.* at 314–316

Element	'004 Claim Recitation	Exemplary Citations in Reference
		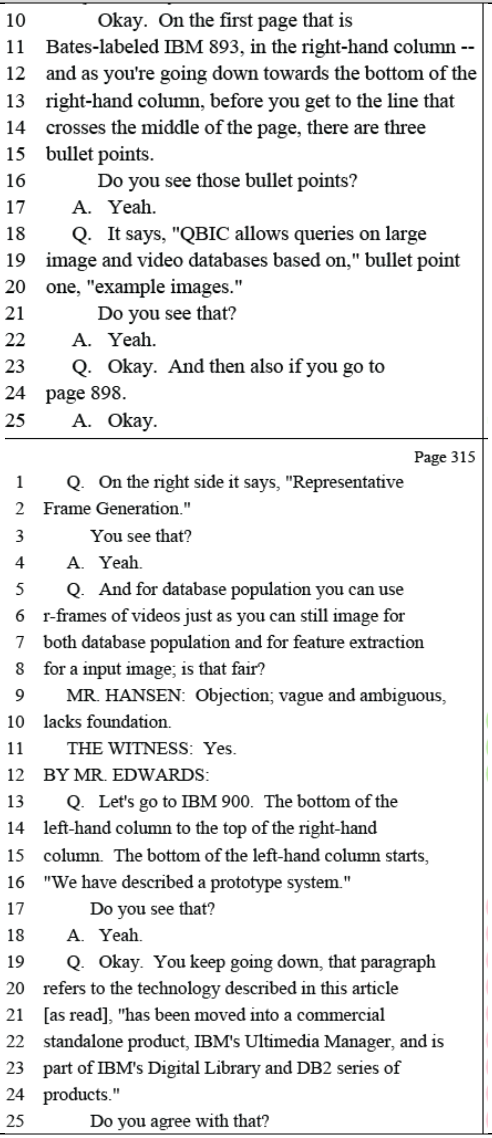 Okay. On the first page that is Bates-labeled IBM 893, in the right-hand column -- and as you're going down towards the bottom of the right-hand column, before you get to the line that crosses the middle of the page, there are three bullet points. Do you see those bullet points? A. Yeah. Q. It says, "QBIC allows queries on large image and video databases based on," bullet point one, "example images." Do you see that? A. Yeah. Q. Okay. And then also if you go to page 898. A. Okay. Page 315 Q. On the right side it says, "Representative Frame Generation." You see that? A. Yeah. Q. And for database population you can use r-frames of videos just as you can still image for both database population and for feature extraction for a input image; is that fair? MR. HANSEN: Objection; vague and ambiguous, lacks foundation. THE WITNESS: Yes. BY MR. EDWARDS: Q. Let's go to IBM 900. The bottom of the left-hand column to the top of the right-hand column. The bottom of the left-hand column starts, "We have described a prototype system." Do you see that? A. Yeah. Q. Okay. You keep going down, that paragraph refers to the technology described in this article [as read], "has been moved into a commercial standalone product, IBM's Ultimedia Manager, and is part of IBM's Digital Library and DB2 series of products." Do you agree with that?

Element	'004 Claim Recitation	Exemplary Citations in Reference
		 A. Yeah. IBM0002263 at 2264-65: 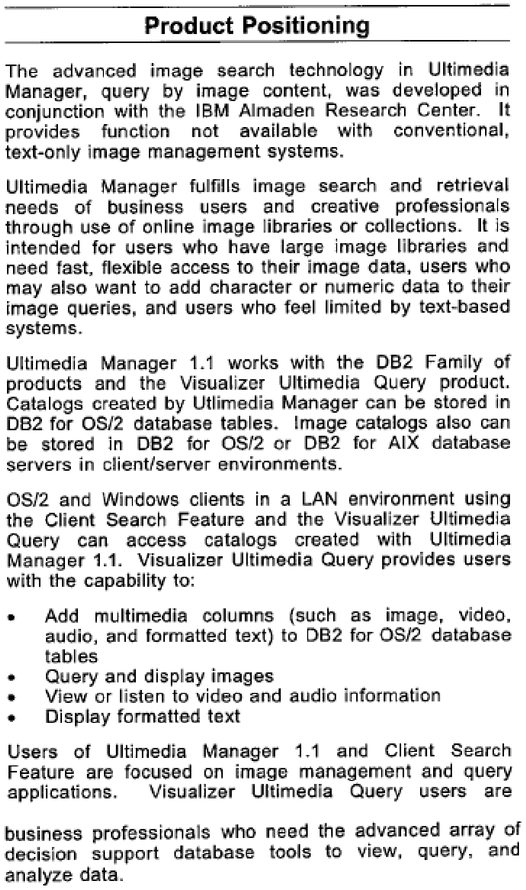 **Product Positioning** The advanced image search technology in Ultimedia Manager, query by image content, was developed in conjunction with the IBM Almaden Research Center. It provides function not available with conventional, text-only image management systems. Ultimedia Manager fulfills image search and retrieval needs of business users and creative professionals through use of online image libraries or collections. It is intended for users who have large image libraries and need fast, flexible access to their image data, users who may also want to add character or numeric data to their image queries, and users who feel limited by text-based systems. Ultimedia Manager 1.1 works with the DB2 Family of products and the Visualizer Ultimedia Query product. Catalogs created by Utlimedia Manager can be stored in DB2 for OS/2 database tables. Image catalogs also can be stored in DB2 for OS/2 or DB2 for AIX database servers in client/server environments. OS/2 and Windows clients in a LAN environment using the Client Search Feature and the Visualizer Ultimedia Query can access catalogs created with Ultimedia Manager 1.1. Visualizer Ultimedia Query provides users with the capability to: • Add multimedia columns (such as image, video, audio, and formatted text) to DB2 for OS/2 database tables • Query and display images • View or listen to video and audio information • Display formatted text Users of Ultimedia Manager 1.1 and Client Search Feature are focused on image management and query applications. Visualizer Ultimedia Query users are business professionals who need the advanced array of decision support database tools to view, query, and analyze data. The QBIC system was implemented in IBM commercial products, including in IBM's Ultimedia Manager, the DB2 series of products such as the Image Extender

Element	'004 Claim Recitation	Exemplary Citations in Reference
		in DB2 Universal Database, and the Enterprise Information Portal. All of the features referenced in this chart were incorporated into these commercial implementations before November 6, 2000, as reflected in a combination of scientific papers, IBM's internal documentation, and the QBIC demo software applications. *See, e.g.*, Flickner 8/29/23 Depo Tr. at 33:9-24, 71:21-72:10, 164:21-166:2, 166:11-175:24, 178:5-186:17, 186:21-192:4, 192:14-194:1, 195:16-196:4, 194:1-308:1; IBM 000001-617; IBM 000747; IBM 000777-880; IBM 000893-902; IBM 0002315-320; IBM 0002390-404; IBM 0002418-430. To the extent it is found that QBIC does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that QBIC does not anticipate this claim, QBIC renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine QBIC with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1A]	analyzing, via a mobile device, a scene represented by the video stream for at least one object;	QBIC uses still images and video (frames) for query and database population: Query by Image and Video Content: The QBIC System (IBM 0000893–902): "We have developed the QBIC (Query by Image Content) system to explore content-based retrieval methods. QBIC allows queries on large image and video databases based on • example images, • user-constructed sketches and drawings, • selected color and texture patterns, • camera and object motion, and • other graphical information." *Id.* at pp. 1–3.

Element	'004 Claim Recitation	Exemplary Citations in Reference
		"To achieve this functionality, QBIC has two main components: database population (the process of creating an image database) and database query. During the population, images and videos are processed to extract features describing their content-colors, textures, shapes, and camera and object motion-and the features are stored in a database. During the query, the user composes a query graphically. Features are generated from the graphical query and then input to a matching engine that finds images or videos from the database with similar features. Figure 2 shows the system architecture." *Id.* at p. 3. "For both population and query, the QBIC data model has • still images or scenes (full images) that contain objects (subsets of an image), and • video shots that consist of sets of contiguous frames and contain motion objects. For still images, the QBIC data model distinguishes between 'scenes' (or images) and 'objects.' A scene is an image or single representative frame of video. An object is a part of a scene-for example, the fox in Figure 3-or a moving entity in a video. For still image database population, features are extracted from images and objects and stored in a database as shown in the top left part of Figure 2. Videos are broken into clips called shots. Representative frames, or r-frames, are generated for each extracted shot. R-frames are treated as still images, and features are extracted and stored in the database. Further processing of shots generates motion objects-for example, a car moving across the screen. Queries are allowed on objects ('Find images with a red, round object'), scenes ('Find images that have approximately 30-percent red and IS-percent blue colors'), shots ('Find all shots panning from left to right'), or any combination ('Find images that have 30 percent red and contain a blue textured object')." *Id.* at p. 3. "In still image database population, the images are reduced to a standard-sized icon called a thumbnail and annotated with any available text information. Object identification is an optional but key part of this step. It lets users manually, semiautomatically, or fully automatically identify interesting regions-which we call objects-in the images. . . .

Element	'004 Claim Recitation	Exemplary Citations in Reference
		For video data, database population has three major components: • shot detection, • representative frame creation for each shot, and • derivation of a layered representation of coherently moving structures/ objects. . . . Once the shot boundaries have been detected, each shot is represented using an r-frame. R-frames are used for several purposes. First, during database population, r-frames are treated as still images in which objects can be identified by using the previously described methods." *Id*. at pp. 3–6. Image, Audio, and Video Extenders Administration and Programming (IBM 000001–617): "Your applications can even search for images by content. Imagine an application that uses visual examples to search for images. With such an application, users could select an example image and have the application find other images that have colors or textures similar to those in the example. With DB2 extenders' Query by Image Content (QBIC) capability, you can create applications that search for images in this visual way." *Id*. at p. 4. "To make an image available for searching by content, you catalog the image. When you catalog an image, the Image Extender analyzes the image, by computing the feature values for the image, and stores the values in a QBIC catalog. When you search for an image by content, your query identifies one or more features for the search (such as average color), a source for each feature (such as an example image), and a target set of cataloged images. The Image Extender computes the feature value of the source and compares it to the cataloged feature values for the target images. It then computes a score that indicates how similar the feature values of the target images are to the source. You can have the Image Extender return the images whose features are most similar to the source. The Image Extender will return the handle of each image and the image score. You can also have the Image Extender return only the score of a single image." *Id*. at pp. 20–21.

Element	'004 Claim Recitation	Exemplary Citations in Reference
		Flickner Depo at 88: 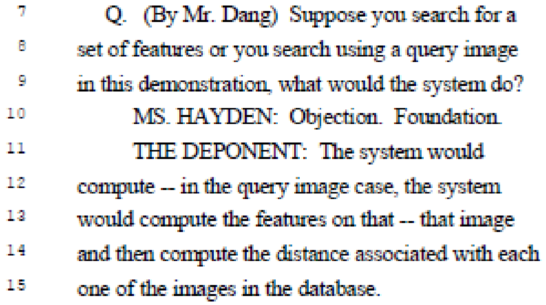 Q. (By Mr. Dang) Suppose you search for a set of features or you search using a query image in this demonstration, what would the system do? MS. HAYDEN: Objection. Foundation. THE DEPONENT: The system would compute -- in the query image case, the system would compute the features on that -- that image and then compute the distance associated with each one of the images in the database. Flickner Depo at 89-90: 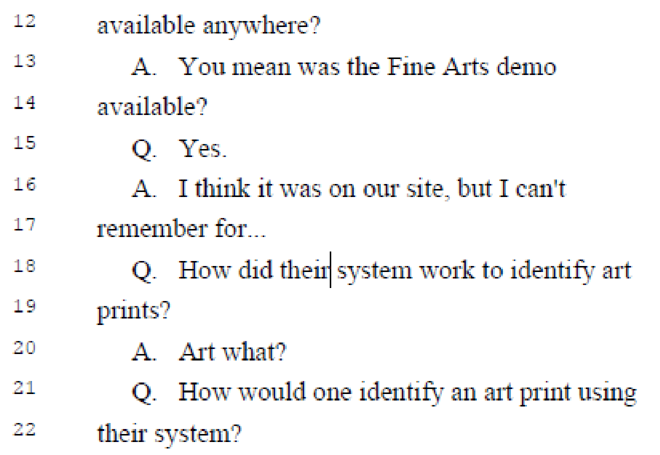 Q. Was that used to the QBIC system publicly available anywhere? A. You mean was the Fine Arts demo available? Q. Yes. A. I think it was on our site, but I can't remember for... Q. How did their system work to identify art prints? A. Art what? Q. How would one identify an art print using their system?

Element	'004 Claim Recitation	Exemplary Citations in Reference
		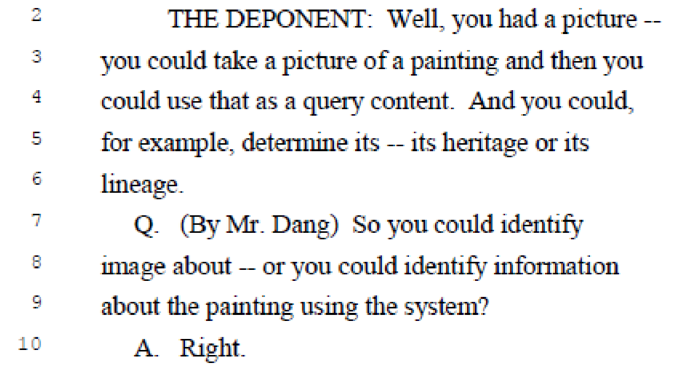 THE DEPONENT: Well, you had a picture -- you could take a picture of a painting and then you could use that as a query content. And you could, for example, determine its -- its heritage or its lineage. Q. (By Mr. Dang) So you could identify image about -- or you could identify information about the painting using the system? A. Right. Flickner Depo at 39-40: 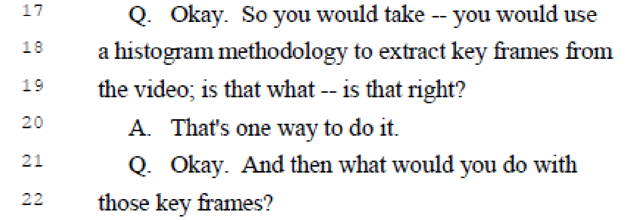 Q. Okay. So you would take -- you would use a histogram methodology to extract key frames from the video; is that what -- is that right? A. That's one way to do it. Q. Okay. And then what would you do with those key frames?

Element	'004 Claim Recitation	Exemplary Citations in Reference
		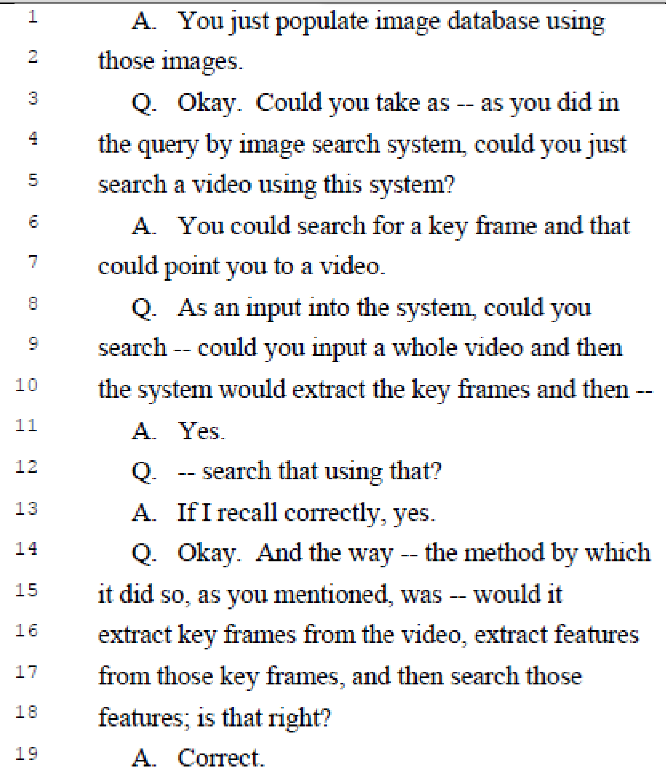 A. You just populate image database using those images. Q. Okay. Could you take as -- as you did in the query by image search system, could you just search a video using this system? A. You could search for a key frame and that could point you to a video. Q. As an input into the system, could you search -- could you input a whole video and then the system would extract the key frames and then -- A. Yes. Q. -- search that using that? A. If I recall correctly, yes. Q. Okay. And the way -- the method by which it did so, as you mentioned, was -- would it extract key frames from the video, extract features from those key frames, and then search those features; is that right? A. Correct. Flickner Depo at 54-55:

Element	'004 Claim Recitation	Exemplary Citations in Reference
		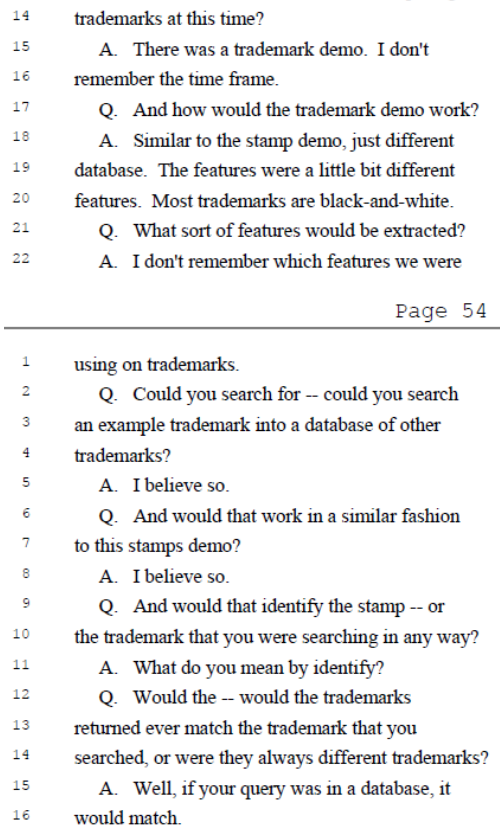 Q. Was there also a demo online comparing trademarks at this time? A. There was a trademark demo. I don't remember the time frame. Q. And how would the trademark demo work? A. Similar to the stamp demo, just different database. The features were a little bit different features. Most trademarks are black-and-white. Q. What sort of features would be extracted? A. I don't remember which features we were Page 54 using on trademarks. Q. Could you search for -- could you search an example trademark into a database of other trademarks? A. I believe so. Q. And would that work in a similar fashion to this stamps demo? A. I believe so. Q. And would that identify the stamp -- or the trademark that you were searching in any way? A. What do you mean by identify? Q. Would the -- would the trademarks returned ever match the trademark that you searched, or were they always different trademarks? A. Well, if your query was in a database, it would match. QBIC could integrate with a camera, such as a portable camera; a user could take an image using a portable camera and submit that image to the QBIC system for query. Flickner Depo at, e.g., 24, 41-42, 62-63 (user could take a picture or video and submit it to the QBIC system):

Element	'004 Claim Recitation	Exemplary Citations in Reference
		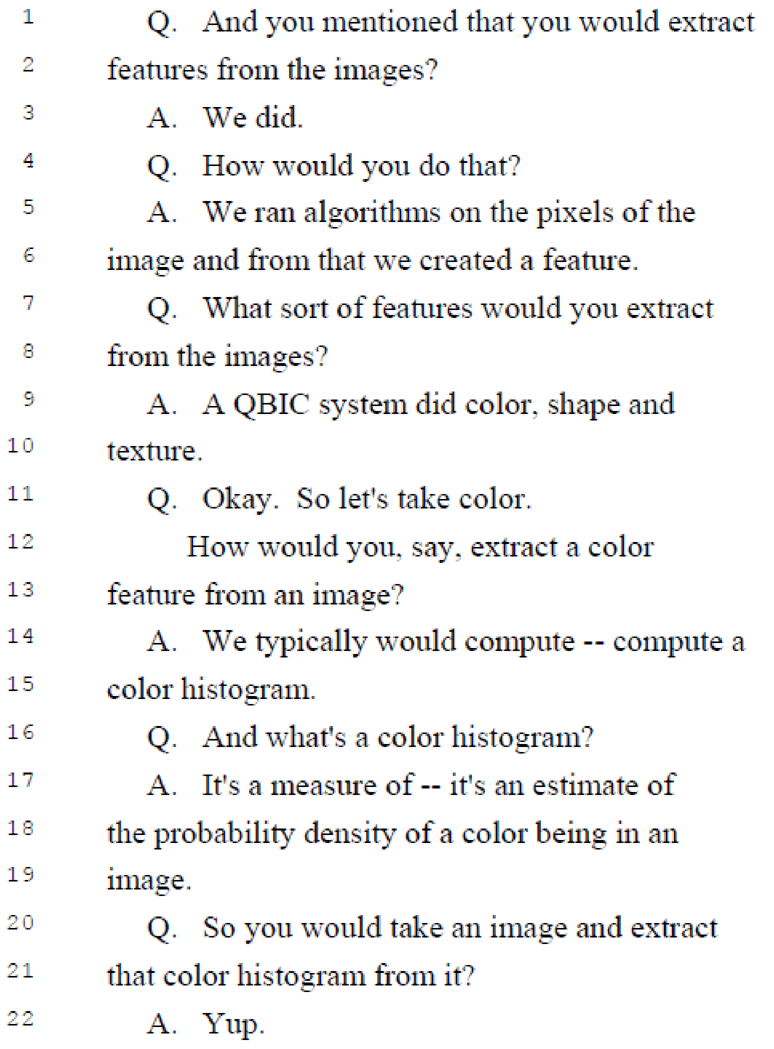 Q. And you mentioned that you would extract features from the images? A. We did. Q. How would you do that? A. We ran algorithms on the pixels of the image and from that we created a feature. Q. What sort of features would you extract from the images? A. A QBIC system did color, shape and texture. Q. Okay. So let's take color. How would you, say, extract a color feature from an image? A. We typically would compute -- compute a color histogram. Q. And what's a color histogram? A. It's a measure of -- it's an estimate of the probability density of a color being in an image. Q. So you would take an image and extract that color histogram from it? A. Yup.

		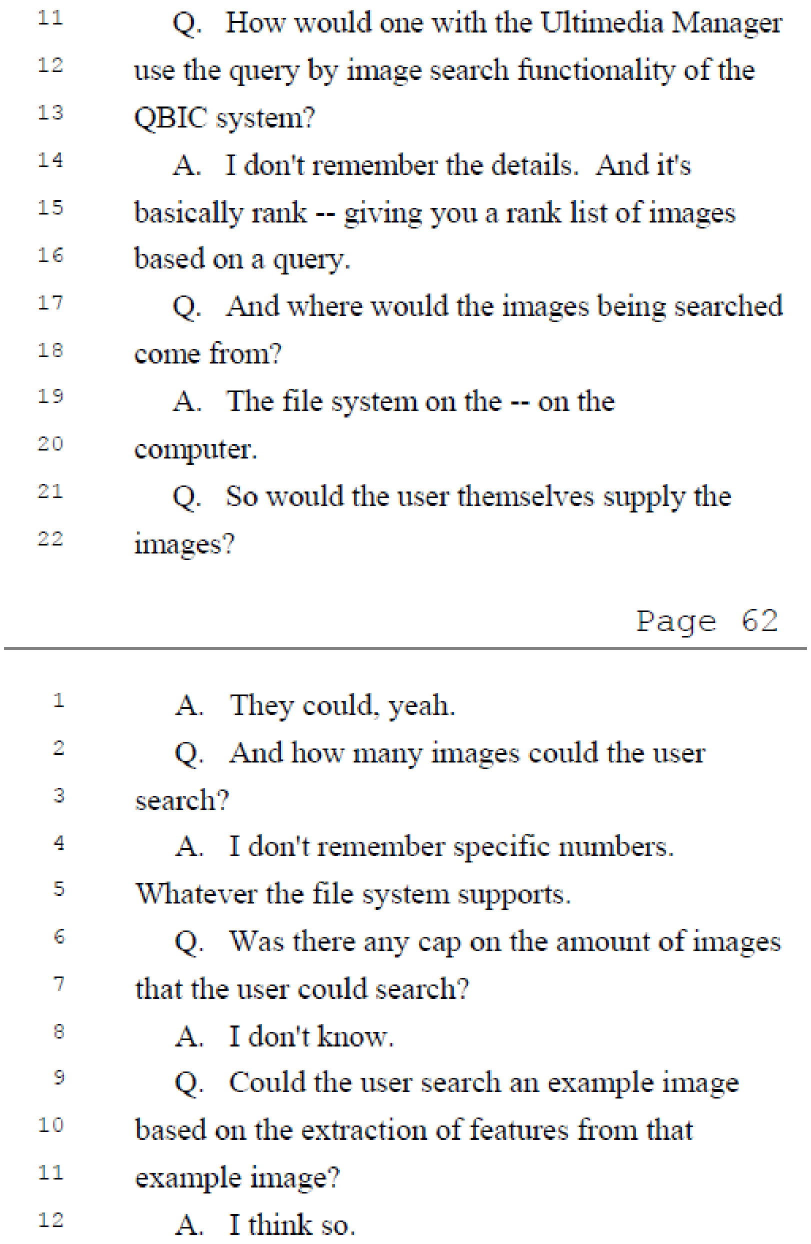 Q. How would one with the Ultimedia Manager use the query by image search functionality of the QBIC system? A. I don't remember the details. And it's basically rank -- giving you a rank list of images based on a query. Q. And where would the images being searched come from? A. The file system on the -- on the computer. Q. So would the user themselves supply the images? Page 62 A. They could, yeah. Q. And how many images could the user search? A. I don't remember specific numbers. Whatever the file system supports. Q. Was there any cap on the amount of images that the user could search? A. I don't know. Q. Could the user search an example image based on the extraction of features from that example image? A. I think so.

Element	'004 Claim Recitation	Exemplary Citations in Reference
		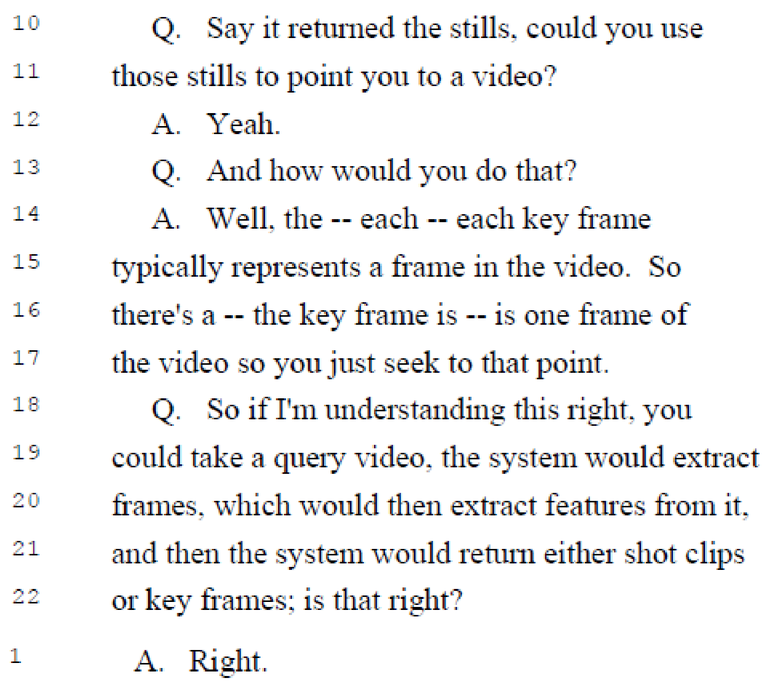 Q. Say it returned the stills, could you use those stills to point you to a video? A. Yeah. Q. And how would you do that? A. Well, the -- each -- each key frame typically represents a frame in the video. So there's a -- the key frame is -- is one frame of the video so you just seek to that point. Q. So if I'm understanding this right, you could take a query video, the system would extract frames, which would then extract features from it, and then the system would return either shot clips or key frames; is that right? A. Right. Flickner 8/29/23 Depo at 34:

Element	'004 Claim Recitation	Exemplary Citations in Reference
		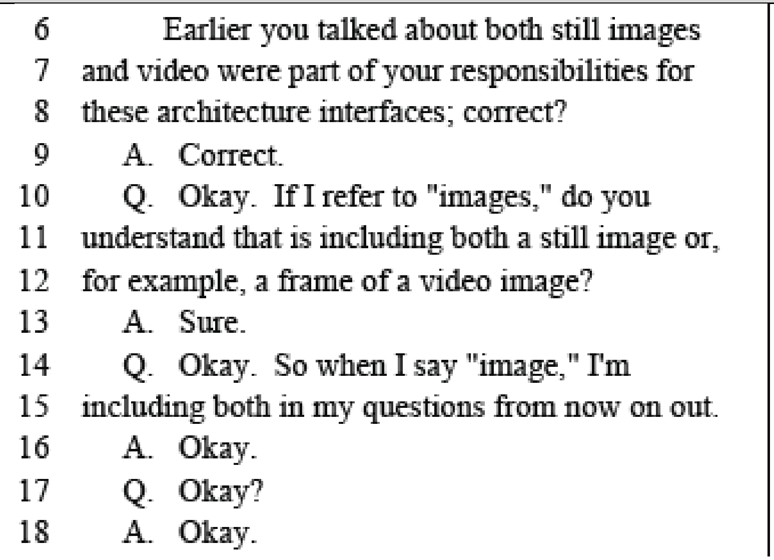 Earlier you talked about both still images and video were part of your responsibilities for these architecture interfaces; correct? A. Correct. Q. Okay. If I refer to "images," do you understand that is including both a still image or, for example, a frame of a video image? A. Sure. Q. Okay. So when I say "image," I'm including both in my questions from now on out. A. Okay. Q. Okay? A. Okay. *Id.* at 53: 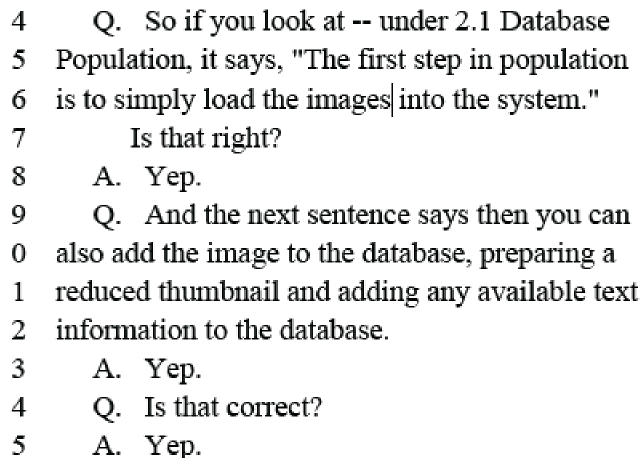 Q. So if you look at -- under 2.1 Database Population, it says, "The first step in population is to simply load the images into the system." Is that right? A. Yep. Q. And the next sentence says then you can also add the image to the database, preparing a reduced thumbnail and adding any available text information to the database. A. Yep. Q. Is that correct? A. Yep.

Element	'004 Claim Recitation	Exemplary Citations in Reference
		Id. at 60–61: 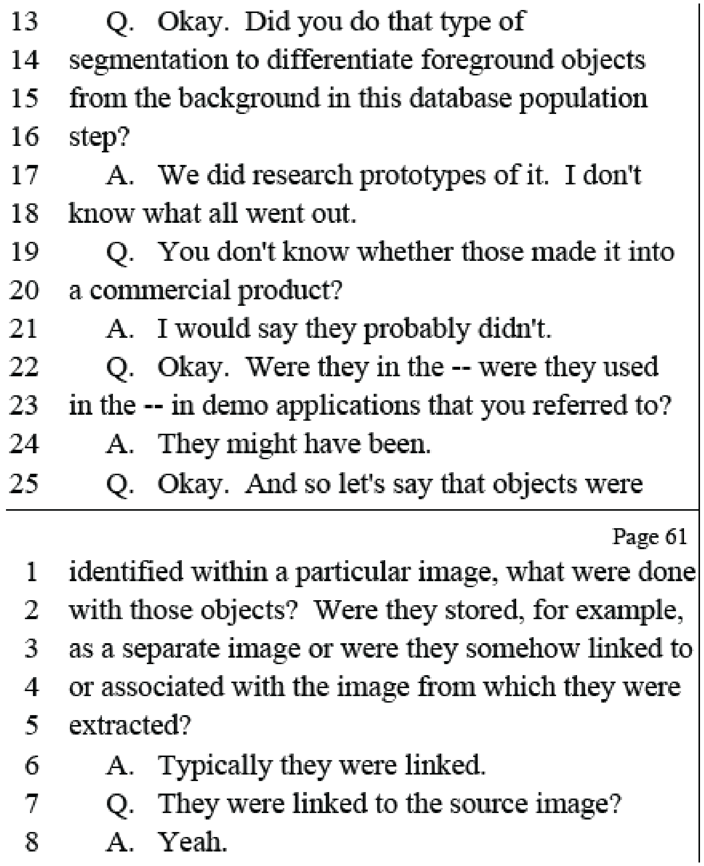 Q. Okay. Did you do that type of segmentation to differentiate foreground objects from the background in this database population step? A. We did research prototypes of it. I don't know what all went out. Q. You don't know whether those made it into a commercial product? A. I would say they probably didn't. Q. Okay. Were they in the -- were they used in the -- in demo applications that you referred to? A. They might have been. Q. Okay. And so let's say that objects were identified within a particular image, what were done with those objects? Were they stored, for example, as a separate image or were they somehow linked to or associated with the image from which they were extracted? Page 61 A. Typically they were linked. Q. They were linked to the source image? A. Yeah. *Id.* at 32–33:

Element	'004 Claim Recitation	Exemplary Citations in Reference
		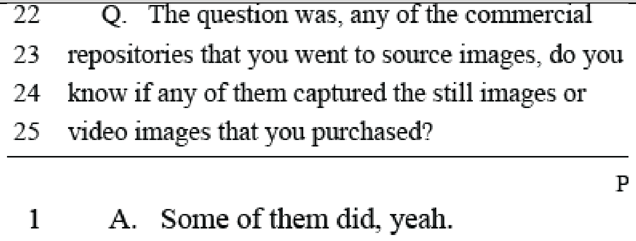 Q. The question was, any of the commercial repositories that you went to source images, do you know if any of them captured the still images or video images that you purchased? P A. Some of them did, yeah. *Id.* at 74–75:

			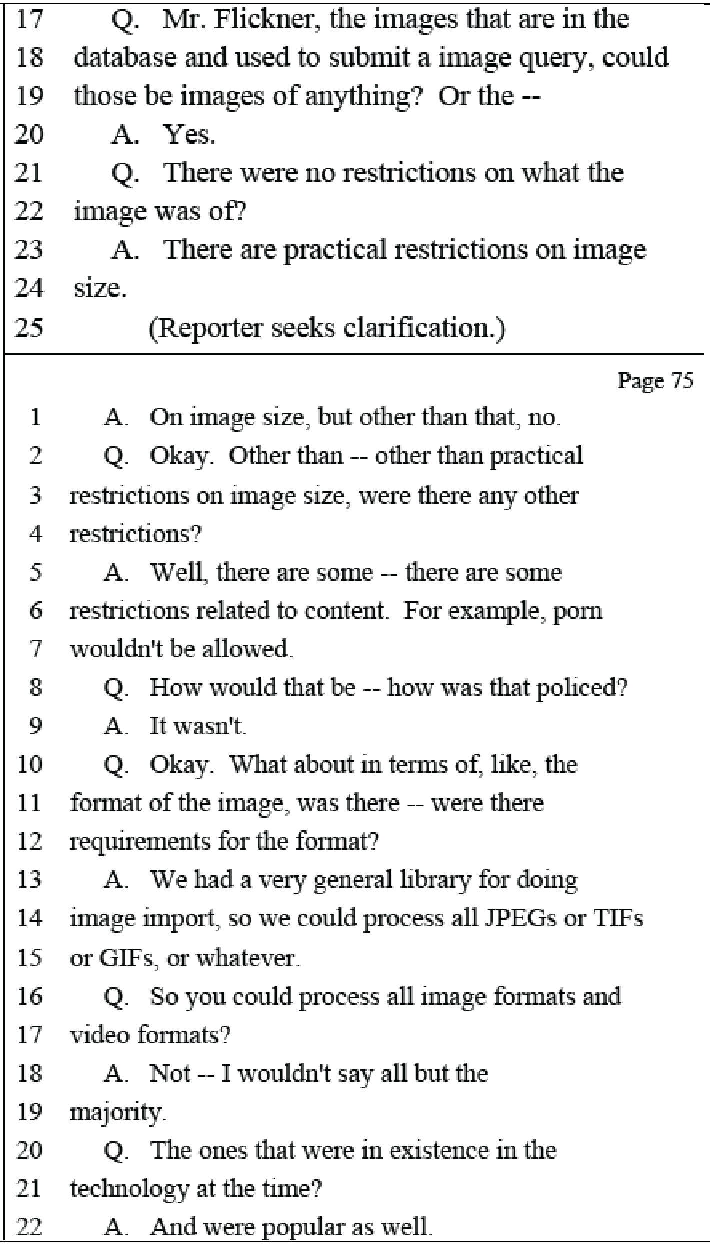 Q. Mr. Flickner, the images that are in the database and used to submit a image query, could those be images of anything? Or the -- A. Yes. Q. There were no restrictions on what the image was of? A. There are practical restrictions on image size. (Reporter seeks clarification.) Page 75 A. On image size, but other than that, no. Q. Okay. Other than -- other than practical restrictions on image size, were there any other restrictions? A. Well, there are some -- there are some restrictions related to content. For example, porn wouldn't be allowed. Q. How would that be -- how was that policed? A. It wasn't. Q. Okay. What about in terms of, like, the format of the image, was there -- were there requirements for the format? A. We had a very general library for doing image import, so we could process all JPEGs or TIFs or GIFs, or whatever. Q. So you could process all image formats and video formats? A. Not -- I wouldn't say all but the majority. Q. The ones that were in existence in the technology at the time? A. And were popular as well.

Element	'004 Claim Recitation	Exemplary Citations in Reference
		Id. at 114: 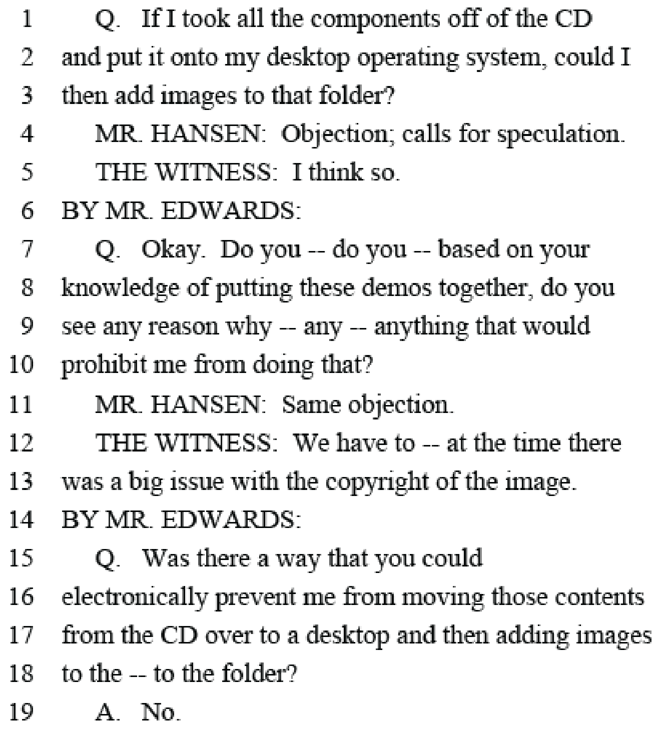 Q. If I took all the components off of the CD and put it onto my desktop operating system, could I then add images to that folder? MR. HANSEN: Objection; calls for speculation. THE WITNESS: I think so. BY MR. EDWARDS: Q. Okay. Do you -- do you -- based on your knowledge of putting these demos together, do you see any reason why -- any -- anything that would prohibit me from doing that? MR. HANSEN: Same objection. THE WITNESS: We have to -- at the time there was a big issue with the copyright of the image. BY MR. EDWARDS: Q. Was there a way that you could electronically prevent me from moving those contents from the CD over to a desktop and then adding images to the -- to the folder? A. No. *Id.* at 202:

Element	'004 Claim Recitation	Exemplary Citations in Reference
		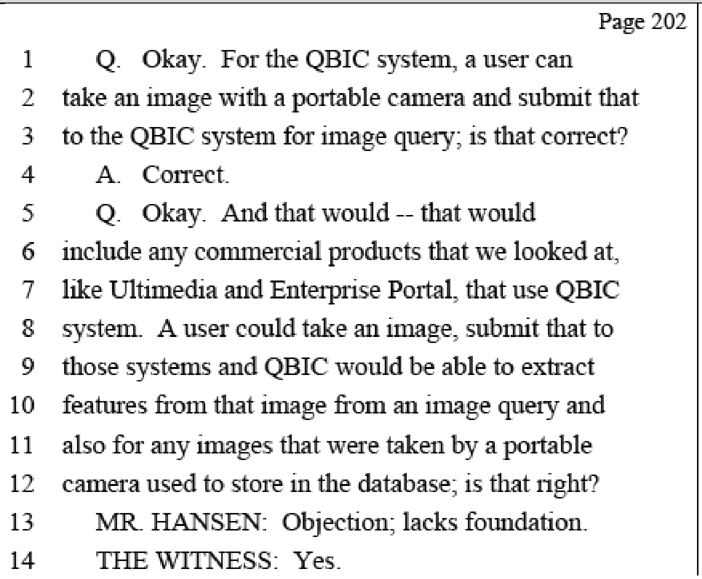 Page 202 Q. Okay. For the QBIC system, a user can take an image with a portable camera and submit that to the QBIC system for image query; is that correct? A. Correct. Q. Okay. And that would -- that would include any commercial products that we looked at, like Ultimedia and Enterprise Portal, that use QBIC system. A user could take an image, submit that to those systems and QBIC would be able to extract features from that image from an image query and also for any images that were taken by a portable camera used to store in the database; is that right? MR. HANSEN: Objection; lacks foundation. THE WITNESS: Yes. *Id.* at 300: 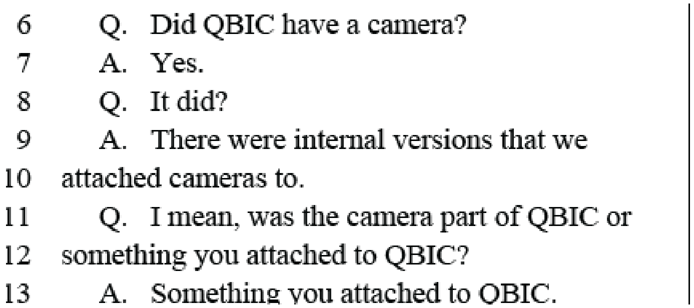 Q. Did QBIC have a camera? A. Yes. Q. It did? A. There were internal versions that we attached cameras to. Q. I mean, was the camera part of QBIC or something you attached to QBIC? A. Something you attached to QBIC. *Id.* at 312:

Element	'004 Claim Recitation	Exemplary Citations in Reference
		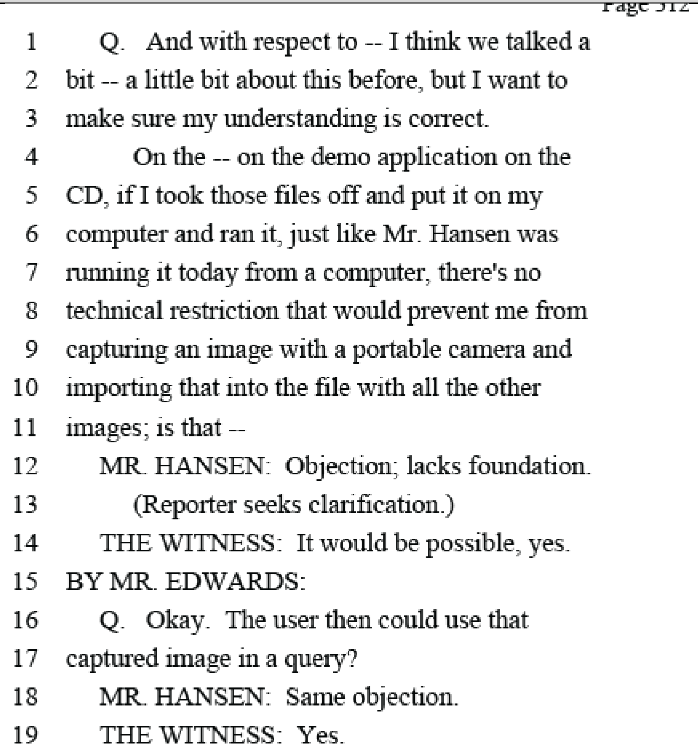 Page 312 Q. And with respect to -- I think we talked a bit -- a little bit about this before, but I want to make sure my understanding is correct. On the -- on the demo application on the CD, if I took those files off and put it on my computer and ran it, just like Mr. Hansen was running it today from a computer, there's no technical restriction that would prevent me from capturing an image with a portable camera and importing that into the file with all the other images; is that -- MR. HANSEN: Objection; lacks foundation. (Reporter seeks clarification.) THE WITNESS: It would be possible, yes. BY MR. EDWARDS: Q. Okay. The user then could use that captured image in a query? MR. HANSEN: Same objection. THE WITNESS: Yes.

Element	'004 Claim Recitation	Exemplary Citations in Reference
		QBIC Demo (IBM000747) – showing compatibility with images captured by a mobile phone camera  (Use a Java-capable Web browser to enable additional functionality.) **Query was:** Example: =AS_test/check1.jpg Query Type: Color Layout

Element	'004 Claim Recitation	Exemplary Citations in Reference
		

Element	’004 Claim Recitation	Exemplary Citations in Reference
		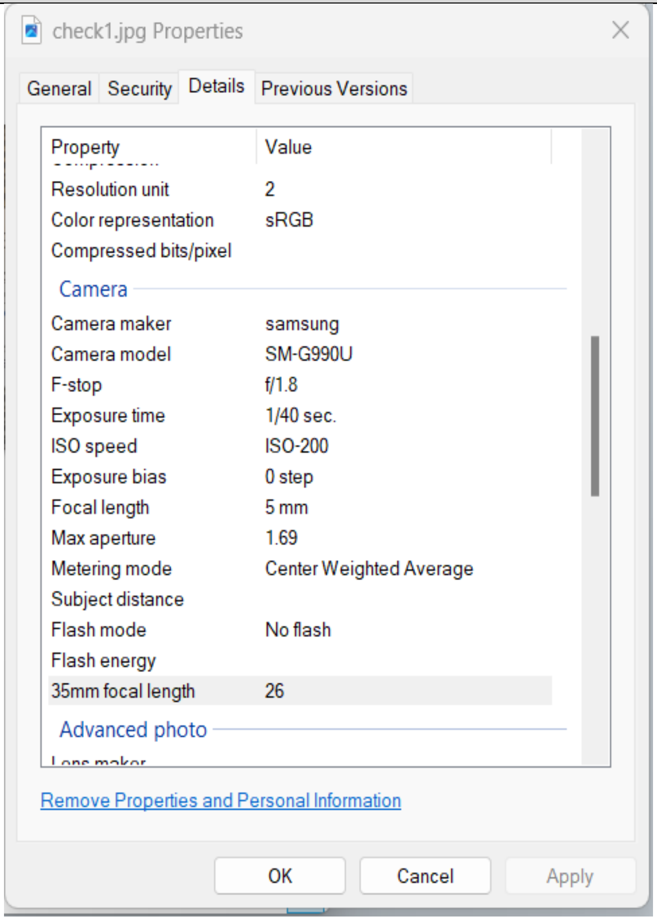 check1.jpg Properties Property: Value Resolution unit: 2 Color representation: sRGB Compressed bits/pixel: Camera Camera maker: samsung Camera model: SM-G990U F-stop: f/1.8 Exposure time: 1/40 sec. ISO speed: ISO-200 Exposure bias: 0 step Focal length: 5 mm Max aperture: 1.69 Metering mode: Center Weighted Average Subject distance: Flash mode: No flash Flash energy: 35mm focal length: 26 Advanced photo

Element	’004 Claim Recitation	Exemplary Citations in Reference
		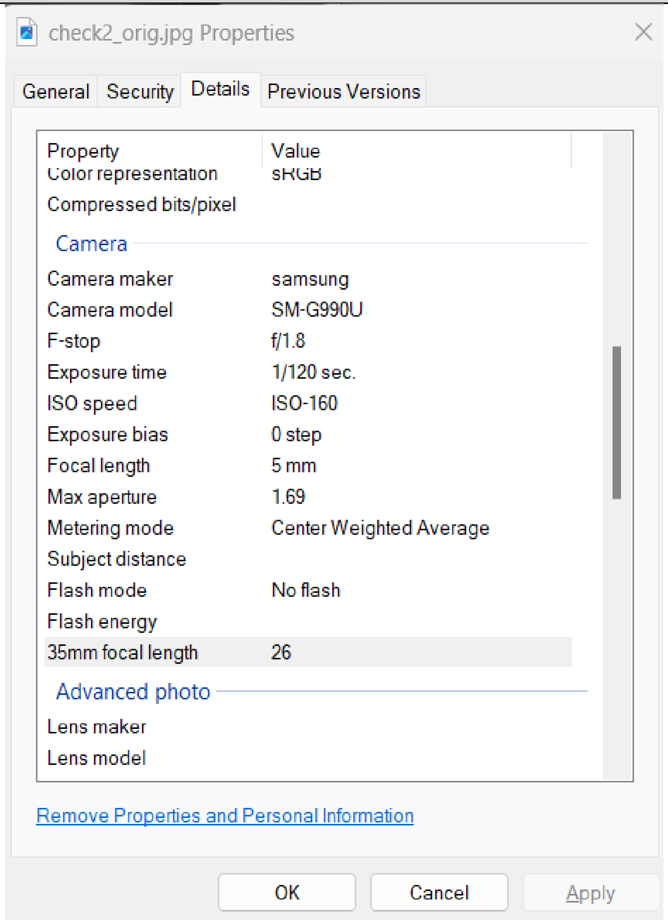

check2_orig.jpg Properties

Property	Value
Color representation	sRGB
Compressed bits/pixel	
Camera	
Camera maker	samsung
Camera model	SM-G990U
F-stop	f/1.8
Exposure time	1/120 sec.
ISO speed	ISO-160
Exposure bias	0 step
Focal length	5 mm
Max aperture	1.69
Metering mode	Center Weighted Average
Subject distance	
Flash mode	No flash
Flash energy	
35mm focal length	26
Advanced photo	
Lens maker	
Lens model	

Element	'004 Claim Recitation	Exemplary Citations in Reference
		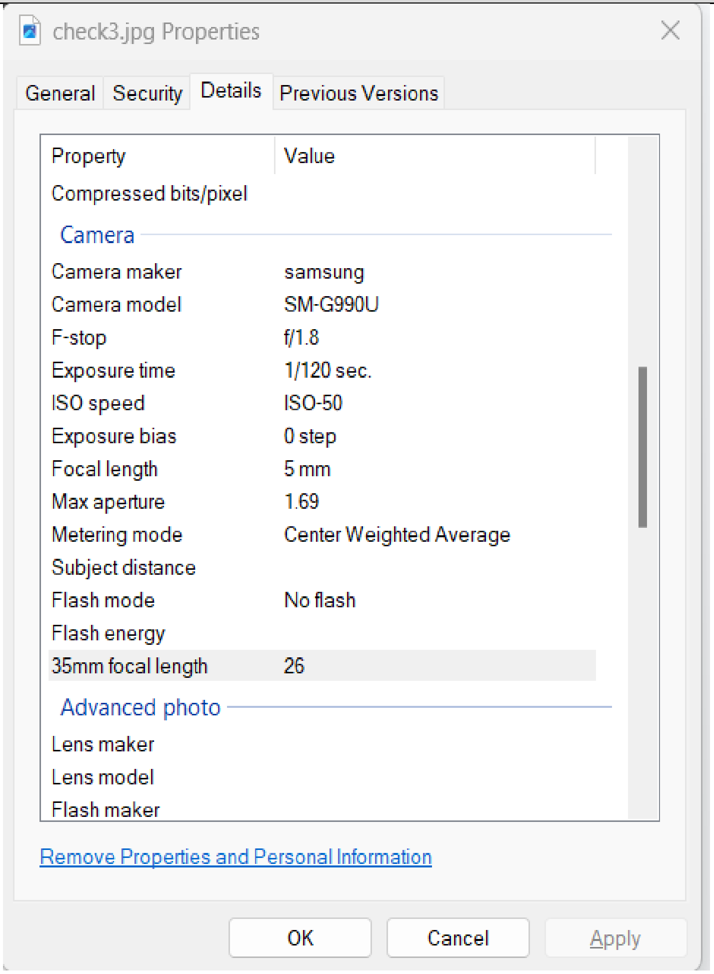 Property Value Compressed bits/pixel Camera Camera maker samsung Camera model SM-G990U F-stop f/1.8 Exposure time 1/120 sec. ISO speed ISO-50 Exposure bias 0 step Focal length 5 mm Max aperture 1.69 Metering mode Center Weighted Average Subject distance Flash mode No flash Flash energy 35mm focal length 26 Advanced photo Lens maker Lens model Flash maker

Element	'004 Claim Recitation	Exemplary Citations in Reference
		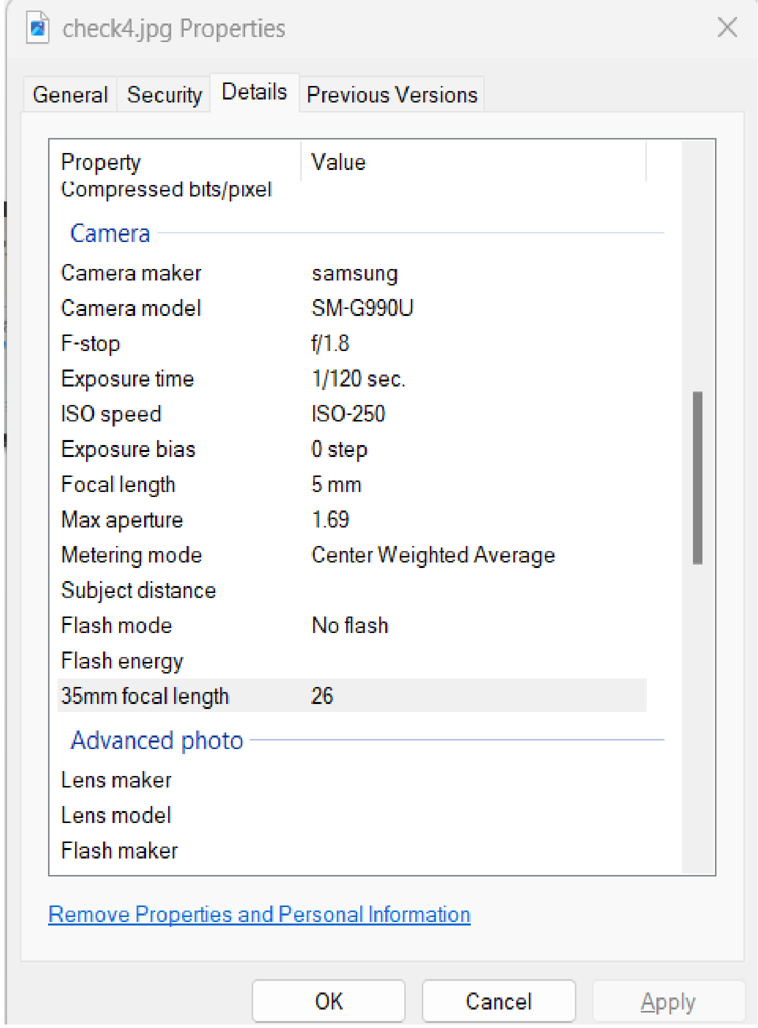

Property	Value
Compressed bits/pixel	
Camera	
Camera maker	samsung
Camera model	SM-G990U
F-stop	f/1.8
Exposure time	1/120 sec.
ISO speed	ISO-250
Exposure bias	0 step
Focal length	5 mm
Max aperture	1.69
Metering mode	Center Weighted Average
Subject distance	
Flash mode	No flash
Flash energy	
35mm focal length	26
Advanced photo	
Lens maker	
Lens model	
Flash maker	

Element	'004 Claim Recitation	Exemplary Citations in Reference
		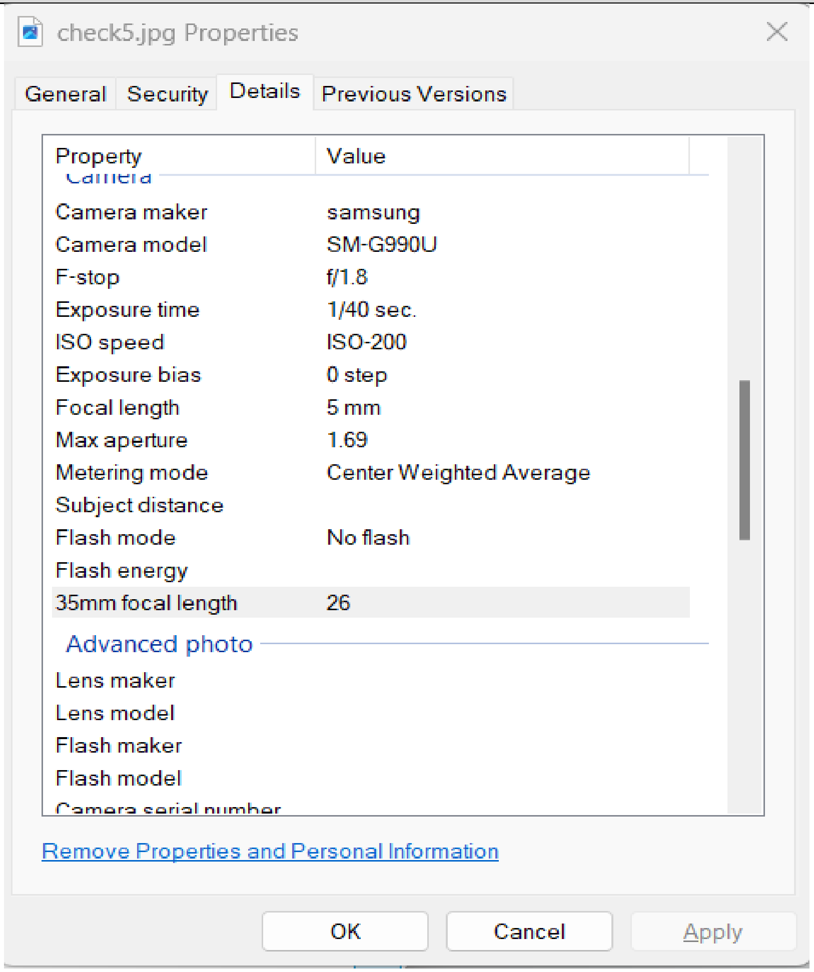 Property Value Camera Camera maker samsung Camera model SM-G990U F-stop f/1.8 Exposure time 1/40 sec. ISO speed ISO-200 Exposure bias 0 step Focal length 5 mm Max aperture 1.69 Metering mode Center Weighted Average Subject distance Flash mode No flash Flash energy 35mm focal length 26 Advanced photo Lens maker Lens model Flash maker Flash model Camera serial number To the extent it is found that QBIC does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that QBIC does not anticipate

Element	'004 Claim Recitation	Exemplary Citations in Reference
		this claim, QBIC renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine QBIC with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1B]	deriving at least one characteristic of the video stream;	QBIC discloses the limitation. The QBIC Project: Querying Images by Content Using Color, Texture, and Shape (IBM 0002390–404) at 1 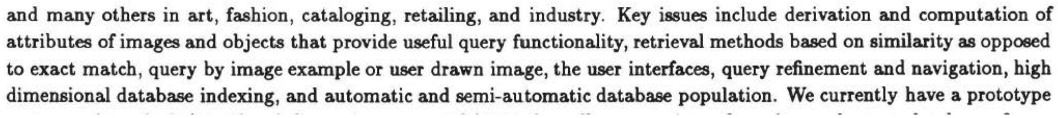 and many others in art, fashion, cataloging, retailing, and industry. Key issues include derivation and computation of attributes of images and objects that provide useful query functionality, retrieval methods based on similarity as opposed to exact match, query by image example or user drawn image, the user interfaces, query refinement and navigation, high dimensional database indexing, and automatic and semi-automatic database population. We currently have a prototype The QBIC Project: Querying Images by Content Using Color, Texture, and Shape (IBM 0002390–404) at 3 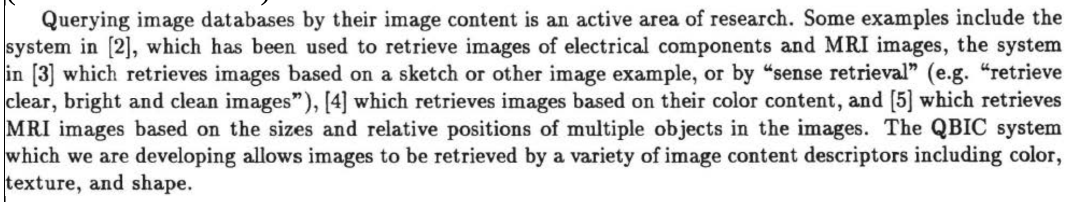 Querying image databases by their image content is an active area of research. Some examples include the system in [2], which has been used to retrieve images of electrical components and MRI images, the system in [3] which retrieves images based on a sketch or other image example, or by "sense retrieval" (e.g. "retrieve clear, bright and clean images"), [4] which retrieves images based on their color content, and [5] which retrieves MRI images based on the sizes and relative positions of multiple objects in the images. The QBIC system which we are developing allows images to be retrieved by a variety of image content descriptors including color, texture, and shape. The QBIC Project: Querying Images by Content Using Color, Texture, and Shape (IBM 0002390–404) at 4–5

Element	'004 Claim Recitation	Exemplary Citations in Reference
		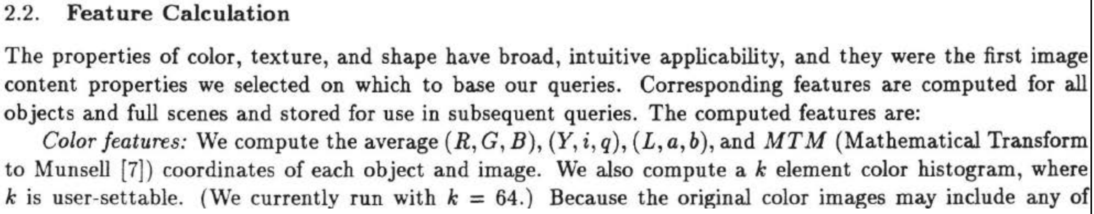 2.2. **Feature Calculation** The properties of color, texture, and shape have broad, intuitive applicability, and they were the first image content properties we selected on which to base our queries. Corresponding features are computed for all objects and full scenes and stored for use in subsequent queries. The computed features are: *Color features:* We compute the average (R,G,B), (Y,i,q), (L,a,b), and MTM (Mathematical Transform to Munsell [7]) coordinates of each object and image. We also compute a k element color histogram, where k is user-settable. (We currently run with $k = 64$.) Because the original color images may include any of 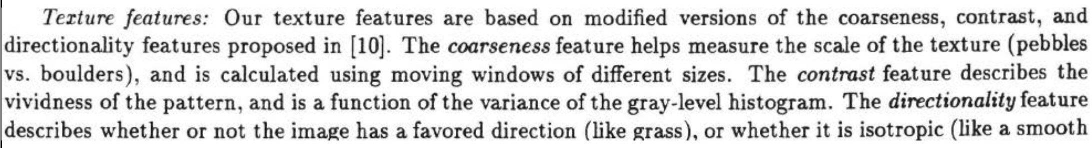 *Texture features:* Our texture features are based on modified versions of the coarseness, contrast, and directionality features proposed in [10]. The *coarseness* feature helps measure the scale of the texture (pebbles vs. boulders), and is calculated using moving windows of different sizes. The *contrast* feature describes the vividness of the pattern, and is a function of the variance of the gray-level histogram. The *directionality* feature describes whether or not the image has a favored direction (like grass), or whether it is isotropic (like a smooth 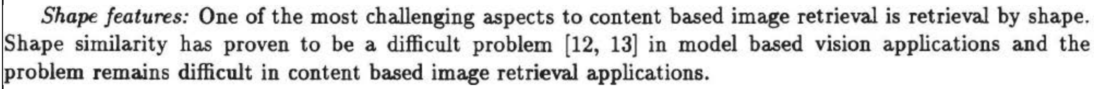 *Shape features:* One of the most challenging aspects to content based image retrieval is retrieval by shape. Shape similarity has proven to be a difficult problem [12, 13] in model based vision applications and the problem remains difficult in content based image retrieval applications. 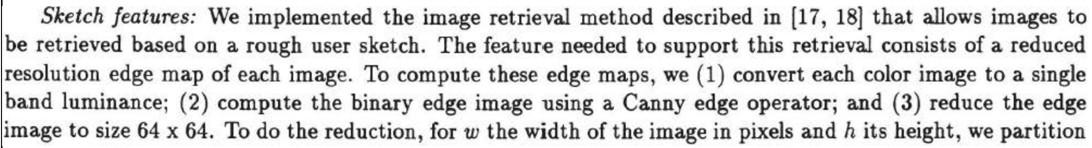 *Sketch features:* We implemented the image retrieval method described in [17, 18] that allows images to be retrieved based on a rough user sketch. The feature needed to support this retrieval consists of a reduced resolution edge map of each image. To compute these edge maps, we (1) convert each color image to a single band luminance; (2) compute the binary edge image using a Canny edge operator; and (3) reduce the edge image to size 64 x 64. To do the reduction, for w the width of the image in pixels and h its height, we partition Query by Image and Video Content: The QBIC System (IBM 0000893–902): "To achieve this functionality, QBIC has two main components: database population (the process of creating an image database) and database query. During the population, images and videos are processed to extract features describing their content-colors, textures, shapes, and camera and object motion-and the features are stored in a database. During the query, the user composes a query graphically. Features are generated from the graphical query and then input to a matching engine that finds images or videos from the database with similar features. Figure 2 shows the system architecture." *Id*. at p. 3. "For both population and query, the QBIC data model has • still images or scenes (full images) that contain objects (subsets of an image), and • video shots

Element	'004 Claim Recitation	Exemplary Citations in Reference
		that consist of sets of contiguous frames and contain motion objects. For still images, the QBIC data model distinguishes between 'scenes' (or images) and 'objects.' A scene is an image or single representative frame of video. An object is a part of a scene-for example, the fox in Figure 3-or a moving entity in a video. For still image database population, features are extracted from images and objects and stored in a database as shown in the top left part of Figure 2. Videos are broken into clips called shots. Representative frames, or r-frames, are generated for each extracted shot. R-frames are treated as still images, and features are extracted and stored in the database. Further processing of shots generates motion objects-for example, a car moving across the screen. Queries are allowed on objects ('Find images with a red, round object'), scenes ('Find images that have approximately 30-percent red and IS-percent blue colors'), shots ('Find all shots panning from left to right'), or any combination ('Find images that have 30 percent red and contain a blue textured object')." *Id*. at p. 3. "**DATABASE POPULATION** In still image database population, the images are reduced to a standard-sized icon called a thumbnail and annotated with any available text information. Object identification is an optional but key part of this step. It lets users manually, semiautomatically, or fully automatically identify interesting regions—which we call objects—in the images. Internally, each object is represented as a binary mask. There may be an arbitrary number of objects per image. Objects can overlap and can consist of multiple disconnected components like the set of dots on a polka-dot dress. Text, like 'baby on beach,' can be associated with an outlined object or with the scene as a whole. **Object-outlining tools** Ideally, object identification would be automatic, but this is generally difficult. The alternative—manual identification—is tedious and can inhibit query-by-content applications. As a result, we have devoted considerable effort to developing tools to aid in this step. In recent work, we have successfully used fully automatic unsupervised segmentation methods along with a

Element	'004 Claim Recitation	Exemplary Citations in Reference
		foreground/background model to identify objects in a restricted class of images. The images, typical of museums and retail catalogs, have a small number of foreground objects on a generally separable background. Figure 4 shows example results. Even in this domain, robust algorithms are required because of the textured and variegated backgrounds. We also provide semiautomatic tools for identifying objects. One is an enhanced flood-fill technique. Flood-fill methods, found in most photo-editing programs, start from a single object pixel and repeatedly add adjacent pixels whose values are within some given threshold of the original pixel. Selecting the threshold, which must change from image to image and object to object, is tedious. We automatically calculate a dynamic threshold by having the user click on background as well as object points. For reasonably uniform objects that are distinct from the background, this operation allows fast object identification without manually adjusting a threshold. The example in Figure 3 shows an object, a fox, identified by using only a few clicks. We designed another outlining tool to help users track object edges. This tool takes a user-drawn curve and automatically aligns it with nearby image edges. Based on the 'snakes' concept developed in recent computer vision research, the tool finds the curve that maximizes the image gradient magnitude along the curve. The spline snake formulation we use allows for smooth solutions to the resulting nonlinear minimization problem. The computation is done at interactive speeds so that, as the user draws a curve, it is 'rubber-banded' to lie along object boundaries. . . . **Video data** For video data, database population has three major components: • shot detection, • representative frame creation for each shot, and • derivation of a layered representation of coherently moving structures/ objects. . . . Once the shot boundaries have been detected, each shot is represented using an r-frame. R-frames are used for several purposes. First, during database population, r-frames are treated as still images in which objects can be identified by using the previously described methods." *Id*. at pp. 3–6.

Element	'004 Claim Recitation	Exemplary Citations in Reference
		Image, Audio, and Video Extenders Administration and Programming (IBM 000001–617): "Your applications can even search for images by content. Imagine an application that uses visual examples to search for images. With such an application, users could select an example image and have the application find other images that have colors or textures similar to those in the example. With DB2 extenders' Query by Image Content (QBIC) capability, you can create applications that search for images in this visual way." *Id*. at p. 4. "A QBIC catalog is a set of files that hold data about the visual features of images. The Image Extender uses this data to search for images by content. You create a QBIC catalog for each column of images in a user table that you want to make available for searching by content. When you create a QBIC catalog you identify the features for which you want the Image Extender to analyze, store, and later query data. You can also add or drop features from a QBIC catalog after the catalog is created. A QBIC catalog can hold data for the following image features: **Average color** The sum of the color values for all pixels in an image divided by the number of pixels in the image. (A pixel is the smallest element of an image that can be assigned color and intensity.) For example, if 50% of an image consists of blue pixels and the other 50% red pixels, the image has an average color value of purple. Average color is used to search for images that have a predominant color. If an image has a predominant color, the average color will be similar to the predominant color. **Histogram color** Measures the distribution of colors in an image against a spectrum of 64 colors. For each of the 64 colors, histogram color identifies the percentage of pixels in an image that have that color. For example, the histogram color of an image might be 40% white pixels, 50% blue, and 10% red; none of the pixels in the image have any of the remaining colors in the histogram spectrum. Histogram

Element	'004 Claim Recitation	Exemplary Citations in Reference
		color is used to search for images that have a variety of colors. **Positional color** The average color value for the pixels in a specified area in an image. For example, the upper right-hand corner of an image might show a bright yellow sun; the positional color of this area of the image is bright yellow. Positional color is used to search for images that have a predominant color in a particular area. **Texture** Measures the coarseness, contrast, and directionality of an image. Coarseness indicates the size of repeating items in an image (for example, pebbles versus boulders). Contrast identifies the brightness variations in an image (light versus dark). Directionality indicates whether a direction predominates in an image (as in the vertical direction of a picket fence) or does not predominate (as in an image of sand). Texture is used to search for images that have a particular pattern. To make an image available for searching by content, you catalog the image. When you catalog an image, the Image Extender analyzes the image, by computing the feature values for the image, and stores the values in a QBIC catalog. When you search for an image by content, your query identifies one or more features for the search (such as average color), a source for each feature (such as an example image), and a target set of cataloged images. The Image Extender computes the feature value of the source and compares it to the cataloged feature values for the target images. It then computes a score that indicates how similar the feature values of the target images are to the source. You can have the Image Extender return the images whose features are most similar to the source. The Image Extender will return the handle of each image and the image score. You can also have the Image Extender return only the score of a single image." *Id.* at pp. 20–21. QBIC Demo (IBM000747)

Element	'004 Claim Recitation	Exemplary Citations in Reference
		 QBIC(TM) -- IBM's Query By Image Content Goto: How to Use QBIC IBM Image Search Demo Stop QBIC Http Server Official QBIC QBIC allows you to make queries of image databases large and small based on visual characteristics -- properties such as color percentages, color layout and textures occurring in the images. The queries match image colors, textures and their positions without having to describe them with in words. QBIC's content-based queries can also be combined with keywords to provide a powerful multimedia search and retrieval tool. As image collections grow larger and become more prolific, QBIC is the best tool to more efficiently search and navigate through them. In this demo , we use QBIC to index about 470 color images based on their color and texture information. You can evaluate the QBIC Search capability by following the link located at left called IBM Image Search Demo. For license and business related questins, please contact Ted Loewenberg by phone 408-927-1202, or by email tedl@almaden.ibm.com. For technical questions, please contact Xiaoming Zhu by phone 408-927-1427 or by email QBIC Demo (IBM000747) 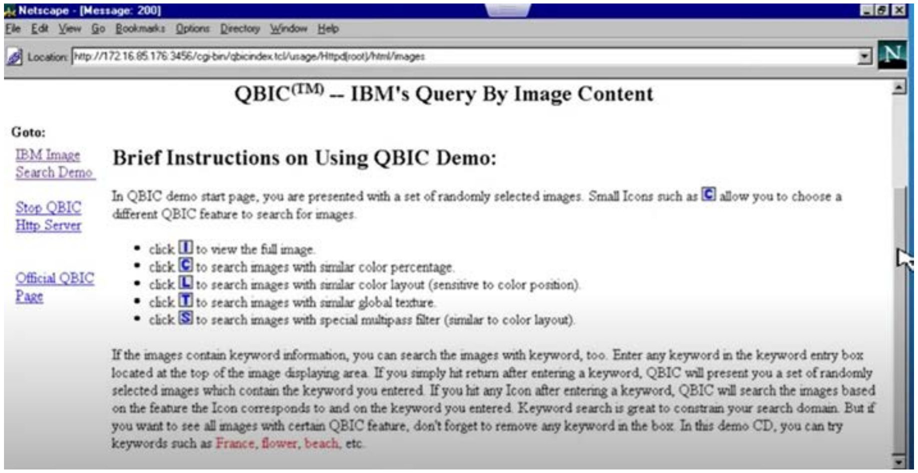 QBIC(TM) -- IBM's Query By Image Content Goto: IBM Image Search Demo Stop QBIC Http Server Official QBIC Page **Brief Instructions on Using QBIC Demo:** In QBIC demo start page, you are presented with a set of randomly selected images. Small Icons such as C allow you to choose a different QBIC feature to search for images. - click I to view the full image. - click C to search images with similar color percentage. - click L to search images with similar color layout (sensitive to color position). - click T to search images with similar global texture. - click S to search images with special multipass filter (similar to color layout). If the images contain keyword information, you can search the images with keyword, too. Enter any keyword in the keyword entry box located at the top of the image displaying area. If you simply hit return after entering a keyword, QBIC will present you a set of randomly selected images which contain the keyword you entered. If you hit any Icon after entering a keyword, QBIC will search the images based on the feature the Icon corresponds to and on the keyword you entered. Keyword search is great to constrain your search domain. But if you want to see all images with certain QBIC feature, don't forget to remove any keyword in the box. In this demo CD, you can try keywords such as France, flower, beach, etc.

Nantworks, LLC v. Bank of America Corporation
Central District of California, No. 2:20-cv-7872
Invalidity Chart for U.S. Patent 9,324,004 based on the Query by Image and Video Content ("QBIC") System

Element	'004 Claim Recitation	Exemplary Citations in Reference
		 http://172.16.85.176:3456/cgi-bin/ibm.tcl/search?fname=Color+Layout&france77%2Ejpg.x=0 Flickner Depo at 88-89

Element	'004 Claim Recitation	Exemplary Citations in Reference
		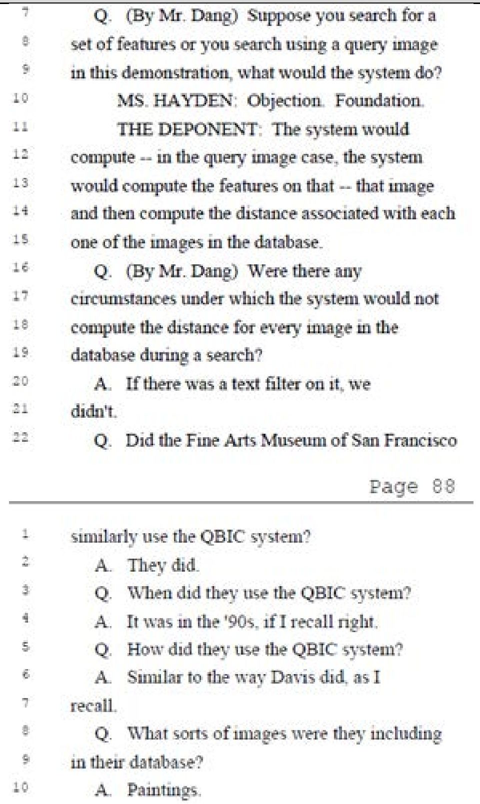 Q. (By Mr. Dang) Suppose you search for a set of features or you search using a query image in this demonstration, what would the system do? MS. HAYDEN: Objection. Foundation. THE DEPONENT: The system would compute -- in the query image case, the system would compute the features on that -- that image and then compute the distance associated with each one of the images in the database. Q. (By Mr. Dang) Were there any circumstances under which the system would not compute the distance for every image in the database during a search? A. If there was a text filter on it, we didn't. Q. Did the Fine Arts Museum of San Francisco Page 88 similarly use the QBIC system? A. They did. Q. When did they use the QBIC system? A. It was in the '90s, if I recall right. Q. How did they use the QBIC system? A. Similar to the way Davis did, as I recall. Q. What sorts of images were they including in their database? A. Paintings. Flickner Depo at 20

Element	'004 Claim Recitation	Exemplary Citations in Reference
		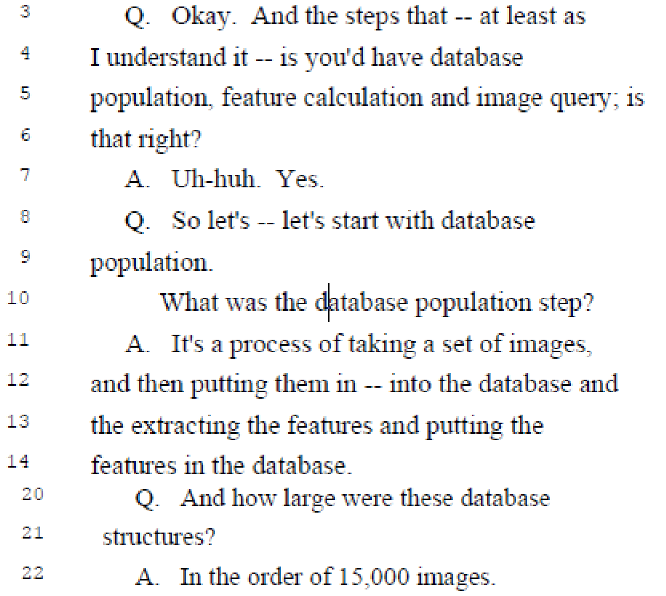 Q. Okay. And the steps that -- at least as I understand it -- is you'd have database population, feature calculation and image query; is that right? A. Uh-huh. Yes. Q. So let's -- let's start with database population. What was the database population step? A. It's a process of taking a set of images, and then putting them in -- into the database and the extracting the features and putting the features in the database. Q. And how large were these database structures? A. In the order of 15,000 images. Flickner Depo at 25-26

Element	’004 Claim Recitation	Exemplary Citations in Reference
		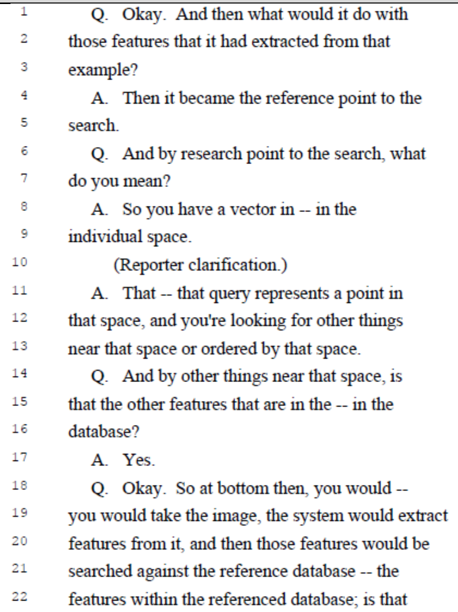 Q. Okay. And then what would it do with those features that it had extracted from that example? A. Then it became the reference point to the search. Q. And by research point to the search, what do you mean? A. So you have a vector in -- in the individual space. (Reporter clarification.) A. That -- that query represents a point in that space, and you're looking for other things near that space or ordered by that space. Q. And by other things near that space, is that the other features that are in the -- in the database? A. Yes. Q. Okay. So at bottom then, you would -- you would take the image, the system would extract features from it, and then those features would be searched against the reference database -- the features within the referenced database; is that 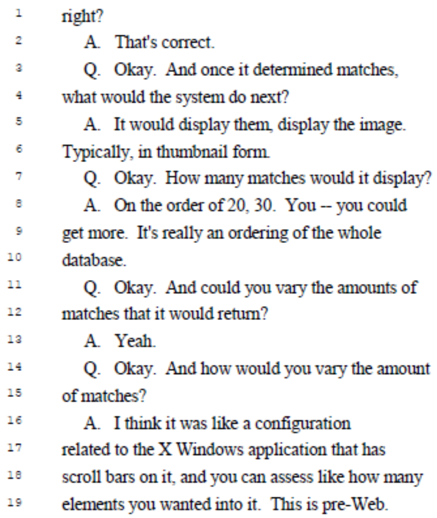 right? A. That's correct. Q. Okay. And once it determined matches, what would the system do next? A. It would display them, display the image. Typically, in thumbnail form. Q. Okay. How many matches would it display? A. On the order of 20, 30. You -- you could get more. It's really an ordering of the whole database. Q. Okay. And could you vary the amounts of matches that it would return? A. Yeah. Q. Okay. And how would you vary the amount of matches? A. I think it was like a configuration related to the X Windows application that has scroll bars on it, and you can assess like how many elements you wanted into it. This is pre-Web.

Element	'004 Claim Recitation	Exemplary Citations in Reference
		Flickner Depo at 38-41 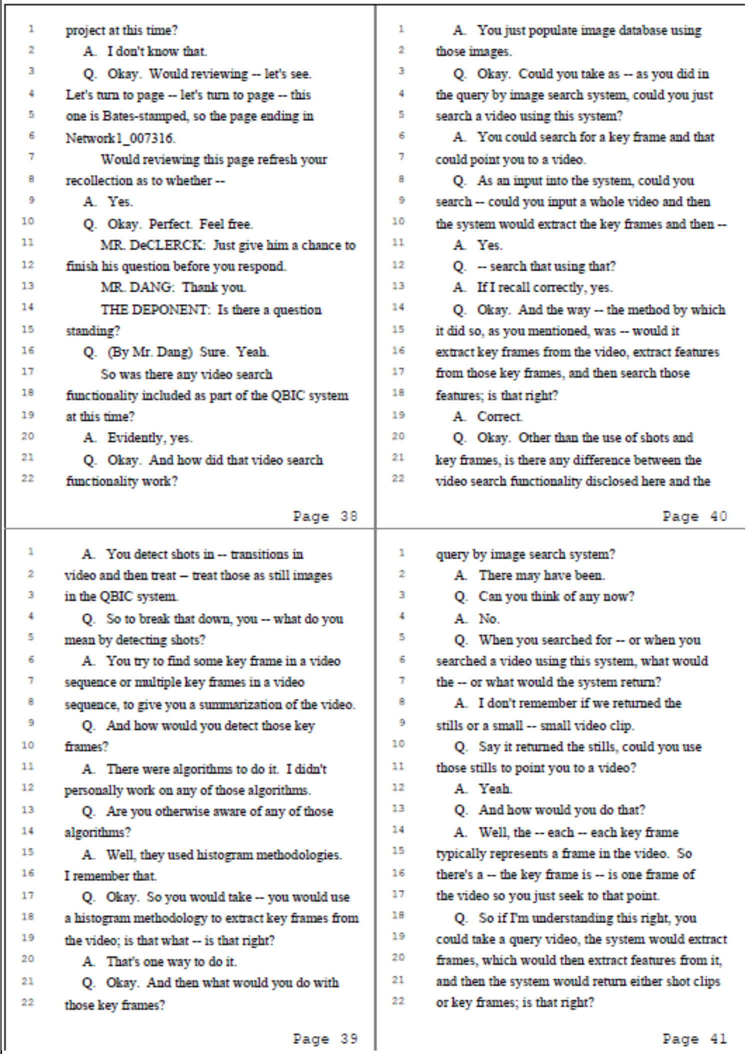 project at this time? A. I don't know that. Q. Okay. Would reviewing -- let's see. Let's turn to page -- let's turn to page -- this one is Bates-stamped, so the page ending in Network1_007316. Would reviewing this page refresh your recollection as to whether -- A. Yes. Q. Okay. Perfect. Feel free. MR. DeCLERCK: Just give him a chance to finish his question before you respond. MR. DANG: Thank you. THE DEPONENT: Is there a question standing? Q. (By Mr. Dang) Sure. Yeah. So was there any video search functionality included as part of the QBIC system at this time? A. Evidently, yes. Q. Okay. And how did that video search functionality work? Page 38 A. You detect shots in -- transitions in video and then treat -- treat those as still images in the QBIC system. Q. So to break that down, you -- what do you mean by detecting shots? A. You try to find some key frame in a video sequence or multiple key frames in a video sequence, to give you a summarization of the video. Q. And how would you detect those key frames? A. There were algorithms to do it. I didn't personally work on any of those algorithms. Q. Are you otherwise aware of any of those algorithms? A. Well, they used histogram methodologies. I remember that. Q. Okay. So you would take -- you would use a histogram methodology to extract key frames from the video; is that what -- is that right? A. That's one way to do it. Q. Okay. And then what would you do with those key frames? Page 39 A. You just populate image database using those images. Q. Okay. Could you take as -- as you did in the query by image search system, could you just search a video using this system? A. You could search for a key frame and that could point you to a video. Q. As an input into the system, could you search -- could you input a whole video and then the system would extract the key frames and then -- A. Yes. Q. -- search that using that? A. If I recall correctly, yes. Q. Okay. And the way -- the method by which it did so, as you mentioned, was -- would it extract key frames from the video, extract features from those key frames, and then search those features; is that right? A. Correct. Q. Okay. Other than the use of shots and key frames, is there any difference between the video search functionality disclosed here and the Page 40 query by image search system? A. There may have been. Q. Can you think of any now? A. No. Q. When you searched for -- or when you searched a video using this system, what would the -- or what would the system return? A. I don't remember if we returned the stills or a small -- small video clip. Q. Say it returned the stills, could you use those stills to point you to a video? A. Yeah. Q. And how would you do that? A. Well, the -- each -- each key frame typically represents a frame in the video. So there's a -- the key frame is -- is one frame of the video so you just seek to that point. Q. So if I'm understanding this right, you could take a query video, the system would extract frames, which would then extract features from it, and then the system would return either shot clips or key frames; is that right? Page 41

Element	’004 Claim Recitation	Exemplary Citations in Reference
		Flickner Depo at 40-42

		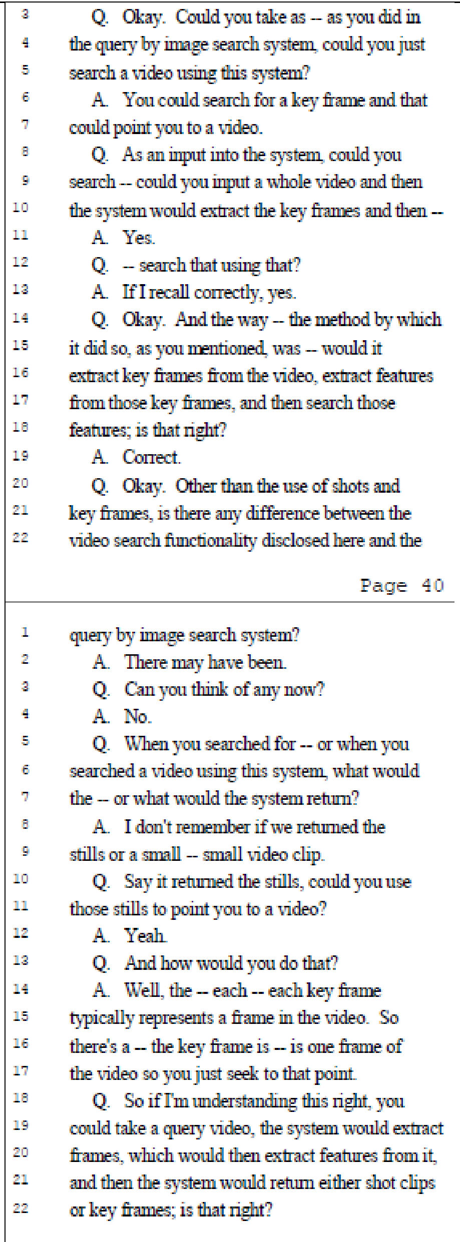 Q. Okay. Could you take as -- as you did in the query by image search system, could you just search a video using this system? A. You could search for a key frame and that could point you to a video. Q. As an input into the system, could you search -- could you input a whole video and then the system would extract the key frames and then -- A. Yes. Q. -- search that using that? A. If I recall correctly, yes. Q. Okay. And the way -- the method by which it did so, as you mentioned, was -- would it extract key frames from the video, extract features from those key frames, and then search those features; is that right? A. Correct. Q. Okay. Other than the use of shots and key frames, is there any difference between the video search functionality disclosed here and the Page 40 query by image search system? A. There may have been. Q. Can you think of any now? A. No. Q. When you searched for -- or when you searched a video using this system, what would the -- or what would the system return? A. I don't remember if we returned the stills or a small -- small video clip. Q. Say it returned the stills, could you use those stills to point you to a video? A. Yeah. Q. And how would you do that? A. Well, the -- each -- each key frame typically represents a frame in the video. So there's a -- the key frame is -- is one frame of the video so you just seek to that point. Q. So if I'm understanding this right, you could take a query video, the system would extract frames, which would then extract features from it, and then the system would return either shot clips or key frames; is that right?

Element	'004 Claim Recitation	Exemplary Citations in Reference
		 A. Right. Flickner Depo at 87. 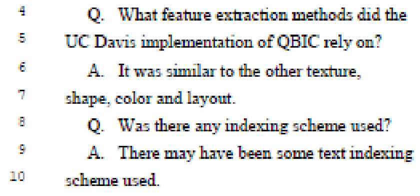 Q. What feature extraction methods did the UC Davis implementation of QBIC rely on? A. It was similar to the other texture, shape, color and layout. Q. Was there any indexing scheme used? A. There may have been some text indexing scheme used. Flickner 8/29/23 Depo at 33–37

Element	'004 Claim Recitation	Exemplary Citations in Reference
		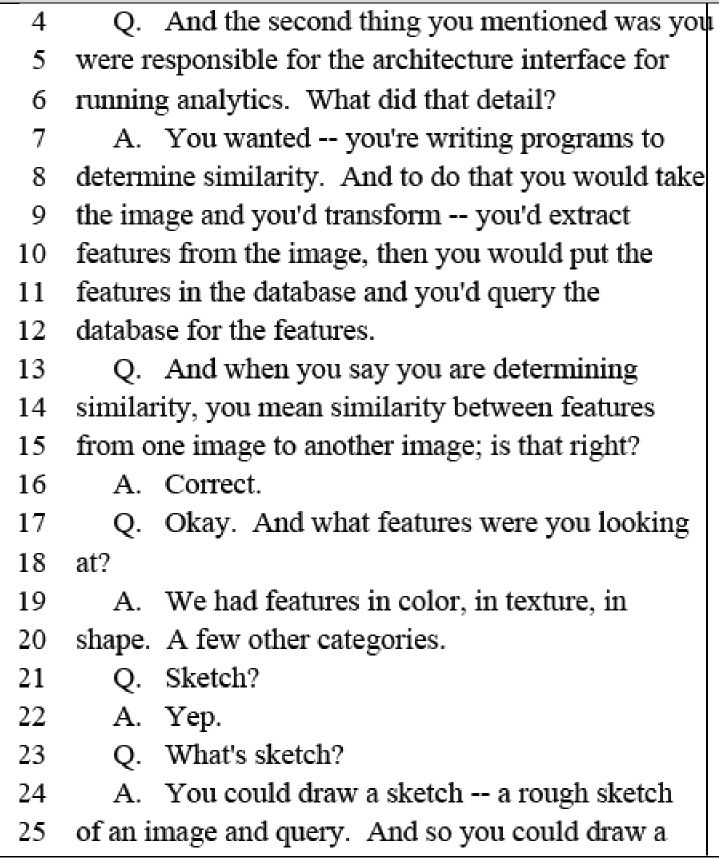 Q. And the second thing you mentioned was you were responsible for the architecture interface for running analytics. What did that detail? A. You wanted -- you're writing programs to determine similarity. And to do that you would take the image and you'd transform -- you'd extract features from the image, then you would put the features in the database and you'd query the database for the features. Q. And when you say you are determining similarity, you mean similarity between features from one image to another image; is that right? A. Correct. Q. Okay. And what features were you looking at? A. We had features in color, in texture, in shape. A few other categories. Q. Sketch? A. Yep. Q. What's sketch? A. You could draw a sketch -- a rough sketch of an image and query. And so you could draw a

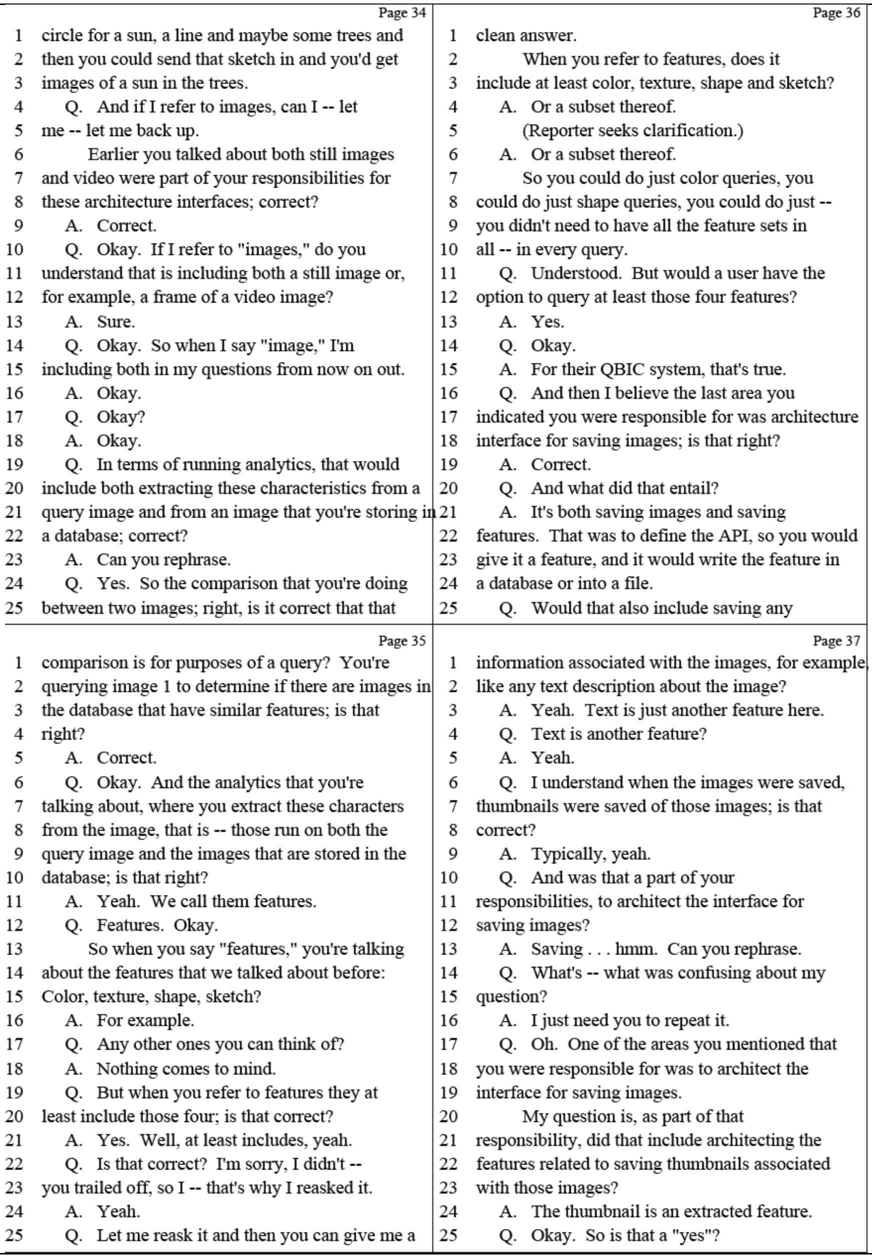

Page 34

circle for a sun, a line and maybe some trees and then you could send that sketch in and you'd get images of a sun in the trees.

Q. And if I refer to images, can I -- let me -- let me back up.

Earlier you talked about both still images and video were part of your responsibilities for these architecture interfaces; correct?

A. Correct.

Q. Okay. If I refer to "images," do you understand that is including both a still image or, for example, a frame of a video image?

A. Sure.

Q. Okay. So when I say "image," I'm including both in my questions from now on out.

A. Okay.

Q. Okay?

A. Okay.

Q. In terms of running analytics, that would include both extracting these characteristics from a query image and from an image that you're storing in a database; correct?

A. Can you rephrase.

Q. Yes. So the comparison that you're doing between two images; right, is it correct that that

Page 35

comparison is for purposes of a query? You're querying image 1 to determine if there are images in the database that have similar features; is that right?

A. Correct.

Q. Okay. And the analytics that you're talking about, where you extract these characters from the image, that is -- those run on both the query image and the images that are stored in the database; is that right?

A. Yeah. We call them features.

Q. Features. Okay.

So when you say "features," you're talking about the features that we talked about before: Color, texture, shape, sketch?

A. For example.

Q. Any other ones you can think of?

A. Nothing comes to mind.

Q. But when you refer to features they at least include those four; is that correct?

A. Yes. Well, at least includes, yeah.

Q. Is that correct? I'm sorry, I didn't -- you trailed off, so I -- that's why I reasked it.

A. Yeah.

Q. Let me reask it and then you can give me a

Page 36

clean answer.

When you refer to features, does it include at least color, texture, shape and sketch?

A. Or a subset thereof.

(Reporter seeks clarification.)

A. Or a subset thereof.

So you could do just color queries, you could do just shape queries, you could do just -- you didn't need to have all the feature sets in all -- in every query.

Q. Understood. But would a user have the option to query at least those four features?

A. Yes.

Q. Okay.

A. For their QBIC system, that's true.

Q. And then I believe the last area you indicated you were responsible for was architecture interface for saving images; is that right?

A. Correct.

Q. And what did that entail?

A. It's both saving images and saving features. That was to define the API, so you would give it a feature, and it would write the feature in a database or into a file.

Q. Would that also include saving any

Page 37

information associated with the images, for example, like any text description about the image?

A. Yeah. Text is just another feature here.

Q. Text is another feature?

A. Yeah.

Q. I understand when the images were saved, thumbnails were saved of those images; is that correct?

A. Typically, yeah.

Q. And was that a part of your responsibilities, to architect the interface for saving images?

A. Saving . . . hmm. Can you rephrase.

Q. What's -- what was confusing about my question?

A. I just need you to repeat it.

Q. Oh. One of the areas you mentioned that you were responsible for was to architect the interface for saving images.

My question is, as part of that responsibility, did that include architecting the features related to saving thumbnails associated with those images?

A. The thumbnail is an extracted feature.

Q. Okay. So is that a "yes"?

Element	'004 Claim Recitation	Exemplary Citations in Reference
		Id. at 77–78 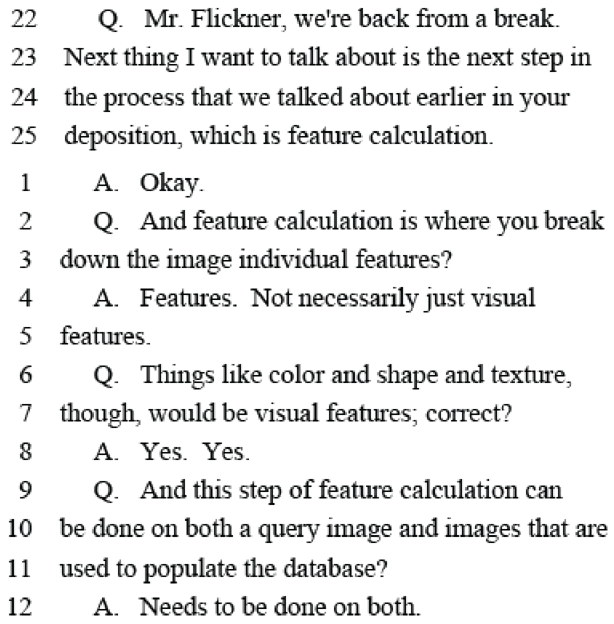 Q. Mr. Flickner, we're back from a break. Next thing I want to talk about is the next step in the process that we talked about earlier in your deposition, which is feature calculation. A. Okay. Q. And feature calculation is where you break down the image individual features? A. Features. Not necessarily just visual features. Q. Things like color and shape and texture, though, would be visual features; correct? A. Yes. Yes. Q. And this step of feature calculation can be done on both a query image and images that are used to populate the database? A. Needs to be done on both. *Id.* at 85–99 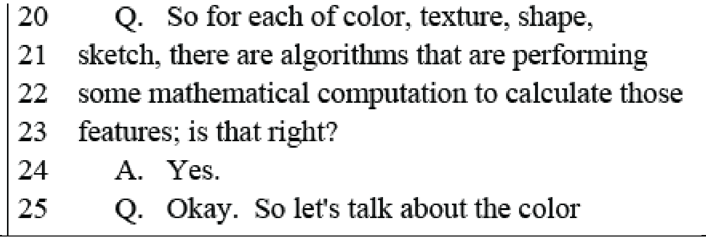 Q. So for each of color, texture, shape, sketch, there are algorithms that are performing some mathematical computation to calculate those features; is that right? A. Yes. Q. Okay. So let's talk about the color

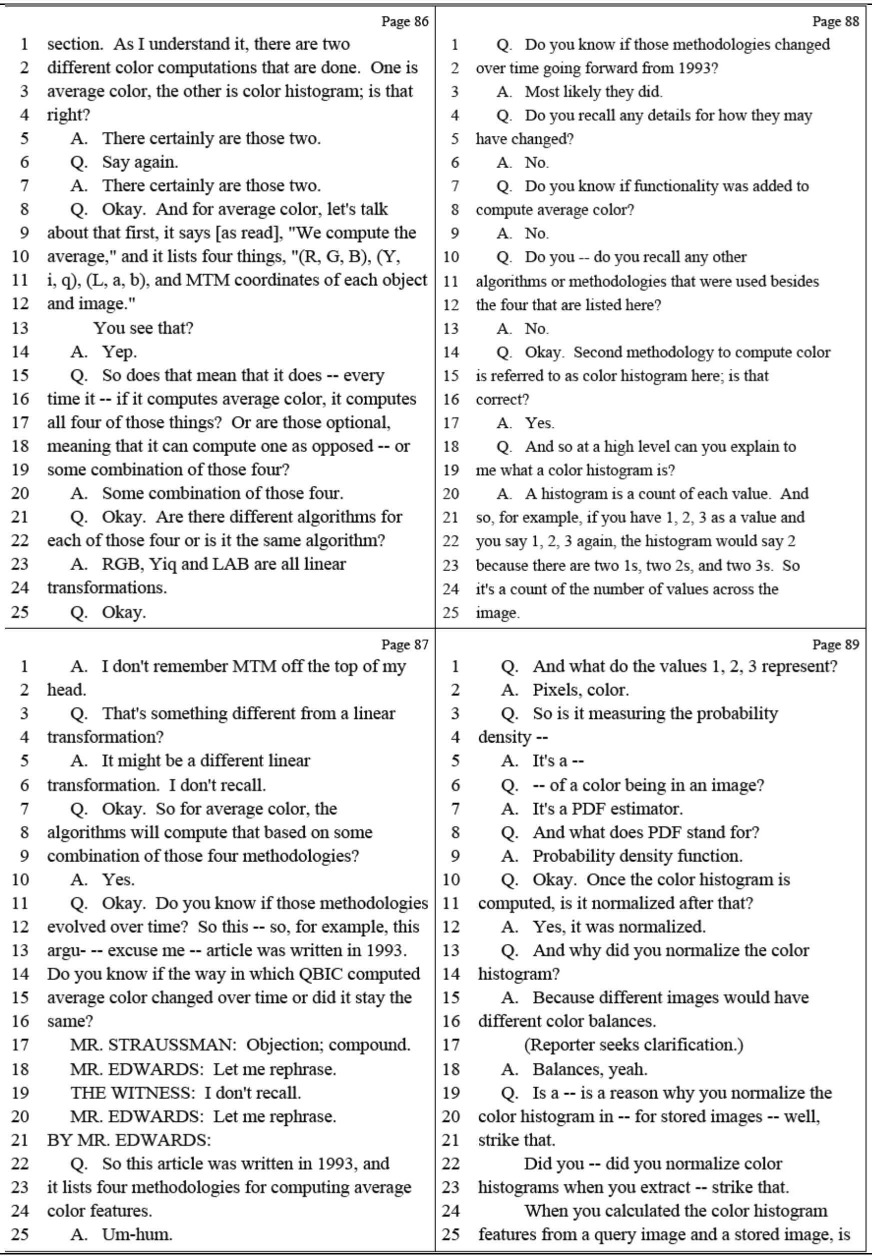

Page 86

section. As I understand it, there are two different color computations that are done. One is average color, the other is color histogram; is that right?
A. There certainly are those two.
Q. Say again.
A. There certainly are those two.
Q. Okay. And for average color, let's talk about that first, it says [as read], "We compute the average," and it lists four things, "(R, G, B), (Y, i, q), (L, a, b), and MTM coordinates of each object and image."
You see that?
A. Yep.
Q. So does that mean that it does -- every time it -- if it computes average color, it computes all four of those things? Or are those optional, meaning that it can compute one as opposed -- or some combination of those four?
A. Some combination of those four.
Q. Okay. Are there different algorithms for each of those four or is it the same algorithm?
A. RGB, Yiq and LAB are all linear transformations.
Q. Okay.

Page 87

A. I don't remember MTM off the top of my head.
Q. That's something different from a linear transformation?
A. It might be a different linear transformation. I don't recall.
Q. Okay. So for average color, the algorithms will compute that based on some combination of those four methodologies?
A. Yes.
Q. Okay. Do you know if those methodologies evolved over time? So this -- so, for example, this argu- -- excuse me -- article was written in 1993. Do you know if the way in which QBIC computed average color changed over time or did it stay the same?
MR. STRAUSSMAN: Objection; compound.
MR. EDWARDS: Let me rephrase.
THE WITNESS: I don't recall.
MR. EDWARDS: Let me rephrase.
BY MR. EDWARDS:
Q. So this article was written in 1993, and it lists four methodologies for computing average color features.
A. Um-hum.

Page 88

Q. Do you know if those methodologies changed over time going forward from 1993?
A. Most likely they did.
Q. Do you recall any details for how they may have changed?
A. No.
Q. Do you know if functionality was added to compute average color?
A. No.
Q. Do you -- do you recall any other algorithms or methodologies that were used besides the four that are listed here?
A. No.
Q. Okay. Second methodology to compute color is referred to as color histogram here; is that correct?
A. Yes.
Q. And so at a high level can you explain to me what a color histogram is?
A. A histogram is a count of each value. And so, for example, if you have 1, 2, 3 as a value and you say 1, 2, 3 again, the histogram would say 2 because there are two 1s, two 2s, and two 3s. So it's a count of the number of values across the image.

Page 89

Q. And what do the values 1, 2, 3 represent?
A. Pixels, color.
Q. So is it measuring the probability density --
A. It's a --
Q. -- of a color being in an image?
A. It's a PDF estimator.
Q. And what does PDF stand for?
A. Probability density function.
Q. Okay. Once the color histogram is computed, is it normalized after that?
A. Yes, it was normalized.
Q. And why did you normalize the color histogram?
A. Because different images would have different color balances.
(Reporter seeks clarification.)
A. Balances, yeah.
Q. Is a -- is a reason why you normalize the color histogram in -- for stored images -- well, strike that.
Did you -- did you normalize color histograms when you extract -- strike that.
When you calculated the color histogram features from a query image and a stored image, is

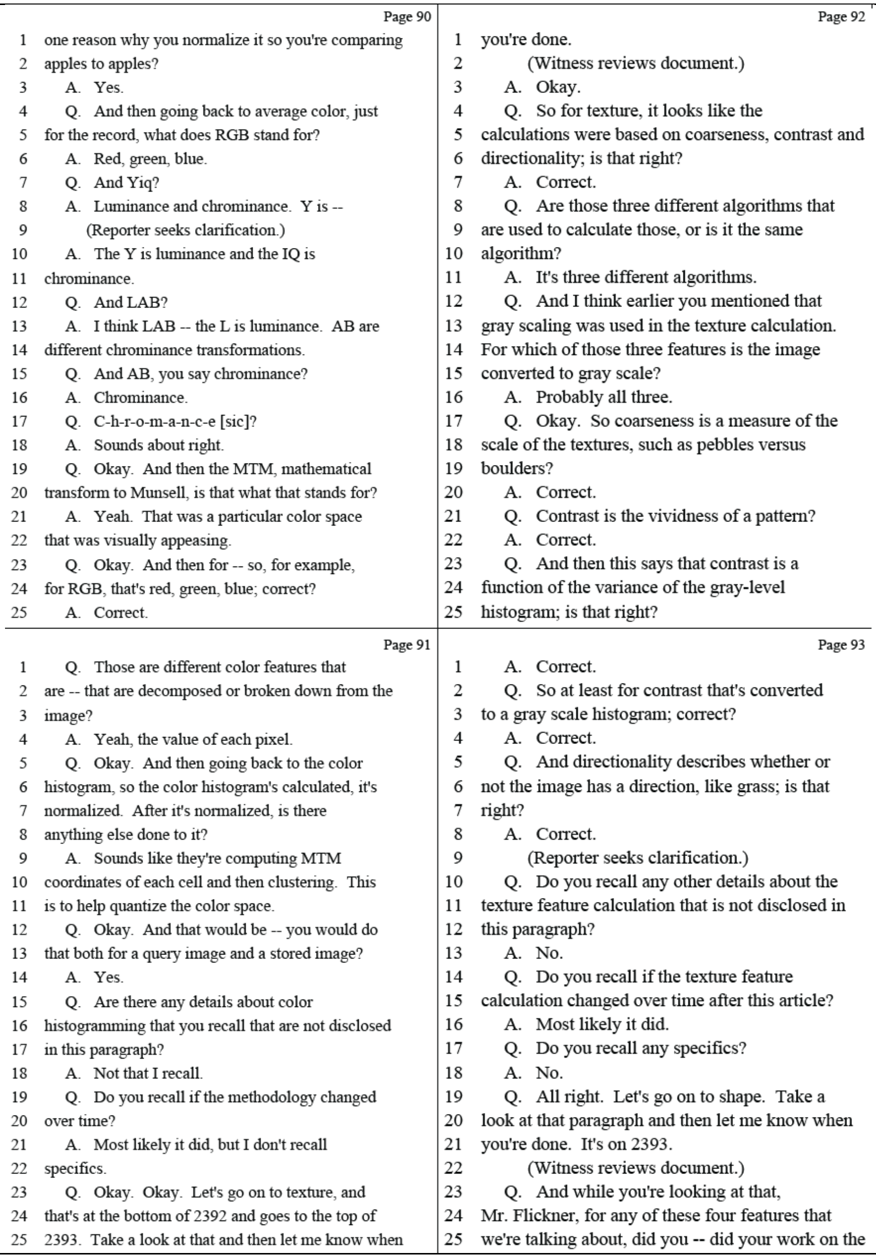

Page 90

one reason why you normalize it so you're comparing apples to apples?
A. Yes.
Q. And then going back to average color, just for the record, what does RGB stand for?
A. Red, green, blue.
Q. And Yiq?
A. Luminance and chrominance. Y is --
(Reporter seeks clarification.)
A. The Y is luminance and the IQ is chrominance.
Q. And LAB?
A. I think LAB -- the L is luminance. AB are different chrominance transformations.
Q. And AB, you say chrominance?
A. Chrominance.
Q. C-h-r-o-m-a-n-c-e [sic]?
A. Sounds about right.
Q. Okay. And then the MTM, mathematical transform to Munsell, is that what that stands for?
A. Yeah. That was a particular color space that was visually appeasing.
Q. Okay. And then for -- so, for example, for RGB, that's red, green, blue; correct?
A. Correct.

Page 91

Q. Those are different color features that are -- that are decomposed or broken down from the image?
A. Yeah, the value of each pixel.
Q. Okay. And then going back to the color histogram, so the color histogram's calculated, it's normalized. After it's normalized, is there anything else done to it?
A. Sounds like they're computing MTM coordinates of each cell and then clustering. This is to help quantize the color space.
Q. Okay. And that would be -- you would do that both for a query image and a stored image?
A. Yes.
Q. Are there any details about color histogramming that you recall that are not disclosed in this paragraph?
A. Not that I recall.
Q. Do you recall if the methodology changed over time?
A. Most likely it did, but I don't recall specifics.
Q. Okay. Okay. Let's go on to texture, and that's at the bottom of 2392 and goes to the top of 2393. Take a look at that and then let me know when

Page 92

you're done.
(Witness reviews document.)
A. Okay.
Q. So for texture, it looks like the calculations were based on coarseness, contrast and directionality; is that right?
A. Correct.
Q. Are those three different algorithms that are used to calculate those, or is it the same algorithm?
A. It's three different algorithms.
Q. And I think earlier you mentioned that gray scaling was used in the texture calculation. For which of those three features is the image converted to gray scale?
A. Probably all three.
Q. Okay. So coarseness is a measure of the scale of the textures, such as pebbles versus boulders?
A. Correct.
Q. Contrast is the vividness of a pattern?
A. Correct.
Q. And then this says that contrast is a function of the variance of the gray-level histogram; is that right?

Page 93

A. Correct.
Q. So at least for contrast that's converted to a gray scale histogram; correct?
A. Correct.
Q. And directionality describes whether or not the image has a direction, like grass; is that right?
A. Correct.
(Reporter seeks clarification.)
Q. Do you recall any other details about the texture feature calculation that is not disclosed in this paragraph?
A. No.
Q. Do you recall if the texture feature calculation changed over time after this article?
A. Most likely it did.
Q. Do you recall any specifics?
A. No.
Q. All right. Let's go on to shape. Take a look at that paragraph and then let me know when you're done. It's on 2393.
(Witness reviews document.)
Q. And while you're looking at that, Mr. Flickner, for any of these four features that we're talking about, did you -- did your work on the

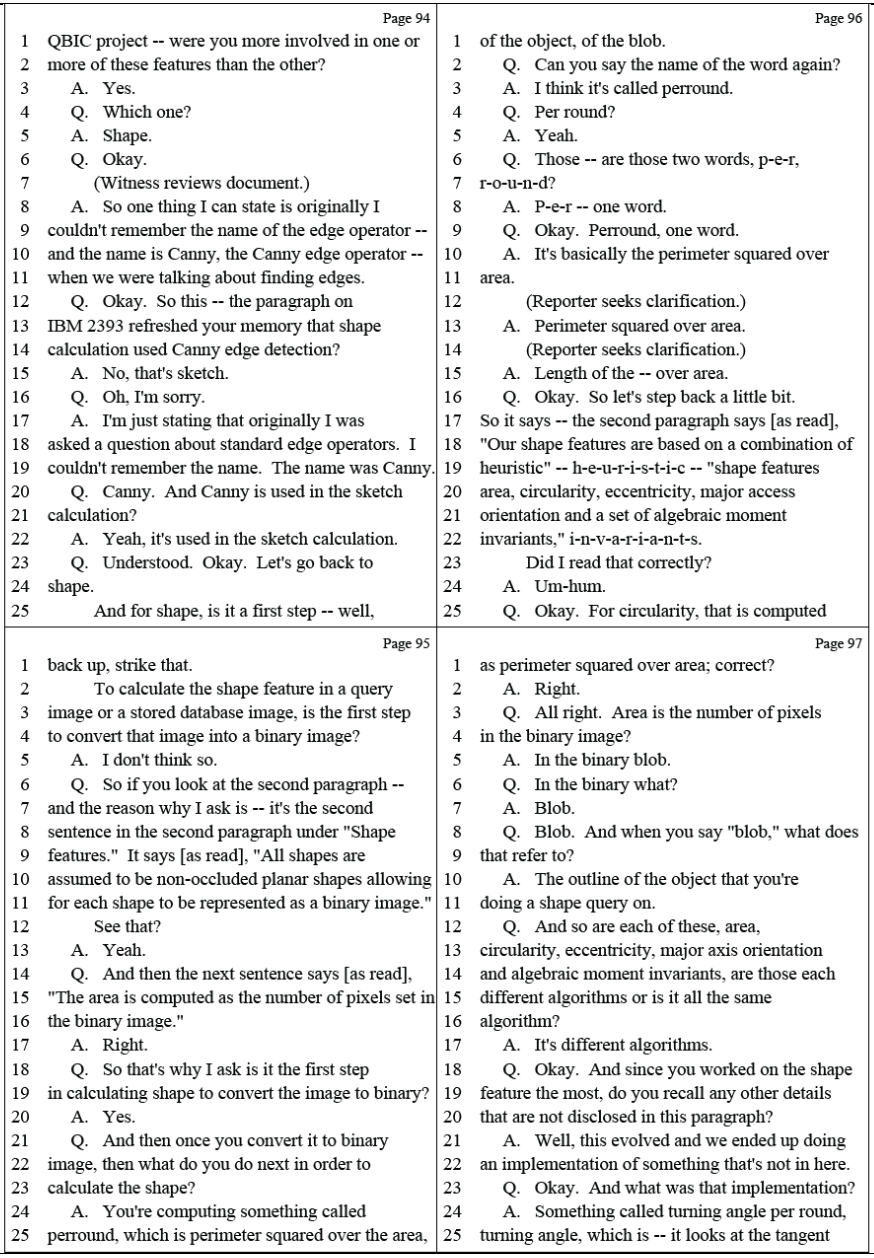

Page 94

QBIC project -- were you more involved in one or more of these features than the other?
A. Yes.
Q. Which one?
A. Shape.
Q. Okay.
(Witness reviews document.)
A. So one thing I can state is originally I couldn't remember the name of the edge operator -- and the name is Canny, the Canny edge operator -- when we were talking about finding edges.
Q. Okay. So this -- the paragraph on IBM 2393 refreshed your memory that shape calculation used Canny edge detection?
A. No, that's sketch.
Q. Oh, I'm sorry.
A. I'm just stating that originally I was asked a question about standard edge operators. I couldn't remember the name. The name was Canny.
Q. Canny. And Canny is used in the sketch calculation?
A. Yeah, it's used in the sketch calculation.
Q. Understood. Okay. Let's go back to shape.
And for shape, is it a first step -- well,

Page 95

back up, strike that.
To calculate the shape feature in a query image or a stored database image, is the first step to convert that image into a binary image?
A. I don't think so.
Q. So if you look at the second paragraph -- and the reason why I ask is -- it's the second sentence in the second paragraph under "Shape features." It says [as read], "All shapes are assumed to be non-occluded planar shapes allowing for each shape to be represented as a binary image."
See that?
A. Yeah.
Q. And then the next sentence says [as read], "The area is computed as the number of pixels set in the binary image."
A. Right.
Q. So that's why I ask is it the first step in calculating shape to convert the image to binary?
A. Yes.
Q. And then once you convert it to binary image, then what do you do next in order to calculate the shape?
A. You're computing something called perround, which is perimeter squared over the area,

Page 96

of the object, of the blob.
Q. Can you say the name of the word again?
A. I think it's called perround.
Q. Per round?
A. Yeah.
Q. Those -- are those two words, p-e-r, r-o-u-n-d?
A. P-e-r -- one word.
Q. Okay. Perround, one word.
A. It's basically the perimeter squared over area.
(Reporter seeks clarification.)
A. Perimeter squared over area.
(Reporter seeks clarification.)
A. Length of the -- over area.
Q. Okay. So let's step back a little bit. So it says -- the second paragraph says [as read], "Our shape features are based on a combination of heuristic" -- h-e-u-r-i-s-t-i-c -- "shape features area, circularity, eccentricity, major access orientation and a set of algebraic moment invariants," i-n-v-a-r-i-a-n-t-s.
Did I read that correctly?
A. Um-hum.
Q. Okay. For circularity, that is computed

Page 97

as perimeter squared over area; correct?
A. Right.
Q. All right. Area is the number of pixels in the binary image?
A. In the binary blob.
Q. In the binary what?
A. Blob.
Q. Blob. And when you say "blob," what does that refer to?
A. The outline of the object that you're doing a shape query on.
Q. And so are each of these, area, circularity, eccentricity, major axis orientation and algebraic moment invariants, are those each different algorithms or is it all the same algorithm?
A. It's different algorithms.
Q. Okay. And since you worked on the shape feature the most, do you recall any other details that are not disclosed in this paragraph?
A. Well, this evolved and we ended up doing an implementation of something that's not in here.
Q. Okay. And what was that implementation?
A. Something called turning angle per round, turning angle, which is -- it looks at the tangent

			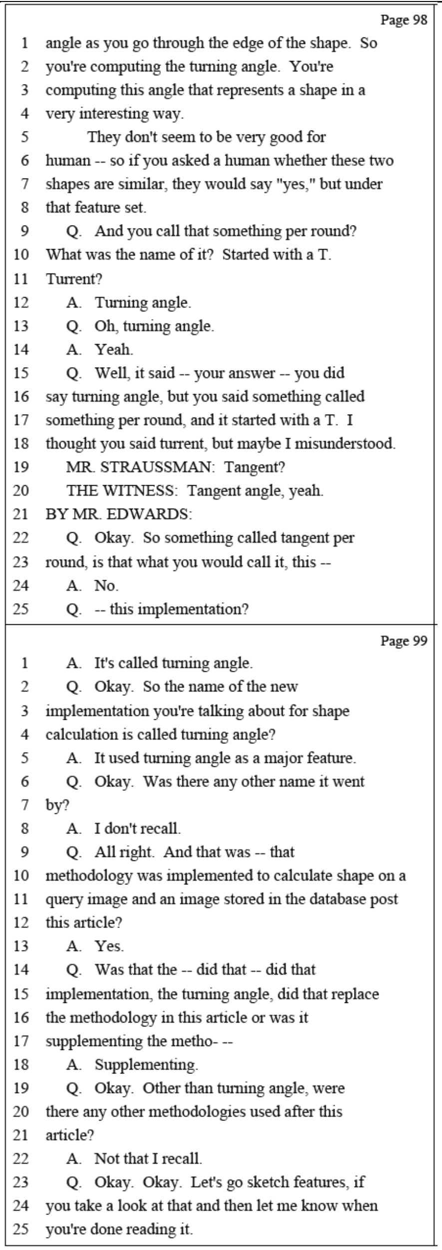 Page 98 angle as you go through the edge of the shape. So you're computing the turning angle. You're computing this angle that represents a shape in a very interesting way. They don't seem to be very good for human -- so if you asked a human whether these two shapes are similar, they would say "yes," but under that feature set. Q. And you call that something per round? What was the name of it? Started with a T. Turrent? A. Turning angle. Q. Oh, turning angle. A. Yeah. Q. Well, it said -- your answer -- you did say turning angle, but you said something called something per round, and it started with a T. I thought you said turrent, but maybe I misunderstood. MR. STRAUSSMAN: Tangent? THE WITNESS: Tangent angle, yeah. BY MR. EDWARDS: Q. Okay. So something called tangent per round, is that what you would call it, this -- A. No. Q. -- this implementation? Page 99 A. It's called turning angle. Q. Okay. So the name of the new implementation you're talking about for shape calculation is called turning angle? A. It used turning angle as a major feature. Q. Okay. Was there any other name it went by? A. I don't recall. Q. All right. And that was -- that methodology was implemented to calculate shape on a query image and an image stored in the database post this article? A. Yes. Q. Was that the -- did that -- did that implementation, the turning angle, did that replace the methodology in this article or was it supplementing the metho- -- A. Supplementing. Q. Okay. Other than turning angle, were there any other methodologies used after this article? A. Not that I recall. Q. Okay. Okay. Let's go sketch features, if you take a look at that and then let me know when you're done reading it.

Element	'004 Claim Recitation	Exemplary Citations in Reference
		Id. at 139–165 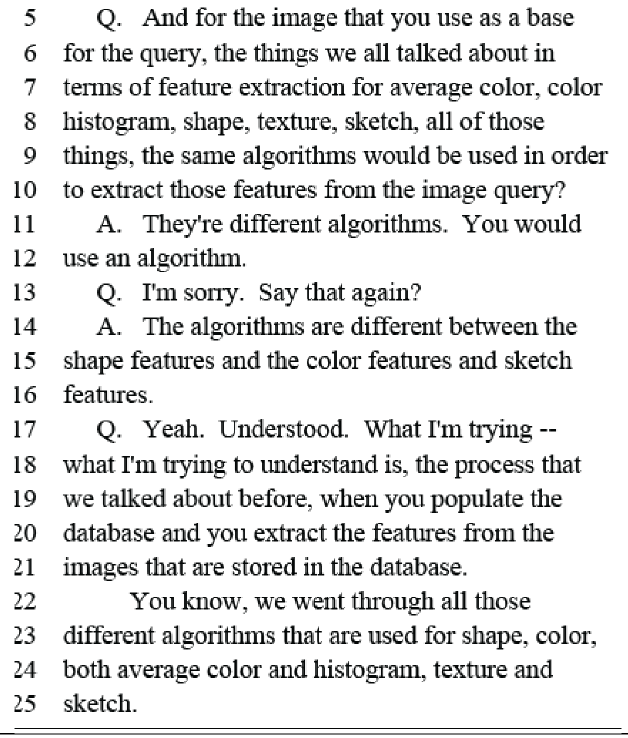 Q. And for the image that you use as a base for the query, the things we all talked about in terms of feature extraction for average color, color histogram, shape, texture, sketch, all of those things, the same algorithms would be used in order to extract those features from the image query? A. They're different algorithms. You would use an algorithm. Q. I'm sorry. Say that again? A. The algorithms are different between the shape features and the color features and sketch features. Q. Yeah. Understood. What I'm trying -- what I'm trying to understand is, the process that we talked about before, when you populate the database and you extract the features from the images that are stored in the database. You know, we went through all those different algorithms that are used for shape, color, both average color and histogram, texture and sketch.

Element	'004 Claim Recitation	Exemplary Citations in Reference
		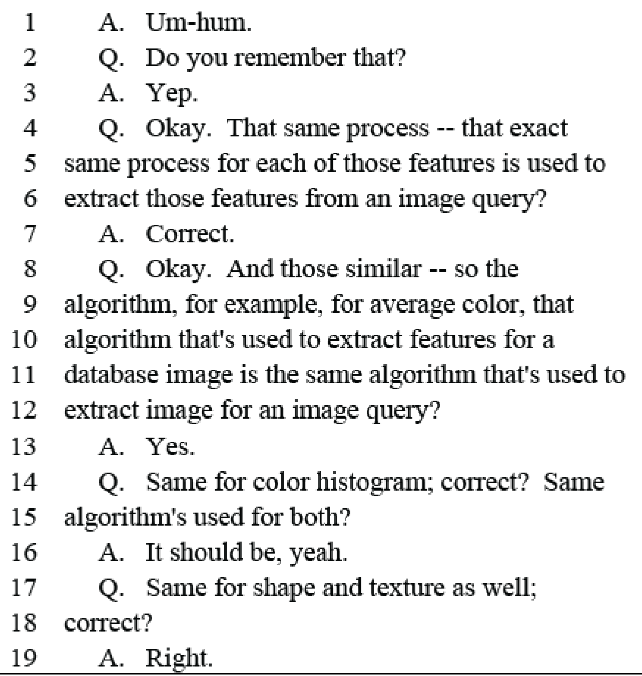 A. Um-hum. Q. Do you remember that? A. Yep. Q. Okay. That same process -- that exact same process for each of those features is used to extract those features from an image query? A. Correct. Q. Okay. And those similar -- so the algorithm, for example, for average color, that algorithm that's used to extract features for a database image is the same algorithm that's used to extract image for an image query? A. Yes. Q. Same for color histogram; correct? Same algorithm's used for both? A. It should be, yeah. Q. Same for shape and texture as well; correct? A. Right.

Element	'004 Claim Recitation	Exemplary Citations in Reference
		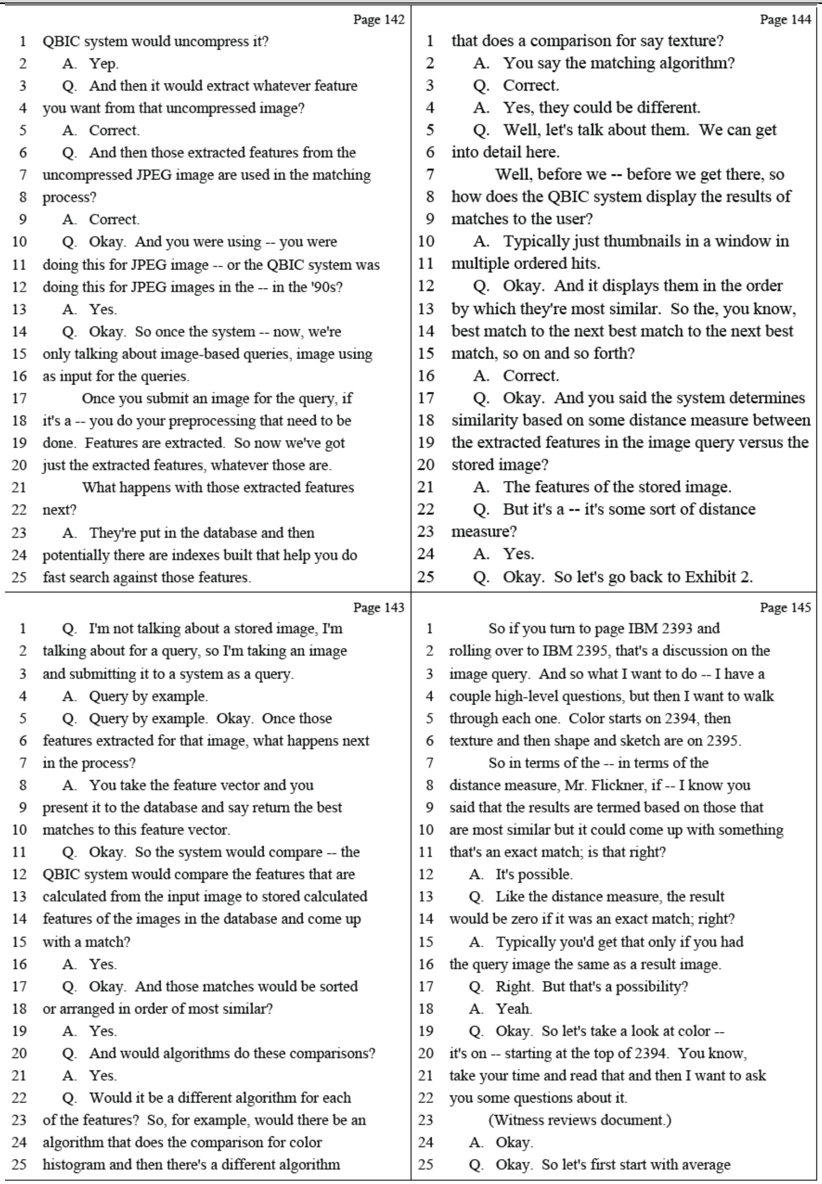 Page 142 QBIC system would uncompress it? A. Yep. Q. And then it would extract whatever feature you want from that uncompressed image? A. Correct. Q. And then those extracted features from the uncompressed JPEG image are used in the matching process? A. Correct. Q. Okay. And you were using -- you were doing this for JPEG image -- or the QBIC system was doing this for JPEG images in the -- in the '90s? A. Yes. Q. Okay. So once the system -- now, we're only talking about image-based queries, image using as input for the queries. Once you submit an image for the query, if it's a -- you do your preprocessing that need to be done. Features are extracted. So now we've got just the extracted features, whatever those are. What happens with those extracted features next? A. They're put in the database and then potentially there are indexes built that help you do fast search against those features. Page 143 Q. I'm not talking about a stored image, I'm talking about for a query, so I'm taking an image and submitting it to a system as a query. A. Query by example. Q. Query by example. Okay. Once those features extracted for that image, what happens next in the process? A. You take the feature vector and you present it to the database and say return the best matches to this feature vector. Q. Okay. So the system would compare -- the QBIC system would compare the features that are calculated from the input image to stored calculated features of the images in the database and come up with a match? A. Yes. Q. Okay. And those matches would be sorted or arranged in order of most similar? A. Yes. Q. And would algorithms do these comparisons? A. Yes. Q. Would it be a different algorithm for each of the features? So, for example, would there be an algorithm that does the comparison for color histogram and then there's a different algorithm Page 144 that does a comparison for say texture? A. You say the matching algorithm? Q. Correct. A. Yes, they could be different. Q. Well, let's talk about them. We can get into detail here. Well, before we -- before we get there, so how does the QBIC system display the results of matches to the user? A. Typically just thumbnails in a window in multiple ordered hits. Q. Okay. And it displays them in the order by which they're most similar. So the, you know, best match to the next best match to the next best match, so on and so forth? A. Correct. Q. Okay. And you said the system determines similarity based on some distance measure between the extracted features in the image query versus the stored image? A. The features of the stored image. Q. But it's a -- it's some sort of distance measure? A. Yes. Q. Okay. So let's go back to Exhibit 2. Page 145 So if you turn to page IBM 2393 and rolling over to IBM 2395, that's a discussion on the image query. And so what I want to do -- I have a couple high-level questions, but then I want to walk through each one. Color starts on 2394, then texture and then shape and sketch are on 2395. So in terms of the -- in terms of the distance measure, Mr. Flickner, if -- I know you said that the results are termed based on those that are most similar but it could come up with something that's an exact match; is that right? A. It's possible. Q. Like the distance measure, the result would be zero if it was an exact match; right? A. Typically you'd get that only if you had the query image the same as a result image. Q. Right. But that's a possibility? A. Yeah. Q. Okay. So let's take a look at color -- it's on -- starting at the top of 2394. You know, take your time and read that and then I want to ask you some questions about it. (Witness reviews document.) A. Okay. Q. Okay. So let's first start with average

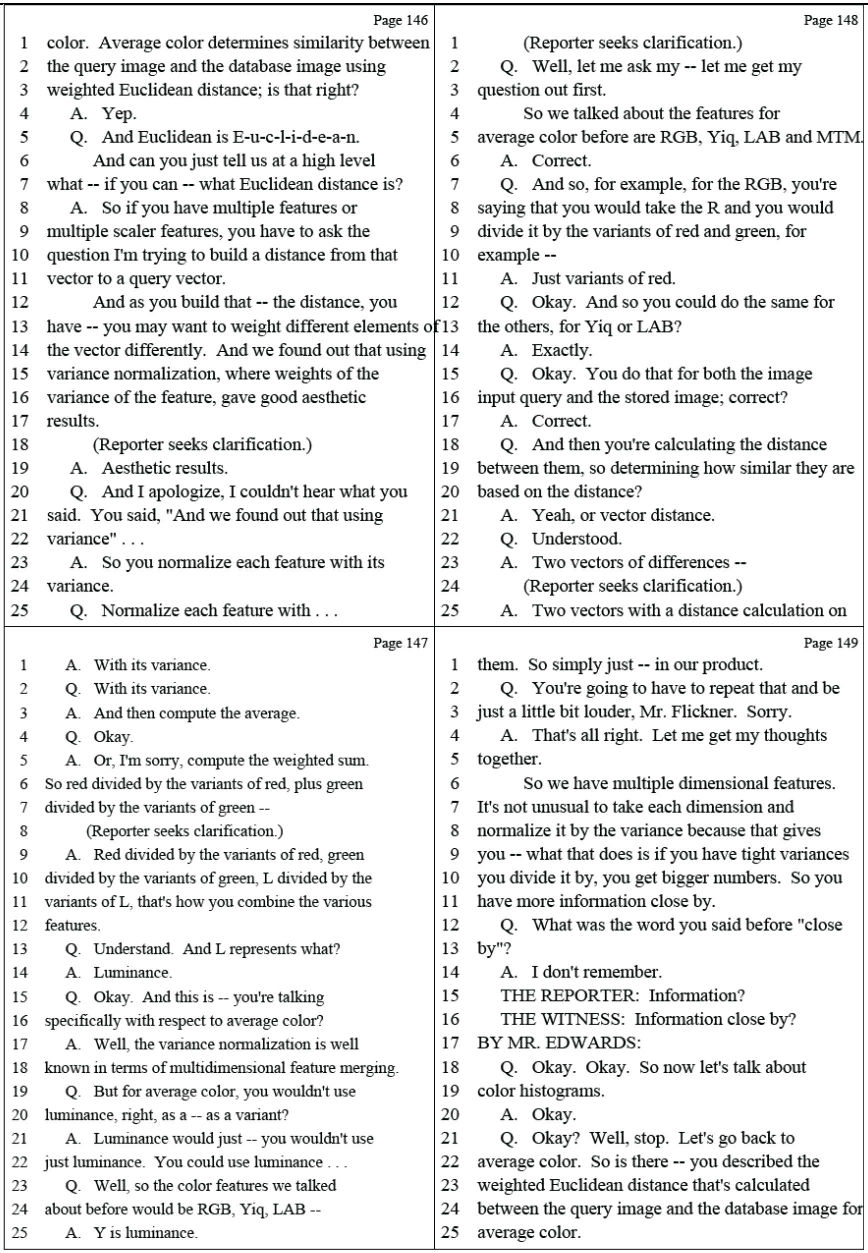

Page 146

color. Average color determines similarity between the query image and the database image using weighted Euclidean distance; is that right?

A. Yep.

Q. And Euclidean is E-u-c-l-i-d-e-a-n.

And can you just tell us at a high level what -- if you can -- what Euclidean distance is?

A. So if you have multiple features or multiple scaler features, you have to ask the question I'm trying to build a distance from that vector to a query vector.

And as you build that -- the distance, you have -- you may want to weight different elements of the vector differently. And we found out that using variance normalization, where weights of the variance of the feature, gave good aesthetic results.

(Reporter seeks clarification.)

A. Aesthetic results.

Q. And I apologize, I couldn't hear what you said. You said, "And we found out that using variance" . . .

A. So you normalize each feature with its variance.

Q. Normalize each feature with . . .

Page 147

A. With its variance.

Q. With its variance.

A. And then compute the average.

Q. Okay.

A. Or, I'm sorry, compute the weighted sum. So red divided by the variants of red, plus green divided by the variants of green --

(Reporter seeks clarification.)

A. Red divided by the variants of red, green divided by the variants of green, L divided by the variants of L, that's how you combine the various features.

Q. Understand. And L represents what?

A. Luminance.

Q. Okay. And this is -- you're talking specifically with respect to average color?

A. Well, the variance normalization is well known in terms of multidimensional feature merging.

Q. But for average color, you wouldn't use luminance, right, as a -- as a variant?

A. Luminance would just -- you wouldn't use just luminance. You could use luminance . . .

Q. Well, so the color features we talked about before would be RGB, Yiq, LAB --

A. Y is luminance.

Page 148

(Reporter seeks clarification.)

Q. Well, let me ask my -- let me get my question out first.

So we talked about the features for average color before are RGB, Yiq, LAB and MTM.

A. Correct.

Q. And so, for example, for the RGB, you're saying that you would take the R and you would divide it by the variants of red and green, for example --

A. Just variants of red.

Q. Okay. And so you could do the same for the others, for Yiq or LAB?

A. Exactly.

Q. Okay. You do that for both the image input query and the stored image; correct?

A. Correct.

Q. And then you're calculating the distance between them, so determining how similar they are based on the distance?

A. Yeah, or vector distance.

Q. Understood.

A. Two vectors of differences --

(Reporter seeks clarification.)

A. Two vectors with a distance calculation on

Page 149

them. So simply just -- in our product.

Q. You're going to have to repeat that and be just a little bit louder, Mr. Flickner. Sorry.

A. That's all right. Let me get my thoughts together.

So we have multiple dimensional features. It's not unusual to take each dimension and normalize it by the variance because that gives you -- what that does is if you have tight variances you divide it by, you get bigger numbers. So you have more information close by.

Q. What was the word you said before "close by"?

A. I don't remember.

THE REPORTER: Information?

THE WITNESS: Information close by?

BY MR. EDWARDS:

Q. Okay. Okay. So now let's talk about color histograms.

A. Okay.

Q. Okay? Well, stop. Let's go back to average color. So is there -- you described the weighted Euclidean distance that's calculated between the query image and the database image for average color.

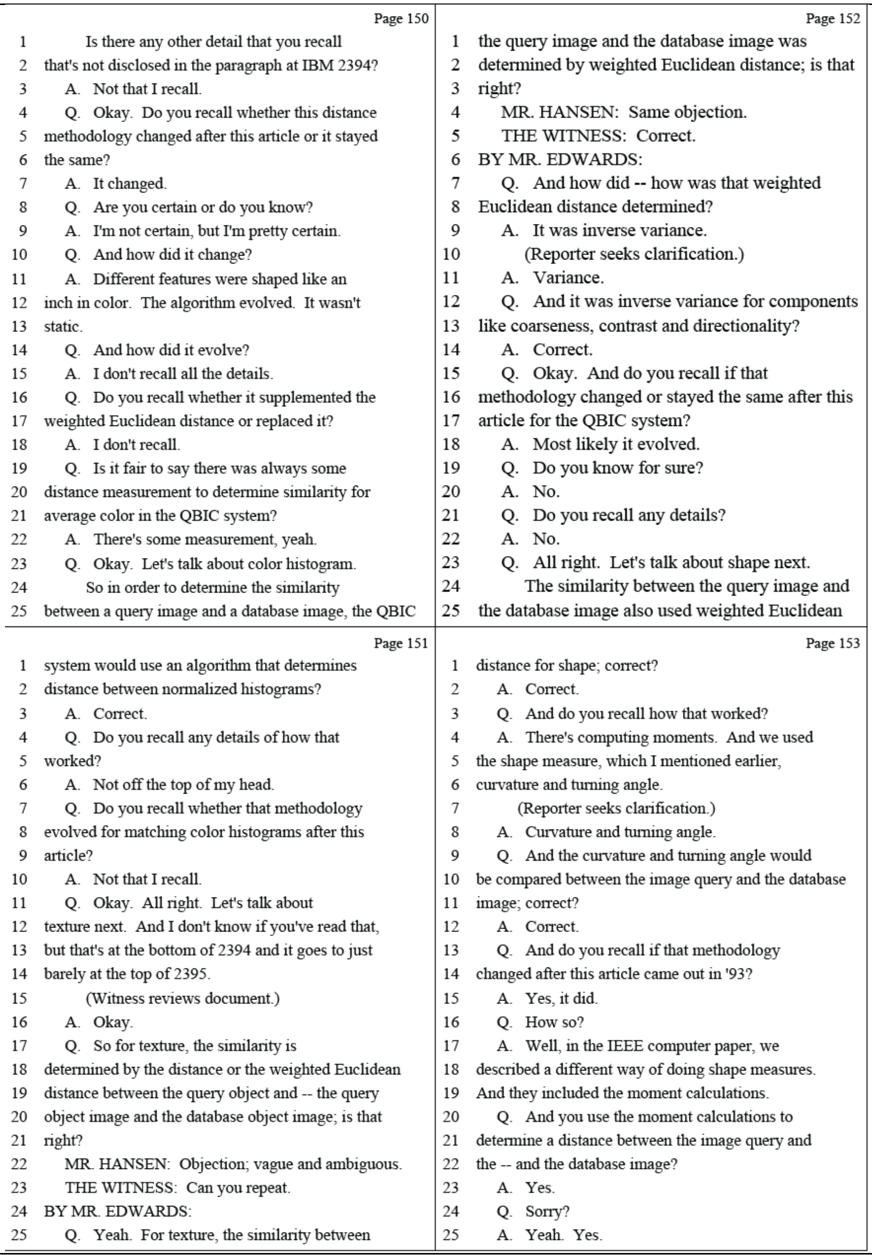

Page 150

Is there any other detail that you recall that's not disclosed in the paragraph at IBM 2394?

A. Not that I recall.

Q. Okay. Do you recall whether this distance methodology changed after this article or it stayed the same?

A. It changed.

Q. Are you certain or do you know?

A. I'm not certain, but I'm pretty certain.

Q. And how did it change?

A. Different features were shaped like an inch in color. The algorithm evolved. It wasn't static.

Q. And how did it evolve?

A. I don't recall all the details.

Q. Do you recall whether it supplemented the weighted Euclidean distance or replaced it?

A. I don't recall.

Q. Is it fair to say there was always some distance measurement to determine similarity for average color in the QBIC system?

A. There's some measurement, yeah.

Q. Okay. Let's talk about color histogram.

So in order to determine the similarity between a query image and a database image, the QBIC

Page 151

system would use an algorithm that determines distance between normalized histograms?

A. Correct.

Q. Do you recall any details of how that worked?

A. Not off the top of my head.

Q. Do you recall whether that methodology evolved for matching color histograms after this article?

A. Not that I recall.

Q. Okay. All right. Let's talk about texture next. And I don't know if you've read that, but that's at the bottom of 2394 and it goes to just barely at the top of 2395.

(Witness reviews document.)

A. Okay.

Q. So for texture, the similarity is determined by the distance or the weighted Euclidean distance between the query object and -- the query object image and the database object image; is that right?

MR. HANSEN: Objection; vague and ambiguous.

THE WITNESS: Can you repeat.

BY MR. EDWARDS:

Q. Yeah. For texture, the similarity between

Page 152

the query image and the database image was determined by weighted Euclidean distance; is that right?

MR. HANSEN: Same objection.

THE WITNESS: Correct.

BY MR. EDWARDS:

Q. And how did -- how was that weighted Euclidean distance determined?

A. It was inverse variance.

(Reporter seeks clarification.)

A. Variance.

Q. And it was inverse variance for components like coarseness, contrast and directionality?

A. Correct.

Q. Okay. And do you recall if that methodology changed or stayed the same after this article for the QBIC system?

A. Most likely it evolved.

Q. Do you know for sure?

A. No.

Q. Do you recall any details?

A. No.

Q. All right. Let's talk about shape next.

The similarity between the query image and the database image also used weighted Euclidean

Page 153

distance for shape; correct?

A. Correct.

Q. And do you recall how that worked?

A. There's computing moments. And we used the shape measure, which I mentioned earlier, curvature and turning angle.

(Reporter seeks clarification.)

A. Curvature and turning angle.

Q. And the curvature and turning angle would be compared between the image query and the database image; correct?

A. Correct.

Q. And do you recall if that methodology changed after this article came out in '93?

A. Yes, it did.

Q. How so?

A. Well, in the IEEE computer paper, we described a different way of doing shape measures. And they included the moment calculations.

Q. And you use the moment calculations to determine a distance between the image query and the -- and the database image?

A. Yes.

Q. Sorry?

A. Yeah. Yes.

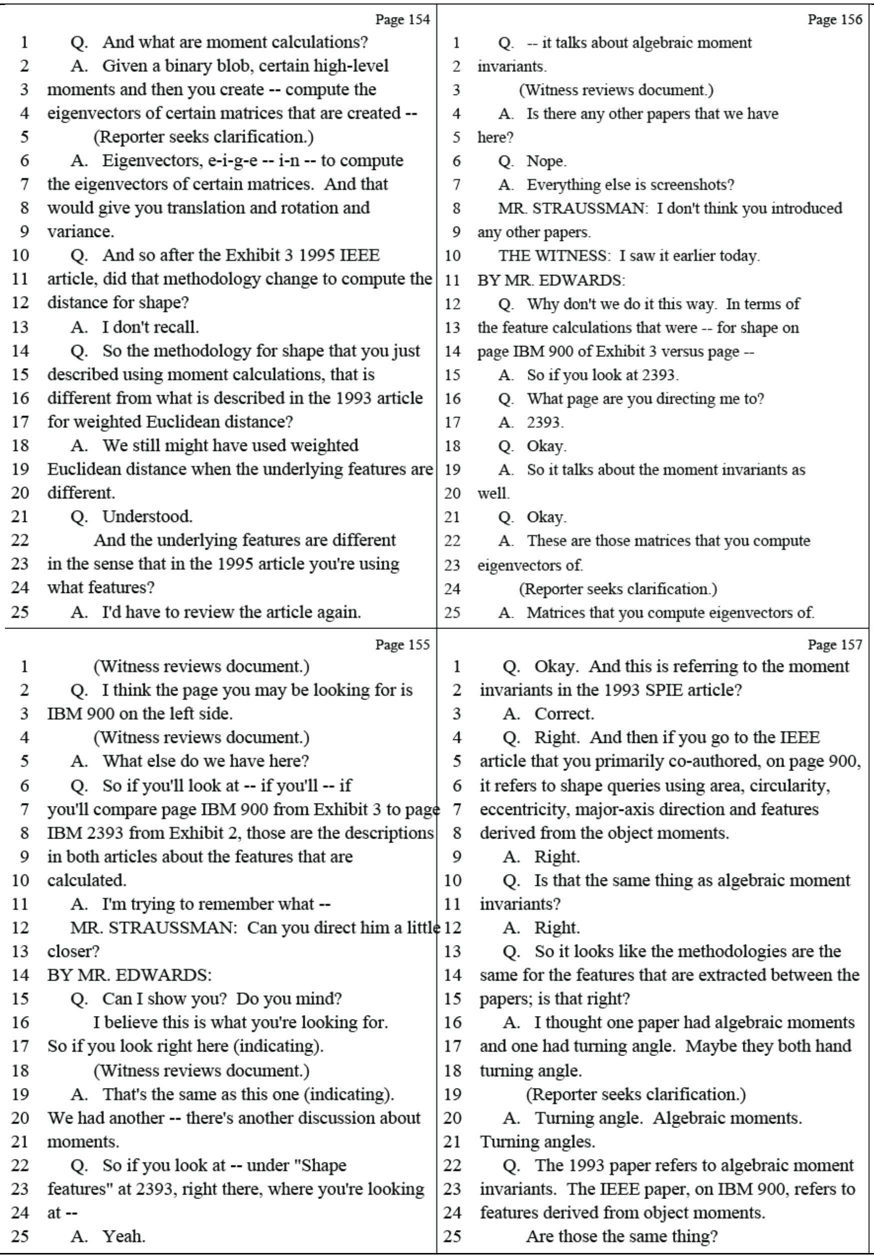

Page 154

Q. And what are moment calculations?
A. Given a binary blob, certain high-level moments and then you create -- compute the eigenvectors of certain matrices that are created --
(Reporter seeks clarification.)
A. Eigenvectors, e-i-g-e -- i-n -- to compute the eigenvectors of certain matrices. And that would give you translation and rotation and variance.
Q. And so after the Exhibit 3 1995 IEEE article, did that methodology change to compute the distance for shape?
A. I don't recall.
Q. So the methodology for shape that you just described using moment calculations, that is different from what is described in the 1993 article for weighted Euclidean distance?
A. We still might have used weighted Euclidean distance when the underlying features are different.
Q. Understood.
And the underlying features are different in the sense that in the 1995 article you're using what features?
A. I'd have to review the article again.

Page 155

(Witness reviews document.)
Q. I think the page you may be looking for is IBM 900 on the left side.
(Witness reviews document.)
A. What else do we have here?
Q. So if you'll look at -- if you'll -- if you'll compare page IBM 900 from Exhibit 3 to page IBM 2393 from Exhibit 2, those are the descriptions in both articles about the features that are calculated.
A. I'm trying to remember what --
MR. STRAUSSMAN: Can you direct him a little closer?
BY MR. EDWARDS:
Q. Can I show you? Do you mind?
I believe this is what you're looking for. So if you look right here (indicating).
(Witness reviews document.)
A. That's the same as this one (indicating). We had another -- there's another discussion about moments.
Q. So if you look at -- under "Shape features" at 2393, right there, where you're looking at --
A. Yeah.

Page 156

Q. -- it talks about algebraic moment invariants.
(Witness reviews document.)
A. Is there any other papers that we have here?
Q. Nope.
A. Everything else is screenshots?
MR. STRAUSSMAN: I don't think you introduced any other papers.
THE WITNESS: I saw it earlier today.
BY MR. EDWARDS:
Q. Why don't we do it this way. In terms of the feature calculations that were -- for shape on page IBM 900 of Exhibit 3 versus page --
A. So if you look at 2393.
Q. What page are you directing me to?
A. 2393.
Q. Okay.
A. So it talks about the moment invariants as well.
Q. Okay.
A. These are those matrices that you compute eigenvectors of.
(Reporter seeks clarification.)
A. Matrices that you compute eigenvectors of.

Page 157

Q. Okay. And this is referring to the moment invariants in the 1993 SPIE article?
A. Correct.
Q. Right. And then if you go to the IEEE article that you primarily co-authored, on page 900, it refers to shape queries using area, circularity, eccentricity, major-axis direction and features derived from the object moments.
A. Right.
Q. Is that the same thing as algebraic moment invariants?
A. Right.
Q. So it looks like the methodologies are the same for the features that are extracted between the papers; is that right?
A. I thought one paper had algebraic moments and one had turning angle. Maybe they both hand turning angle.
(Reporter seeks clarification.)
A. Turning angle. Algebraic moments. Turning angles.
Q. The 1993 paper refers to algebraic moment invariants. The IEEE paper, on IBM 900, refers to features derived from object moments.
Are those the same thing?

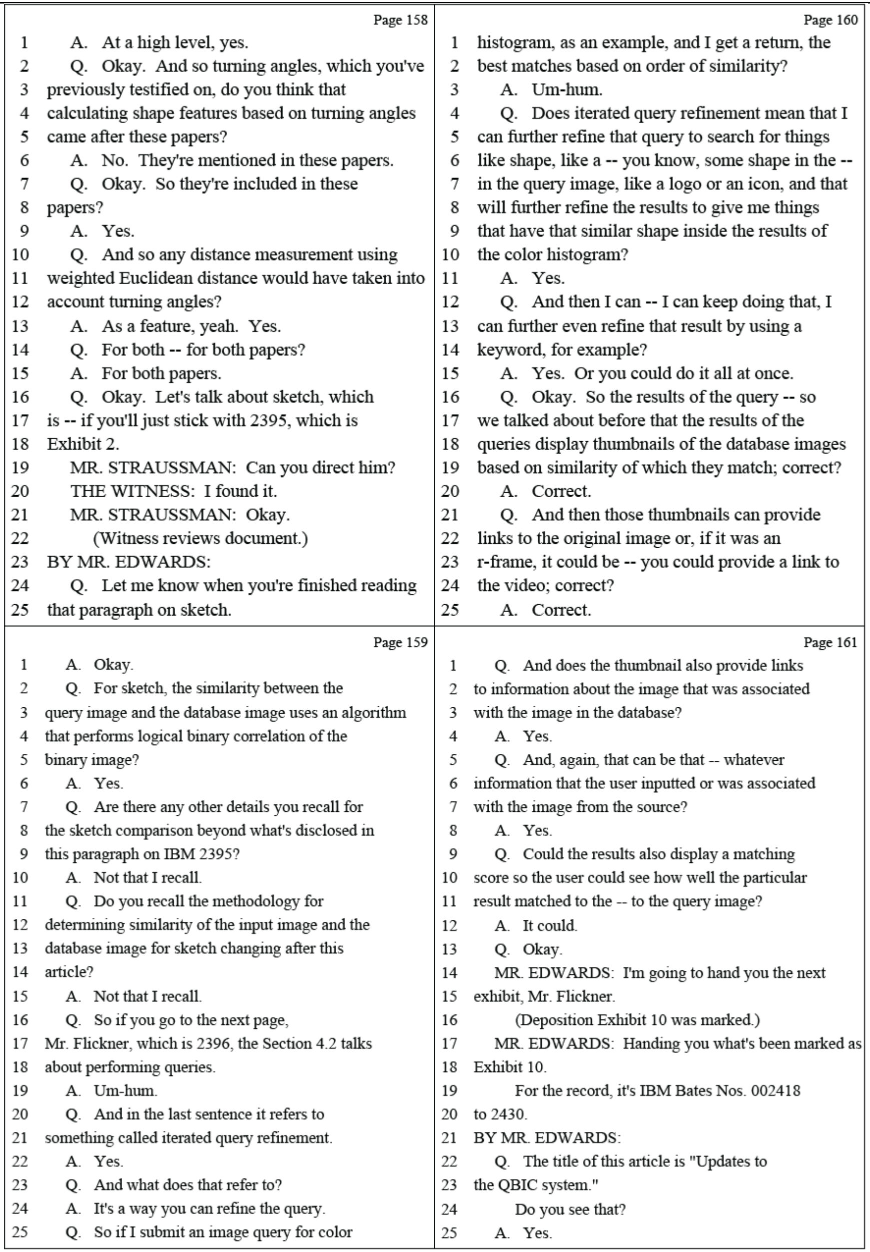

Page 158

A. At a high level, yes.
Q. Okay. And so turning angles, which you've previously testified on, do you think that calculating shape features based on turning angles came after these papers?
A. No. They're mentioned in these papers.
Q. Okay. So they're included in these papers?
A. Yes.
Q. And so any distance measurement using weighted Euclidean distance would have taken into account turning angles?
A. As a feature, yeah. Yes.
Q. For both -- for both papers?
A. For both papers.
Q. Okay. Let's talk about sketch, which is -- if you'll just stick with 2395, which is Exhibit 2.
MR. STRAUSSMAN: Can you direct him?
THE WITNESS: I found it.
MR. STRAUSSMAN: Okay.
(Witness reviews document.)
BY MR. EDWARDS:
Q. Let me know when you're finished reading that paragraph on sketch.

Page 159

A. Okay.
Q. For sketch, the similarity between the query image and the database image uses an algorithm that performs logical binary correlation of the binary image?
A. Yes.
Q. Are there any other details you recall for the sketch comparison beyond what's disclosed in this paragraph on IBM 2395?
A. Not that I recall.
Q. Do you recall the methodology for determining similarity of the input image and the database image for sketch changing after this article?
A. Not that I recall.
Q. So if you go to the next page, Mr. Flickner, which is 2396, the Section 4.2 talks about performing queries.
A. Um-hum.
Q. And in the last sentence it refers to something called iterated query refinement.
A. Yes.
Q. And what does that refer to?
A. It's a way you can refine the query.
Q. So if I submit an image query for color

Page 160

histogram, as an example, and I get a return, the best matches based on order of similarity?
A. Um-hum.
Q. Does iterated query refinement mean that I can further refine that query to search for things like shape, like a -- you know, some shape in the -- in the query image, like a logo or an icon, and that will further refine the results to give me things that have that similar shape inside the results of the color histogram?
A. Yes.
Q. And then I can -- I can keep doing that, I can further even refine that result by using a keyword, for example?
A. Yes. Or you could do it all at once.
Q. Okay. So the results of the query -- so we talked about before that the results of the queries display thumbnails of the database images based on similarity of which they match; correct?
A. Correct.
Q. And then those thumbnails can provide links to the original image or, if it was an r-frame, it could be -- you could provide a link to the video; correct?
A. Correct.

Page 161

Q. And does the thumbnail also provide links to information about the image that was associated with the image in the database?
A. Yes.
Q. And, again, that can be that -- whatever information that the user inputted or was associated with the image from the source?
A. Yes.
Q. Could the results also display a matching score so the user could see how well the particular result matched to the -- to the query image?
A. It could.
Q. Okay.
MR. EDWARDS: I'm going to hand you the next exhibit, Mr. Flickner.
(Deposition Exhibit 10 was marked.)
MR. EDWARDS: Handing you what's been marked as Exhibit 10.
For the record, it's IBM Bates Nos. 002418 to 2430.
BY MR. EDWARDS:
Q. The title of this article is "Updates to the QBIC system."
Do you see that?
A. Yes.

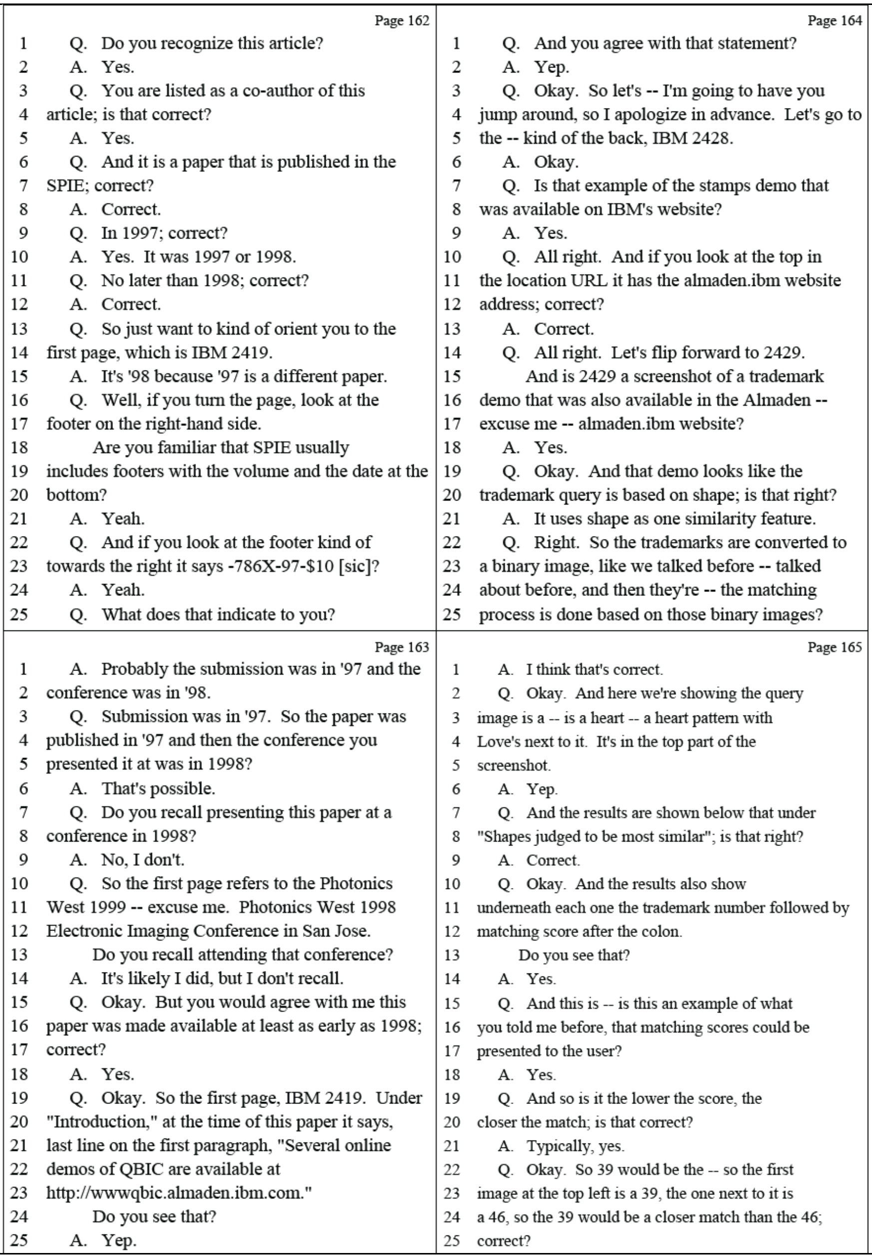

Page 162

Q. Do you recognize this article?
A. Yes.
Q. You are listed as a co-author of this article; is that correct?
A. Yes.
Q. And it is a paper that is published in the SPIE; correct?
A. Correct.
Q. In 1997; correct?
A. Yes. It was 1997 or 1998.
Q. No later than 1998; correct?
A. Correct.
Q. So just want to kind of orient you to the first page, which is IBM 2419.
A. It's '98 because '97 is a different paper.
Q. Well, if you turn the page, look at the footer on the right-hand side.
Are you familiar that SPIE usually includes footers with the volume and the date at the bottom?
A. Yeah.
Q. And if you look at the footer kind of towards the right it says -786X-97-$10 [sic]?
A. Yeah.
Q. What does that indicate to you?

Page 163

A. Probably the submission was in '97 and the conference was in '98.
Q. Submission was in '97. So the paper was published in '97 and then the conference you presented it at was in 1998?
A. That's possible.
Q. Do you recall presenting this paper at a conference in 1998?
A. No, I don't.
Q. So the first page refers to the Photonics West 1999 -- excuse me. Photonics West 1998 Electronic Imaging Conference in San Jose.
Do you recall attending that conference?
A. It's likely I did, but I don't recall.
Q. Okay. But you would agree with me this paper was made available at least as early as 1998; correct?
A. Yes.
Q. Okay. So the first page, IBM 2419. Under "Introduction," at the time of this paper it says, last line on the first paragraph, "Several online demos of QBIC are available at http://wwwqbic.almaden.ibm.com."
Do you see that?
A. Yep.

Page 164

Q. And you agree with that statement?
A. Yep.
Q. Okay. So let's -- I'm going to have you jump around, so I apologize in advance. Let's go to the -- kind of the back, IBM 2428.
A. Okay.
Q. Is that example of the stamps demo that was available on IBM's website?
A. Yes.
Q. All right. And if you look at the top in the location URL it has the almaden.ibm website address; correct?
A. Correct.
Q. All right. Let's flip forward to 2429.
And is 2429 a screenshot of a trademark demo that was also available in the Almaden -- excuse me -- almaden.ibm website?
A. Yes.
Q. Okay. And that demo looks like the trademark query is based on shape; is that right?
A. It uses shape as one similarity feature.
Q. Right. So the trademarks are converted to a binary image, like we talked before -- talked about before, and then they're -- the matching process is done based on those binary images?

Page 165

A. I think that's correct.
Q. Okay. And here we're showing the query image is a -- is a heart -- a heart pattern with Love's next to it. It's in the top part of the screenshot.
A. Yep.
Q. And the results are shown below that under "Shapes judged to be most similar"; is that right?
A. Correct.
Q. Okay. And the results also show underneath each one the trademark number followed by matching score after the colon.
Do you see that?
A. Yes.
Q. And this is -- is this an example of what you told me before, that matching scores could be presented to the user?
A. Yes.
Q. And so is it the lower the score, the closer the match; is that correct?
A. Typically, yes.
Q. Okay. So 39 would be the -- so the first image at the top left is a 39, the one next to it is a 46, so the 39 would be a closer match than the 46; correct?

Element	'004 Claim Recitation	Exemplary Citations in Reference
		Id. at 317–322 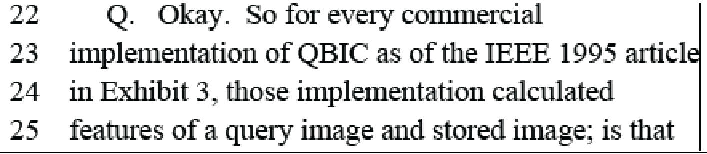 Q. Okay. So for every commercial implementation of QBIC as of the IEEE 1995 article in Exhibit 3, those implementation calculated features of a query image and stored image; is that

		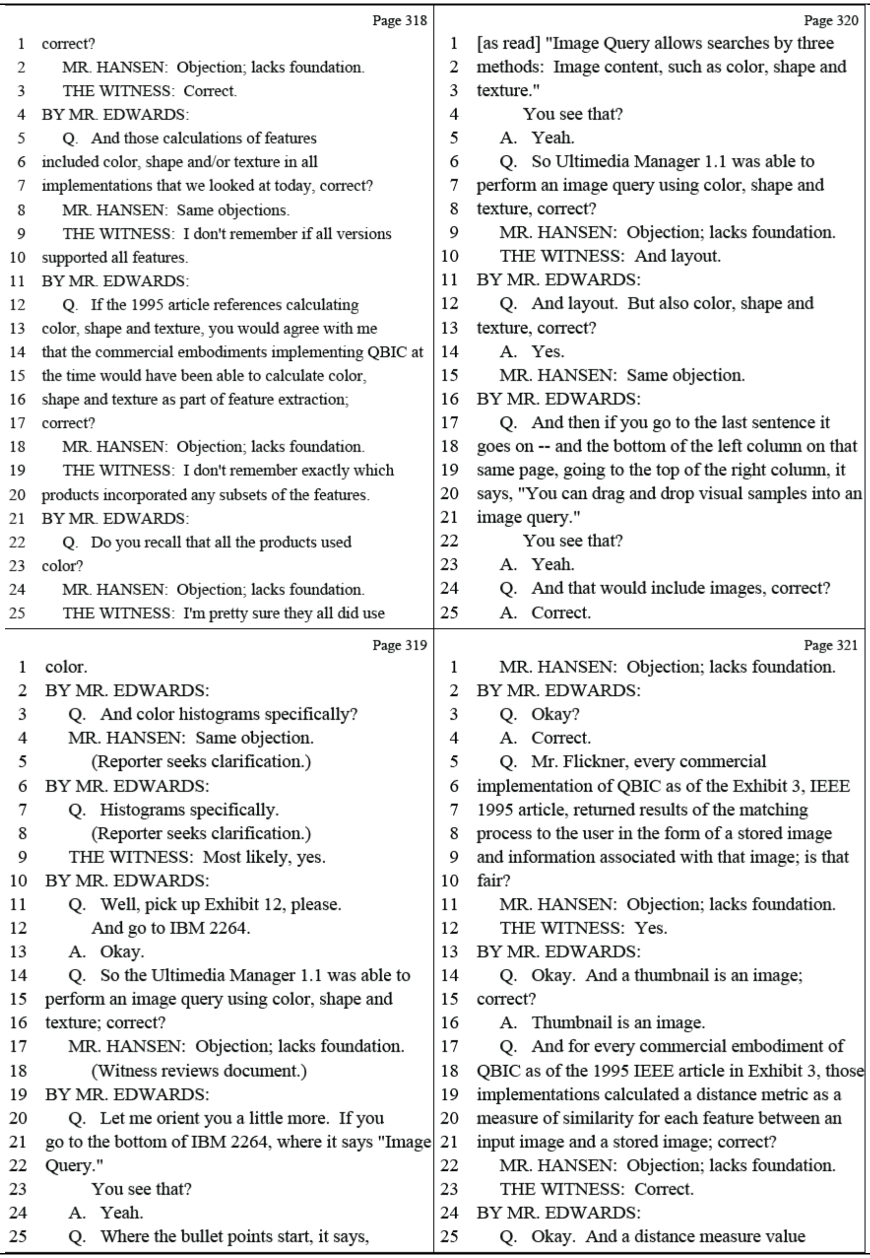

Page 318

correct?

MR. HANSEN: Objection; lacks foundation.

THE WITNESS: Correct.

BY MR. EDWARDS:

Q. And those calculations of features included color, shape and/or texture in all implementations that we looked at today, correct?

MR. HANSEN: Same objections.

THE WITNESS: I don't remember if all versions supported all features.

BY MR. EDWARDS:

Q. If the 1995 article references calculating color, shape and texture, you would agree with me that the commercial embodiments implementing QBIC at the time would have been able to calculate color, shape and texture as part of feature extraction; correct?

MR. HANSEN: Objection; lacks foundation.

THE WITNESS: I don't remember exactly which products incorporated any subsets of the features.

BY MR. EDWARDS:

Q. Do you recall that all the products used color?

MR. HANSEN: Objection; lacks foundation.

THE WITNESS: I'm pretty sure they all did use

Page 319

color.

BY MR. EDWARDS:

Q. And color histograms specifically?

MR. HANSEN: Same objection.

(Reporter seeks clarification.)

BY MR. EDWARDS:

Q. Histograms specifically.

(Reporter seeks clarification.)

THE WITNESS: Most likely, yes.

BY MR. EDWARDS:

Q. Well, pick up Exhibit 12, please.

And go to IBM 2264.

A. Okay.

Q. So the Ultimedia Manager 1.1 was able to perform an image query using color, shape and texture; correct?

MR. HANSEN: Objection; lacks foundation.

(Witness reviews document.)

BY MR. EDWARDS:

Q. Let me orient you a little more. If you go to the bottom of IBM 2264, where it says "Image Query."

You see that?

A. Yeah.

Q. Where the bullet points start, it says,

Page 320

[as read] "Image Query allows searches by three methods: Image content, such as color, shape and texture."

You see that?

A. Yeah.

Q. So Ultimedia Manager 1.1 was able to perform an image query using color, shape and texture, correct?

MR. HANSEN: Objection; lacks foundation.

THE WITNESS: And layout.

BY MR. EDWARDS:

Q. And layout. But also color, shape and texture, correct?

A. Yes.

MR. HANSEN: Same objection.

BY MR. EDWARDS:

Q. And then if you go to the last sentence it goes on -- and the bottom of the left column on that same page, going to the top of the right column, it says, "You can drag and drop visual samples into an image query."

You see that?

A. Yeah.

Q. And that would include images, correct?

A. Correct.

Page 321

MR. HANSEN: Objection; lacks foundation.

BY MR. EDWARDS:

Q. Okay?

A. Correct.

Q. Mr. Flickner, every commercial implementation of QBIC as of the Exhibit 3, IEEE 1995 article, returned results of the matching process to the user in the form of a stored image and information associated with that image; is that fair?

MR. HANSEN: Objection; lacks foundation.

THE WITNESS: Yes.

BY MR. EDWARDS:

Q. Okay. And a thumbnail is an image; correct?

A. Thumbnail is an image.

Q. And for every commercial embodiment of QBIC as of the 1995 IEEE article in Exhibit 3, those implementations calculated a distance metric as a measure of similarity for each feature between an input image and a stored image; correct?

MR. HANSEN: Objection; lacks foundation.

THE WITNESS: Correct.

BY MR. EDWARDS:

Q. Okay. And a distance measure value

Element	'004 Claim Recitation	Exemplary Citations in Reference
		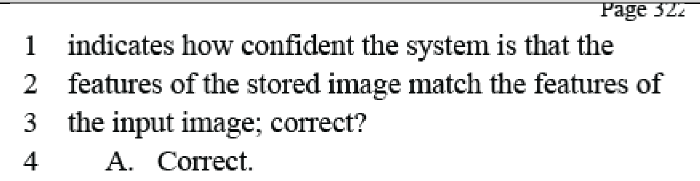 Page 322 indicates how confident the system is that the features of the stored image match the features of the input image; correct? A. Correct. *See also* Flickner Depo at 38-41, 201, 209-210, 214-215. To the extent it is found that QBIC does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that QBIC does not anticipate this claim, QBIC renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine QBIC with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1C]	recognizing the at least one object in the scene as a target object based at least in part on the at least one characteristic of the video stream;	QBIC discloses the limitation. The QBIC Project: Querying Images by Content Using Color, Texture, and Shape (IBM 0002390–404) at 2391 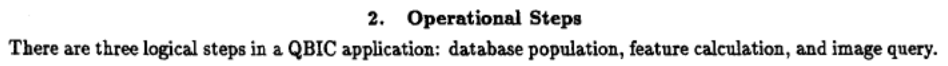 **2. Operational Steps** There are three logical steps in a QBIC application: database population, feature calculation, and image query.

Element	'004 Claim Recitation	Exemplary Citations in Reference
		IBM0002390 at 2391-92: 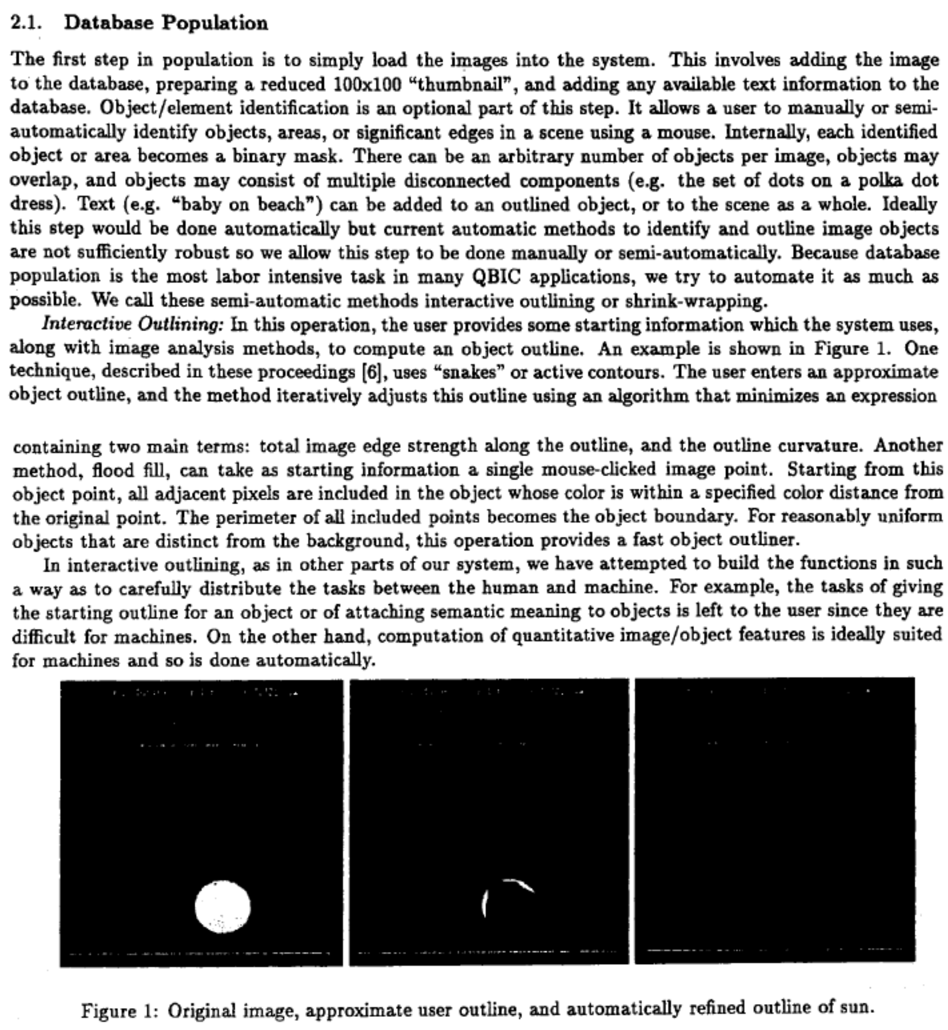 **2.1. Database Population** The first step in population is to simply load the images into the system. This involves adding the image to the database, preparing a reduced 100x100 "thumbnail", and adding any available text information to the database. Object/element identification is an optional part of this step. It allows a user to manually or semi-automatically identify objects, areas, or significant edges in a scene using a mouse. Internally, each identified object or area becomes a binary mask. There can be an arbitrary number of objects per image, objects may overlap, and objects may consist of multiple disconnected components (e.g. the set of dots on a polka dot dress). Text (e.g. "baby on beach") can be added to an outlined object, or to the scene as a whole. Ideally this step would be done automatically but current automatic methods to identify and outline image objects are not sufficiently robust so we allow this step to be done manually or semi-automatically. Because database population is the most labor intensive task in many QBIC applications, we try to automate it as much as possible. We call these semi-automatic methods interactive outlining or shrink-wrapping. *Interactive Outlining:* In this operation, the user provides some starting information which the system uses, along with image analysis methods, to compute an object outline. An example is shown in Figure 1. One technique, described in these proceedings [6], uses "snakes" or active contours. The user enters an approximate object outline, and the method iteratively adjusts this outline using an algorithm that minimizes an expression containing two main terms: total image edge strength along the outline, and the outline curvature. Another method, flood fill, can take as starting information a single mouse-clicked image point. Starting from this object point, all adjacent pixels are included in the object whose color is within a specified color distance from the original point. The perimeter of all included points becomes the object boundary. For reasonably uniform objects that are distinct from the background, this operation provides a fast object outliner. In interactive outlining, as in other parts of our system, we have attempted to build the functions in such a way as to carefully distribute the tasks between the human and machine. For example, the tasks of giving the starting outline for an object or of attaching semantic meaning to objects is left to the user since they are difficult for machines. On the other hand, computation of quantitative image/object features is ideally suited for machines and so is done automatically. Figure 1: Original image, approximate user outline, and automatically refined outline of sun. IBM0002390 at 2392-93

Element	'004 Claim Recitation	Exemplary Citations in Reference
		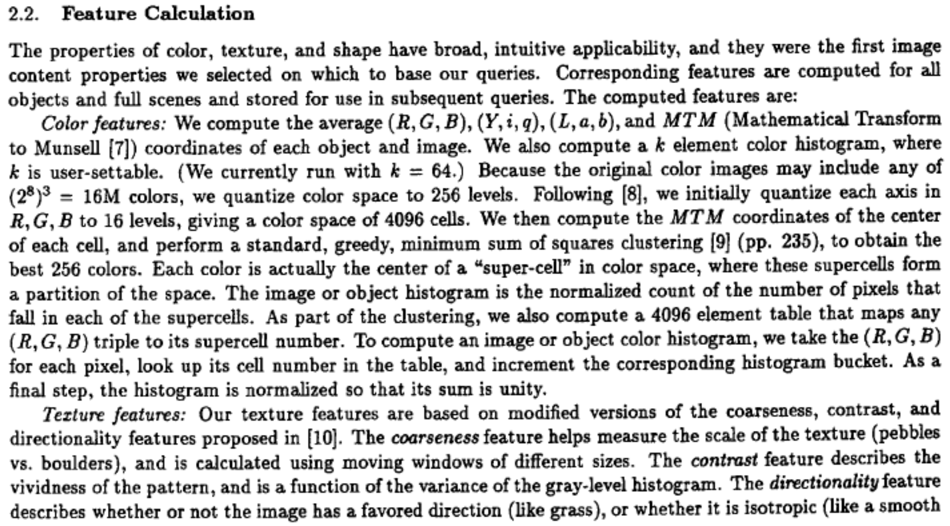 2.2. **Feature Calculation** The properties of color, texture, and shape have broad, intuitive applicability, and they were the first image content properties we selected on which to base our queries. Corresponding features are computed for all objects and full scenes and stored for use in subsequent queries. The computed features are: *Color features:* We compute the average (R,G,B), (Y,i,q), (L,a,b), and MTM (Mathematical Transform to Munsell [7]) coordinates of each object and image. We also compute a k element color histogram, where k is user-settable. (We currently run with $k = 64$.) Because the original color images may include any of $(2^8)^3 = 16$M colors, we quantize color space to 256 levels. Following [8], we initially quantize each axis in R,G,B to 16 levels, giving a color space of 4096 cells. We then compute the MTM coordinates of the center of each cell, and perform a standard, greedy, minimum sum of squares clustering [9] (pp. 235), to obtain the best 256 colors. Each color is actually the center of a "super-cell" in color space, where these supercells form a partition of the space. The image or object histogram is the normalized count of the number of pixels that fall in each of the supercells. As part of the clustering, we also compute a 4096 element table that maps any (R,G,B) triple to its supercell number. To compute an image or object color histogram, we take the (R,G,B) for each pixel, look up its cell number in the table, and increment the corresponding histogram bucket. As a final step, the histogram is normalized so that its sum is unity. *Texture features:* Our texture features are based on modified versions of the coarseness, contrast, and directionality features proposed in [10]. The *coarseness* feature helps measure the scale of the texture (pebbles vs. boulders), and is calculated using moving windows of different sizes. The *contrast* feature describes the vividness of the pattern, and is a function of the variance of the gray-level histogram. The *directionality* feature describes whether or not the image has a favored direction (like grass), or whether it is isotropic (like a smooth

Element	'004 Claim Recitation	Exemplary Citations in Reference
		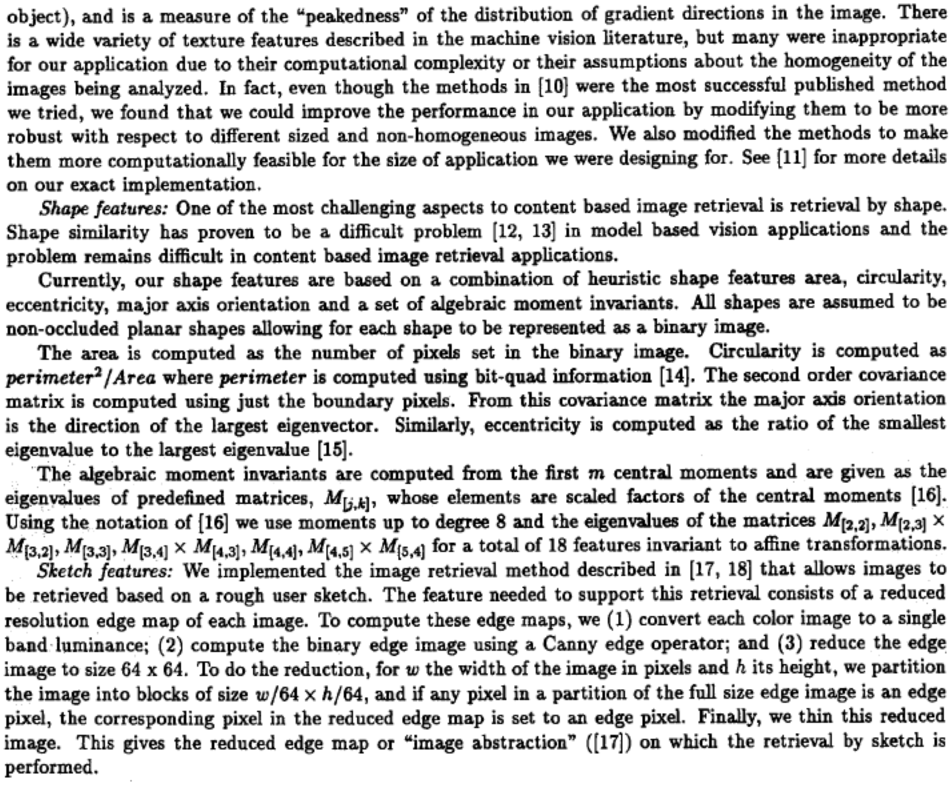 object), and is a measure of the "peakedness" of the distribution of gradient directions in the image. There is a wide variety of texture features described in the machine vision literature, but many were inappropriate for our application due to their computational complexity or their assumptions about the homogeneity of the images being analyzed. In fact, even though the methods in [10] were the most successful published method we tried, we found that we could improve the performance in our application by modifying them to be more robust with respect to different sized and non-homogeneous images. We also modified the methods to make them more computationally feasible for the size of application we were designing for. See [11] for more details on our exact implementation. *Shape features:* One of the most challenging aspects to content based image retrieval is retrieval by shape. Shape similarity has proven to be a difficult problem [12, 13] in model based vision applications and the problem remains difficult in content based image retrieval applications. Currently, our shape features are based on a combination of heuristic shape features area, circularity, eccentricity, major axis orientation and a set of algebraic moment invariants. All shapes are assumed to be non-occluded planar shapes allowing for each shape to be represented as a binary image. The area is computed as the number of pixels set in the binary image. Circularity is computed as $perimeter^2/Area$ where $perimeter$ is computed using bit-quad information [14]. The second order covariance matrix is computed using just the boundary pixels. From this covariance matrix the major axis orientation is the direction of the largest eigenvector. Similarly, eccentricity is computed as the ratio of the smallest eigenvalue to the largest eigenvalue [15]. The algebraic moment invariants are computed from the first m central moments and are given as the eigenvalues of predefined matrices, $M_{[j,k]}$, whose elements are scaled factors of the central moments [16]. Using the notation of [16] we use moments up to degree 8 and the eigenvalues of the matrices $M_{[2,2]}, M_{[2,3]} \times M_{[3,2]}, M_{[3,3]}, M_{[3,4]} \times M_{[4,3]}, M_{[4,4]}, M_{[4,5]} \times M_{[5,4]}$ for a total of 18 features invariant to affine transformations. *Sketch features:* We implemented the image retrieval method described in [17, 18] that allows images to be retrieved based on a rough user sketch. The feature needed to support this retrieval consists of a reduced resolution edge map of each image. To compute these edge maps, we (1) convert each color image to a single band luminance; (2) compute the binary edge image using a Canny edge operator; and (3) reduce the edge image to size 64 x 64. To do the reduction, for w the width of the image in pixels and h its height, we partition the image into blocks of size $w/64 \times h/64$, and if any pixel in a partition of the full size edge image is an edge pixel, the corresponding pixel in the reduced edge map is set to an edge pixel. Finally, we thin this reduced image. This gives the reduced edge map or "image abstraction" ([17]) on which the retrieval by sketch is performed. IBM0002390 at 2393-95

Element	'004 Claim Recitation	Exemplary Citations in Reference
		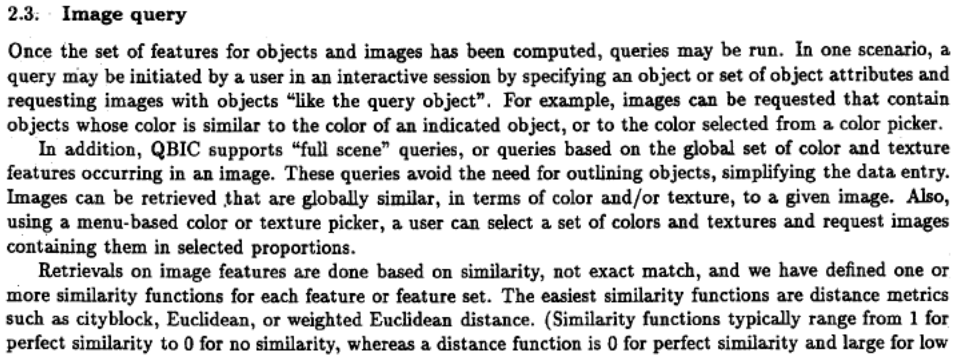 **2.3. Image query** Once the set of features for objects and images has been computed, queries may be run. In one scenario, a query may be initiated by a user in an interactive session by specifying an object or set of object attributes and requesting images with objects "like the query object". For example, images can be requested that contain objects whose color is similar to the color of an indicated object, or to the color selected from a color picker. In addition, QBIC supports "full scene" queries, or queries based on the global set of color and texture features occurring in an image. These queries avoid the need for outlining objects, simplifying the data entry. Images can be retrieved that are globally similar, in terms of color and/or texture, to a given image. Also, using a menu-based color or texture picker, a user can select a set of colors and textures and request images containing them in selected proportions. Retrievals on image features are done based on similarity, not exact match, and we have defined one or more similarity functions for each feature or feature set. The easiest similarity functions are distance metrics such as cityblock, Euclidean, or weighted Euclidean distance. (Similarity functions typically range from 1 for perfect similarity to 0 for no similarity, whereas a distance function is 0 for perfect similarity and large for low

Element	'004 Claim Recitation	Exemplary Citations in Reference
		similarity, so we must convert as necessary to obtain consistent measures.) Different similarity measures may be used, and they may vary from application to application. The matching methods and similarity functions for the different feature types are described below: *Color:* For average color, the distance between a query object and database object is weighted Euclidean distance. The weights are user adjustable so that we may, for example, downweight the L component of an (L, a, b) query. Best results, as judged by our subjective evaluation in on-going experiments, are obtained by representing the average color in the MTM color space, and by inversely weighting each component by its standard deviation over the samples in the database. Improved results are obtained when color matching is done not on average color, but on the distribution of colors that occur in an object or image. In this way, a red and blue object matches other red and blue objects better than to a purple object. To match color histograms, we use a method similar to the one described in [8]. Let X be the query histogram and Y the histogram of an item in the database, both normalized. We compute the element by element difference histogram Z. Then the similarity between X and Y is given by $\|Z\| = Z^T A Z$, where A is a symmetric color similarity matrix with $$a(i,j) \equiv 1 - d(c_i, c_j)/d_{max}$$ where c_i and c_j are the ith and jth colors in the color histograms, and $d(c_i, c_j)$ is the MTM color distance, and d_{max} is the maximum distance between any two colors ([8] used Luv distance.) This metric gives the (weighted) length of the difference vector between X and Y, weighted by A, which accounts for the perceptual distance between different pairs of colors. The result is a measure that accounts for both the difference between the amounts in any given color (e.g. a particular shade of red), as well as differences between similar colors (e.g. red and orange). Based on our experience, we also found it useful for the user to be able to request images with $x\%$ of color 1, $y\%$ of color 2, etc. where both the percentages and the colors are explicitly specified by the user. This allows searches such as for a beach scene (25% white (for the sand) and 50% blue (for the sky and water)), and for many scenes for which the user has seen before, and for which he/she can remember approximate color percentages. The color picker that allows this type of query is shown in Figure 2. 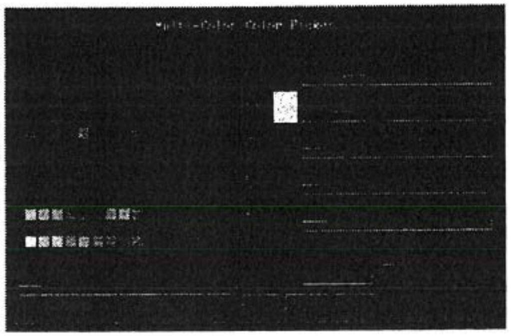 Figure 2: Multi-Color Color Picker with two selected colors. The palette on the left displays the set of selectable colors. Once a color is selected, it appears in a box on the right, and the amount of the color can be specified by the sliders. In this case, 25% light blue and 15% white have been selected. *Texture:* Texture distance is computed as weighted Euclidean distance in the three dimensional texture space. The most common weighting factors are the inverse variances for each component, computed over the

Element	'004 Claim Recitation	Exemplary Citations in Reference
		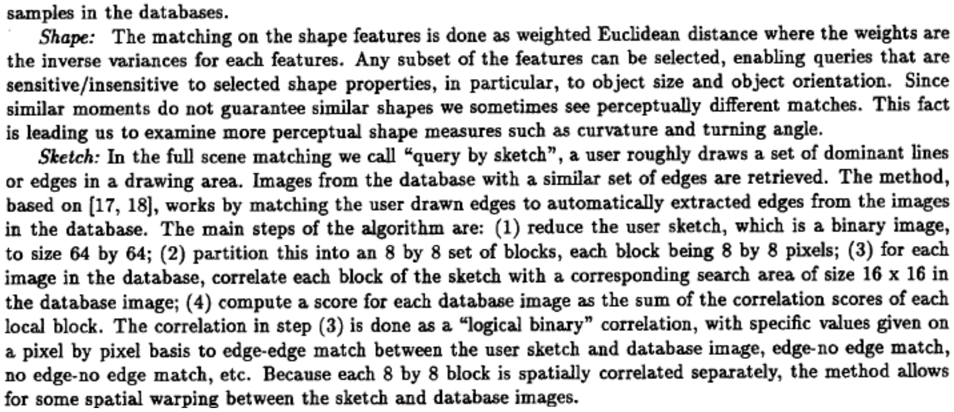 samples in the databases. *Shape:* The matching on the shape features is done as weighted Euclidean distance where the weights are the inverse variances for each features. Any subset of the features can be selected, enabling queries that are sensitive/insensitive to selected shape properties, in particular, to object size and object orientation. Since similar moments do not guarantee similar shapes we sometimes see perceptually different matches. This fact is leading us to examine more perceptual shape measures such as curvature and turning angle. *Sketch:* In the full scene matching we call "query by sketch", a user roughly draws a set of dominant lines or edges in a drawing area. Images from the database with a similar set of edges are retrieved. The method, based on [17, 18], works by matching the user drawn edges to automatically extracted edges from the images in the database. The main steps of the algorithm are: (1) reduce the user sketch, which is a binary image, to size 64 by 64; (2) partition this into an 8 by 8 set of blocks, each block being 8 by 8 pixels; (3) for each image in the database, correlate each block of the sketch with a corresponding search area of size 16 x 16 in the database image; (4) compute a score for each database image as the sum of the correlation scores of each local block. The correlation in step (3) is done as a "logical binary" correlation, with specific values given on a pixel by pixel basis to edge-edge match between the user sketch and database image, edge-no edge match, no edge-no edge match, etc. Because each 8 by 8 block is spatially correlated separately, the method allows for some spatial warping between the sketch and database images. IBM0002390 at 2395 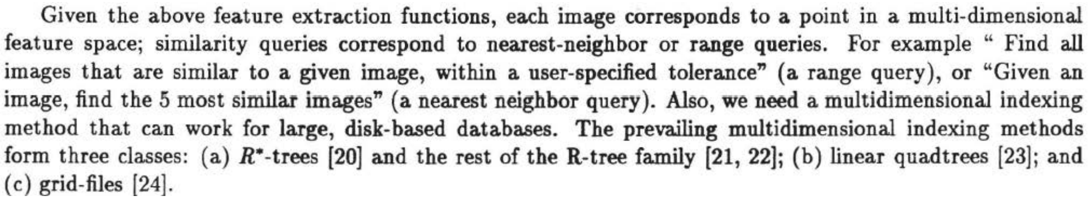 Given the above feature extraction functions, each image corresponds to a point in a multi-dimensional feature space; similarity queries correspond to nearest-neighbor or range queries. For example " Find all images that are similar to a given image, within a user-specified tolerance" (a range query), or "Given an image, find the 5 most similar images" (a nearest neighbor query). Also, we need a multidimensional indexing method that can work for large, disk-based databases. The prevailing multidimensional indexing methods form three classes: (a) R^*-trees [20] and the rest of the R-tree family [21, 22]; (b) linear quadtrees [23]; and (c) grid-files [24]. IBM0002390 at 2396:

Element	'004 Claim Recitation	Exemplary Citations in Reference
		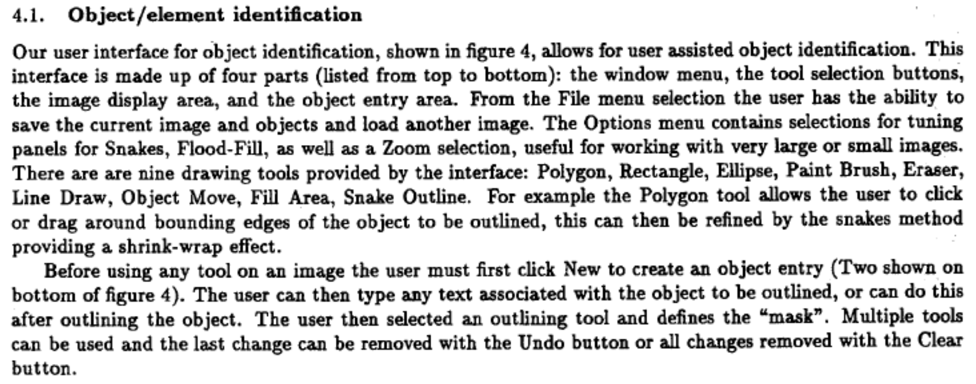 4.1. **Object/element identification** Our user interface for object identification, shown in figure 4, allows for user assisted object identification. This interface is made up of four parts (listed from top to bottom): the window menu, the tool selection buttons, the image display area, and the object entry area. From the File menu selection the user has the ability to save the current image and objects and load another image. The Options menu contains selections for tuning panels for Snakes, Flood-Fill, as well as a Zoom selection, useful for working with very large or small images. There are are nine drawing tools provided by the interface: Polygon, Rectangle, Ellipse, Paint Brush, Eraser, Line Draw, Object Move, Fill Area, Snake Outline. For example the Polygon tool allows the user to click or drag around bounding edges of the object to be outlined, this can then be refined by the snakes method providing a shrink-wrap effect. Before using any tool on an image the user must first click New to create an object entry (Two shown on bottom of figure 4). The user can then type any text associated with the object to be outlined, or can do this after outlining the object. The user then selected an outlining tool and defines the "mask". Multiple tools can be used and the last change can be removed with the Undo button or all changes removed with the Clear button. Query by Image and Video Content: The QBIC System (IBM 0000893–902): "To achieve this functionality, QBIC has two main components: database population (the process of creating an image database) and database query. During the population, images and videos are processed to extract features describing their content-colors, textures, shapes, and camera and object motion-and the features are stored in a database. During the query, the user composes a query graphically. Features are generated from the graphical query and then input to a matching engine that finds images or videos from the database with similar features. Figure 2 shows the system architecture." *Id.* at p. 3. "For both population and query, the QBIC data model has • still images or scenes (full images) that contain objects (subsets of an image), and • video shots that consist of sets of contiguous frames and contain motion objects. For still images, the QBIC data model distinguishes between 'scenes' (or images) and 'objects.' A scene is an image or single representative frame of video. An object is a part of a scene-for example, the fox in Figure 3-or a moving entity in a video. For still image database population, features are extracted from images and objects and stored in a database as shown in the top left part of Figure 2. Videos are broken into clips called shots. Representative frames, or r-frames,

Element	'004 Claim Recitation	Exemplary Citations in Reference
		are generated for each extracted shot. R-frames are treated as still images, and features are extracted and stored in the database. Further processing of shots generates motion objects-for example, a car moving across the screen. Queries are allowed on objects ('Find images with a red, round object'), scenes ('Find images that have approximately 30-percent red and IS-percent blue colors'), shots ('Find all shots panning from left to right'), or any combination ('Find images that have 30 percent red and contain a blue textured object')." *Id*. at p. 3. "In still image database population, the images are reduced to a standard-sized icon called a thumbnail and annotated with any available text information. Object identification is an optional but key part of this step. It lets users manually, semiautomatically, or fully automatically identify interesting regions-which we call objects-in the images. . . . For video data, database population has three major components: • shot detection, • representative frame creation for each shot, and • derivation of a layered representation of coherently moving structures/ objects. . . . Once the shot boundaries have been detected, each shot is represented using an r-frame. R-frames are used for several purposes. First, during database population, r-frames are treated as still images in which objects can be identified by using the previously described methods." *Id*. at pp. 3–6. "**SAMPLE QUERIES** For each full-scene image, identified image object, r-frame, and identified video object resulting from the above processing, a set of features is computed to allow content-based queries. The features are computed and stored during database population. We present a brief description of the features and the associated queries. Mathematical details on the features and matching methods can be found in Ashley et al. and Niblack et al. Average color queries let users find images or objects that are similar to a selected color, say from a color wheel, or to the color of an object. The feature used in the query is a 3D vector of Munsell color coordinates. Histogram color

Element	'004 Claim Recitation	Exemplary Citations in Reference
		queries return items with matching color distributions—say, a fabric pattern with approximately 40 percent red and 20 percent blue. For this case, the underlying feature is a 256-element histogram computed over a quantized version of the color space. Figure 7 shows a histogram query on still images and a color query on video r-frames. Note that in the query specification for the histogram query of Figure 7, the user has selected percentages of two colors (blue and white) by adjusting sliders. Using such a query, an advertising agent could, for example, search for a picture of a beach scene, one predominantly blue (for sky and water) and white (for sand and clouds); or find images with similar color spreads for a uniform ad campaign. The average color query demonstrates a query against a video shot database where the user is searching for red r-frames. Again, the query specification is on the left and the best hits are on the right. Figure 8 shows an example texture query. In this case, the query is specified by selecting from a sampler—a set of prestored example images. The underlying texture features are mathematical representations of coarseness, contrast, and directionality features. Coarseness measures the scale of a texture (pebbles versus boulders), contrast describes its vividness, and directionality describes whether it has a favored direction (like grass) or not (like a smooth object). An object shape query is shown in Figure 8. In this case, the query specification is the drawn shape on the left. Area, circularity, eccentricity, major-axis direction, features derived from the object moments, and a set of tangent angles around the object perimeter are the features used to characterize and match shapes. Figure 9 illustrates query by sketch. In this case, the query specification is a freehand drawing of the dominant lines and edges in the image. The sketch feature is an automatically extracted reduced-resolution 'edge map.' Matching is done by using a template-matching technique. A multiobject query asking for images that contain both a red round object and a green textured object is shown in the bottom of Figure 9. The features are standard color and texture. The matching is done by combining the color and texture distances. Combining distances is applied to arbitrary sets of objects

Element	'004 Claim Recitation	Exemplary Citations in Reference
		and features to implement logical And semantics. . . . . **ANNOTATION AND DATABASE POPULATION TOOLS**. Automatic methods (such as our Positional Color query) that don't rely on object identification, methods that identify objects automatically as in the museum image example, fast and easy-to-use semiautomatic outlining tools, and motion-based segmentation algorithms will enable additional application areas. **FEATURE EXTRACTION AND MATCHING METHODS**. New mathematical representations of video, image, and object attributes that capture 'interesting' features for retrieval are needed. Features that describe new image properties such as alternate texture measures or that are based on fractals or wavelet representations, for example, may offer advantages of representation, indexability, and ease of similarity matching. **INTEGRATION WITH TEXT AND PARAMETRIC ANNOTATION**. Query by visual content complements and extends existing query methods. Systems must be able to integrate queries combining date, subject matter, price, and availability with content properties such as color, texture, and shape." *Id.* at pp. 7–8.

Element	'004 Claim Recitation	Exemplary Citations in Reference
		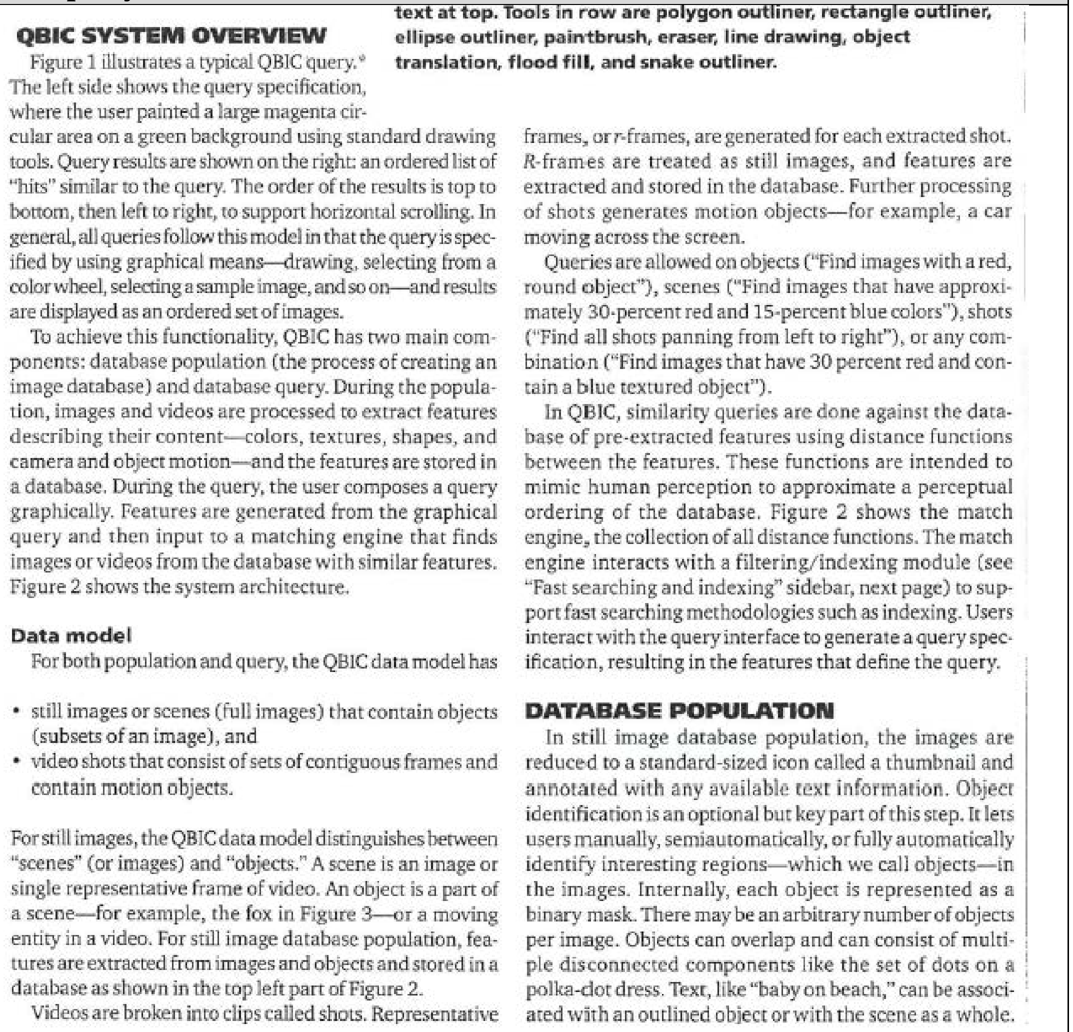 **text at top. Tools in row are polygon outliner, rectangle outliner, ellipse outliner, paintbrush, eraser, line drawing, object translation, flood fill, and snake outliner.** **QBIC SYSTEM OVERVIEW** Figure 1 illustrates a typical QBIC query.* The left side shows the query specification, where the user painted a large magenta circular area on a green background using standard drawing tools. Query results are shown on the right: an ordered list of "hits" similar to the query. The order of the results is top to bottom, then left to right, to support horizontal scrolling. In general, all queries follow this model in that the query is specified by using graphical means—drawing, selecting from a color wheel, selecting a sample image, and so on—and results are displayed as an ordered set of images. To achieve this functionality, QBIC has two main components: database population (the process of creating an image database) and database query. During the population, images and videos are processed to extract features describing their content—colors, textures, shapes, and camera and object motion—and the features are stored in a database. During the query, the user composes a query graphically. Features are generated from the graphical query and then input to a matching engine that finds images or videos from the database with similar features. Figure 2 shows the system architecture. **Data model** For both population and query, the QBIC data model has • still images or scenes (full images) that contain objects (subsets of an image), and • video shots that consist of sets of contiguous frames and contain motion objects. For still images, the QBIC data model distinguishes between "scenes" (or images) and "objects." A scene is an image or single representative frame of video. An object is a part of a scene—for example, the fox in Figure 3—or a moving entity in a video. For still image database population, features are extracted from images and objects and stored in a database as shown in the top left part of Figure 2. Videos are broken into clips called shots. Representative frames, or *r*-frames, are generated for each extracted shot. *R*-frames are treated as still images, and features are extracted and stored in the database. Further processing of shots generates motion objects—for example, a car moving across the screen. Queries are allowed on objects ("Find images with a red, round object"), scenes ("Find images that have approximately 30-percent red and 15-percent blue colors"), shots ("Find all shots panning from left to right"), or any combination ("Find images that have 30 percent red and contain a blue textured object"). In QBIC, similarity queries are done against the database of pre-extracted features using distance functions between the features. These functions are intended to mimic human perception to approximate a perceptual ordering of the database. Figure 2 shows the match engine, the collection of all distance functions. The match engine interacts with a filtering/indexing module (see "Fast searching and indexing" sidebar, next page) to support fast searching methodologies such as indexing. Users interact with the query interface to generate a query specification, resulting in the features that define the query. **DATABASE POPULATION** In still image database population, the images are reduced to a standard-sized icon called a thumbnail and annotated with any available text information. Object identification is an optional but key part of this step. It lets users manually, semiautomatically, or fully automatically identify interesting regions—which we call objects—in the images. Internally, each object is represented as a binary mask. There may be an arbitrary number of objects per image. Objects can overlap and can consist of multiple disconnected components like the set of dots on a polka-dot dress. Text, like "baby on beach," can be associated with an outlined object or with the scene as a whole. IBM0000893 at 895

Element	'004 Claim Recitation	Exemplary Citations in Reference
		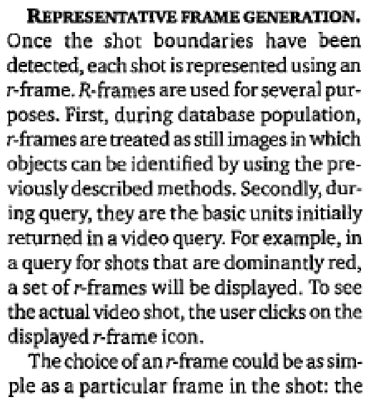 **REPRESENTATIVE FRAME GENERATION.** Once the shot boundaries have been detected, each shot is represented using an *r*-frame. *R*-frames are used for several purposes. First, during database population, *r*-frames are treated as still images in which objects can be identified by using the previously described methods. Secondly, during query, they are the basic units initially returned in a video query. For example, in a query for shots that are dominantly red, a set of *r*-frames will be displayed. To see the actual video shot, the user clicks on the displayed *r*-frame icon. The choice of an *r*-frame could be as simple as a particular frame in the shot: the IBM0000893 at 898 Image, Audio, and Video Extenders Administration and Programming (IBM 000001–617): "Your applications can even search for images by content. Imagine an application that uses visual examples to search for images. With such an application, users could select an example image and have the application find other images that have colors or textures similar to those in the example. With DB2 extenders' Query by Image Content (QBIC) capability, you can create applications that search for images in this visual way." *Id*. at p. 4. "To make an image available for searching by content, you catalog the image. When you catalog an image, the Image Extender analyzes the image, by computing the feature values for the image, and stores the values in a QBIC catalog. When you search for an image by content, your query identifies one or more features for the search (such as average color), a source for each feature (such as an example image), and a target set of cataloged images. The Image Extender computes the feature value of the source and compares it to the cataloged

Element	'004 Claim Recitation	Exemplary Citations in Reference
		feature values for the target images. It then computes a score that indicates how similar the feature values of the target images are to the source. You can have the Image Extender return the images whose features are most similar to the source. The Image Extender will return the handle of each image and the image score. You can also have the Image Extender return only the score of a single image." *Id*. at pp. 20–21. QBIC Demo (IBM000747) 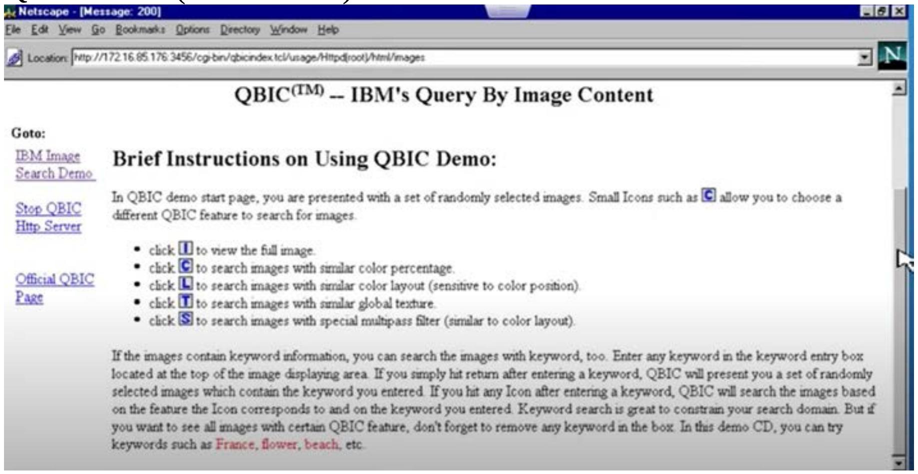 QBIC$^{(TM)}$ -- IBM's Query By Image Content **Brief Instructions on Using QBIC Demo:** In QBIC demo start page, you are presented with a set of randomly selected images. Small Icons such as C allow you to choose a different QBIC feature to search for images. - click I to view the full image. - click C to search images with similar color percentage. - click L to search images with similar color layout (sensitive to color position). - click T to search images with similar global texture. - click S to search images with special multipass filter (similar to color layout). If the images contain keyword information, you can search the images with keyword, too. Enter any keyword in the keyword entry box located at the top of the image displaying area. If you simply hit return after entering a keyword, QBIC will present you a set of randomly selected images which contain the keyword you entered. If you hit any Icon after entering a keyword, QBIC will search the images based on the feature the Icon corresponds to and on the keyword you entered. Keyword search is great to constrain your search domain. But if you want to see all images with certain QBIC feature, don't forget to remove any keyword in the box. In this demo CD, you can try keywords such as France, flower, beach, etc. QBIC Demo (IBM000747)

Element	'004 Claim Recitation	Exemplary Citations in Reference
		 Location: http://172.16.85.176:3456/cgi-bin/ibm.tcl/histpicker Keywords: Red: 59 Green: 131 Blue: 59 9% (255,255,255) 8% (59,131,59) 4% (106,182,106) To specify your color histogram query, "paint" a graphical representation of your pattern of colors and approximate percentages that you want. First select a *percentage* by clicking in the painting area to block off a rectangle of the desired size (percentage of the full area). All new regions are painted white, initially. Then select a color from the color picker to *paint* your rectangle. You can either click in the color regions or enter your desired QBIC Demo (IBM000747) 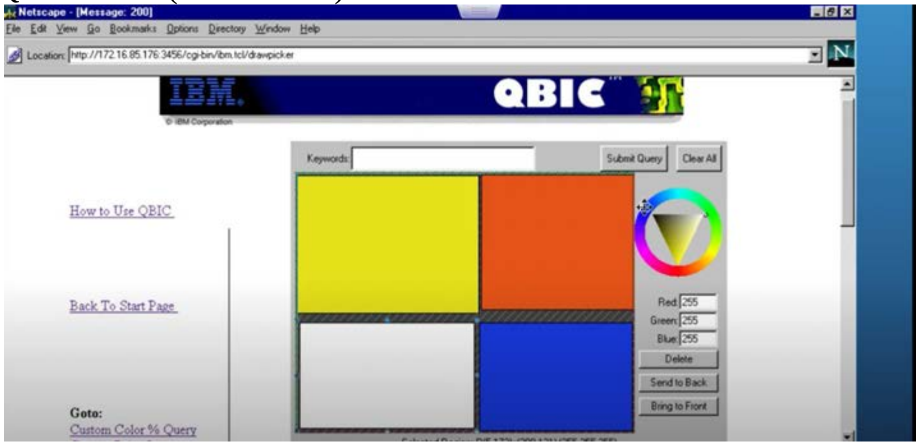 Location: http://172.16.85.176:3456/cgi-bin/ibm.tcl/drawpicker Keywords: Red: 255 Green: 255 Blue: 255

Element	'004 Claim Recitation	Exemplary Citations in Reference
		 Location: http://172.16.85.176:3456/cgi-bin/ibm.tcl/search?fname=HybridColor&france10%2Ejpg.x=0 **Goto:** Custom Color % Query Custom Paint Query Random Images Official QBIC Page **Query was:** Example: =france10.jpg Query Type: QbHybridColorFeatureClass Flickner Depo at 21

Element	'004 Claim Recitation	Exemplary Citations in Reference
		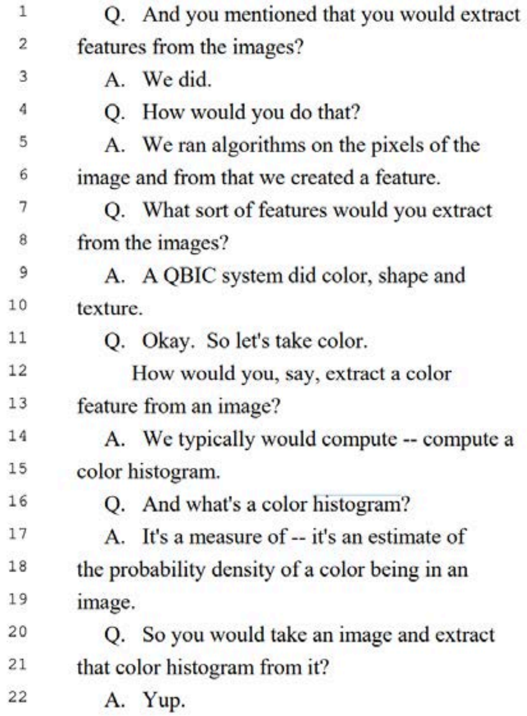 Q. And you mentioned that you would extract features from the images? A. We did. Q. How would you do that? A. We ran algorithms on the pixels of the image and from that we created a feature. Q. What sort of features would you extract from the images? A. A QBIC system did color, shape and texture. Q. Okay. So let's take color. How would you, say, extract a color feature from an image? A. We typically would compute -- compute a color histogram. Q. And what's a color histogram? A. It's a measure of -- it's an estimate of the probability density of a color being in an image. Q. So you would take an image and extract that color histogram from it? A. Yup. Flickner Depo at 88 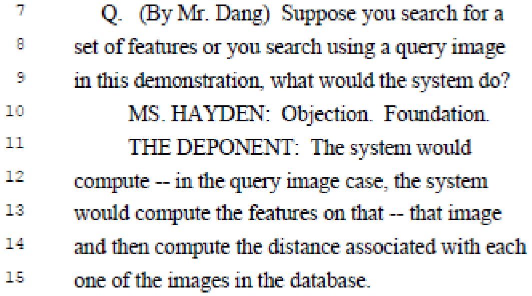 Q. (By Mr. Dang) Suppose you search for a set of features or you search using a query image in this demonstration, what would the system do? MS. HAYDEN: Objection. Foundation. THE DEPONENT: The system would compute -- in the query image case, the system would compute the features on that -- that image and then compute the distance associated with each one of the images in the database.

Element	'004 Claim Recitation	Exemplary Citations in Reference
		Flickner Depo at 22-25 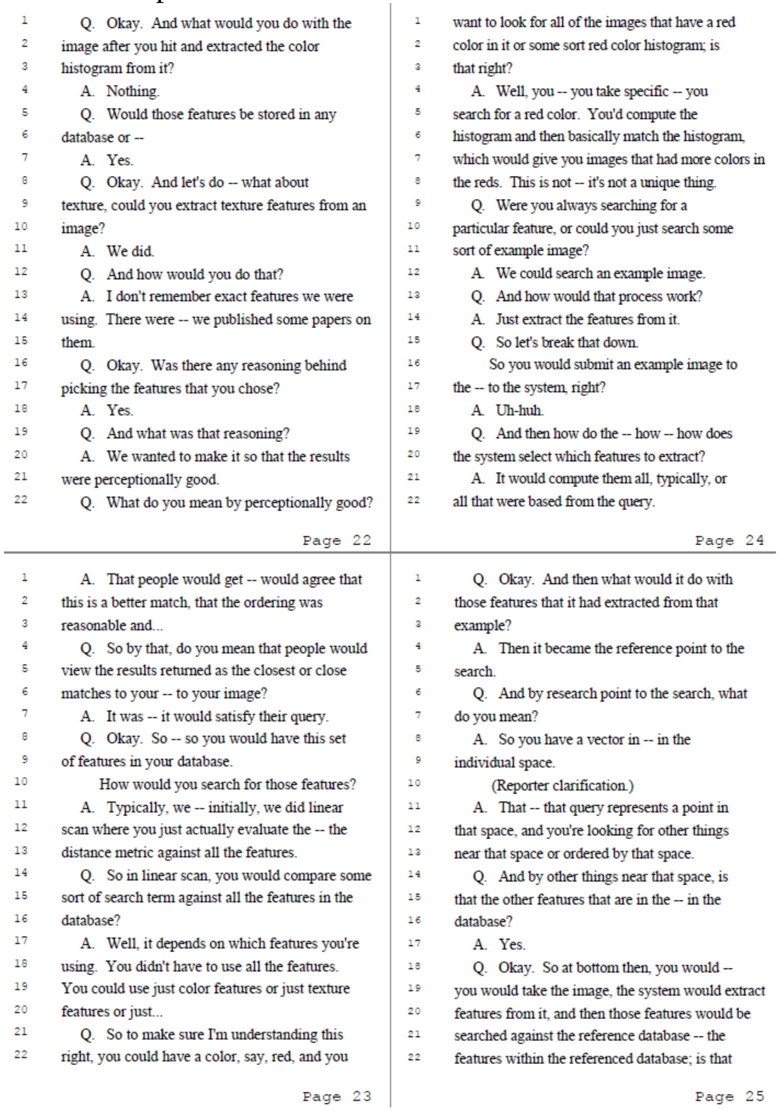 Q. Okay. And what would you do with the image after you hit and extracted the color histogram from it? A. Nothing. Q. Would those features be stored in any database or -- A. Yes. Q. Okay. And let's do -- what about texture, could you extract texture features from an image? A. We did. Q. And how would you do that? A. I don't remember exact features we were using. There were -- we published some papers on them. Q. Okay. Was there any reasoning behind picking the features that you chose? A. Yes. Q. And what was that reasoning? A. We wanted to make it so that the results were perceptionally good. Q. What do you mean by perceptionally good? Page 22 A. That people would get -- would agree that this is a better match, that the ordering was reasonable and... Q. So by that, do you mean that people would view the results returned as the closest or close matches to your -- to your image? A. It was -- it would satisfy their query. Q. Okay. So -- so you would have this set of features in your database. How would you search for those features? A. Typically, we -- initially, we did linear scan where you just actually evaluate the -- the distance metric against all the features. Q. So in linear scan, you would compare some sort of search term against all the features in the database? A. Well, it depends on which features you're using. You didn't have to use all the features. You could use just color features or just texture features or just... Q. So to make sure I'm understanding this right, you could have a color, say, red, and you Page 23 want to look for all of the images that have a red color in it or some sort red color histogram; is that right? A. Well, you -- you take specific -- you search for a red color. You'd compute the histogram and then basically match the histogram, which would give you images that had more colors in the reds. This is not -- it's not a unique thing. Q. Were you always searching for a particular feature, or could you just search some sort of example image? A. We could search an example image. Q. And how would that process work? A. Just extract the features from it. Q. So let's break that down. So you would submit an example image to the -- to the system, right? A. Uh-huh. Q. And then how do the -- how -- how does the system select which features to extract? A. It would compute them all, typically, or all that were based from the query. Page 24 Q. Okay. And then what would it do with those features that it had extracted from that example? A. Then it became the reference point to the search. Q. And by research point to the search, what do you mean? A. So you have a vector in -- in the individual space. (Reporter clarification.) A. That -- that query represents a point in that space, and you're looking for other things near that space or ordered by that space. Q. And by other things near that space, is that the other features that are in the -- in the database? A. Yes. Q. Okay. So at bottom then, you would -- you would take the image, the system would extract features from it, and then those features would be searched against the reference database -- the features within the referenced database; is that Page 25

Element	'004 Claim Recitation	Exemplary Citations in Reference
		Flickner Depo at 36 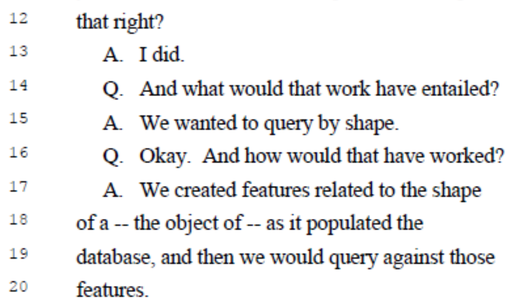 Q. So underlying this paper, you mentioned you did more work on the shape side of things; is that right? A. I did. Q. And what would that work have entailed? A. We wanted to query by shape. Q. Okay. And how would that have worked? A. We created features related to the shape of a -- the object of -- as it populated the database, and then we would query against those features. Flickner 8/29/23 Depo at 33–37

Element	'004 Claim Recitation	Exemplary Citations in Reference
		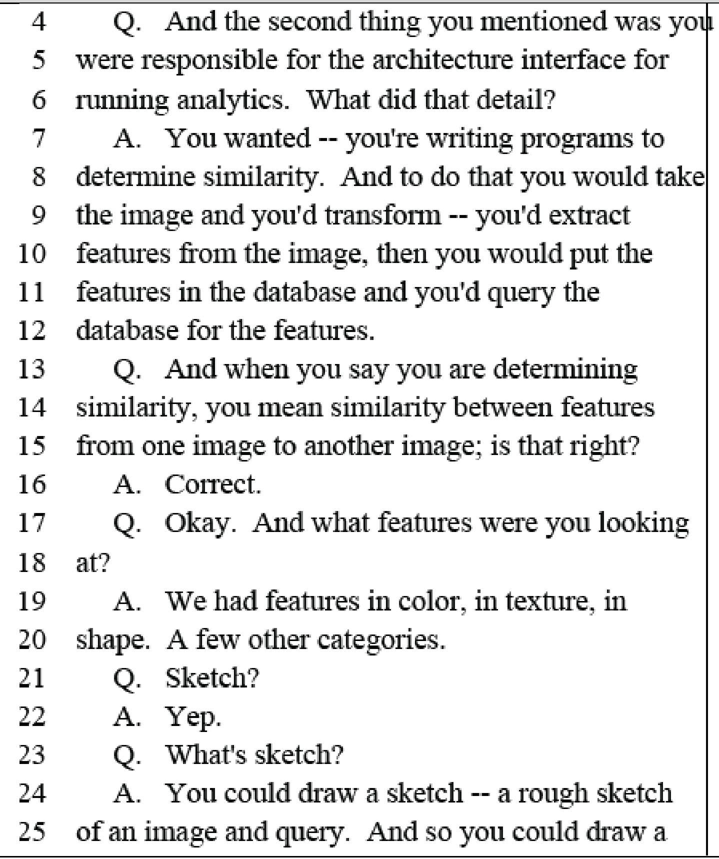 Q. And the second thing you mentioned was you were responsible for the architecture interface for running analytics. What did that detail? A. You wanted -- you're writing programs to determine similarity. And to do that you would take the image and you'd transform -- you'd extract features from the image, then you would put the features in the database and you'd query the database for the features. Q. And when you say you are determining similarity, you mean similarity between features from one image to another image; is that right? A. Correct. Q. Okay. And what features were you looking at? A. We had features in color, in texture, in shape. A few other categories. Q. Sketch? A. Yep. Q. What's sketch? A. You could draw a sketch -- a rough sketch of an image and query. And so you could draw a

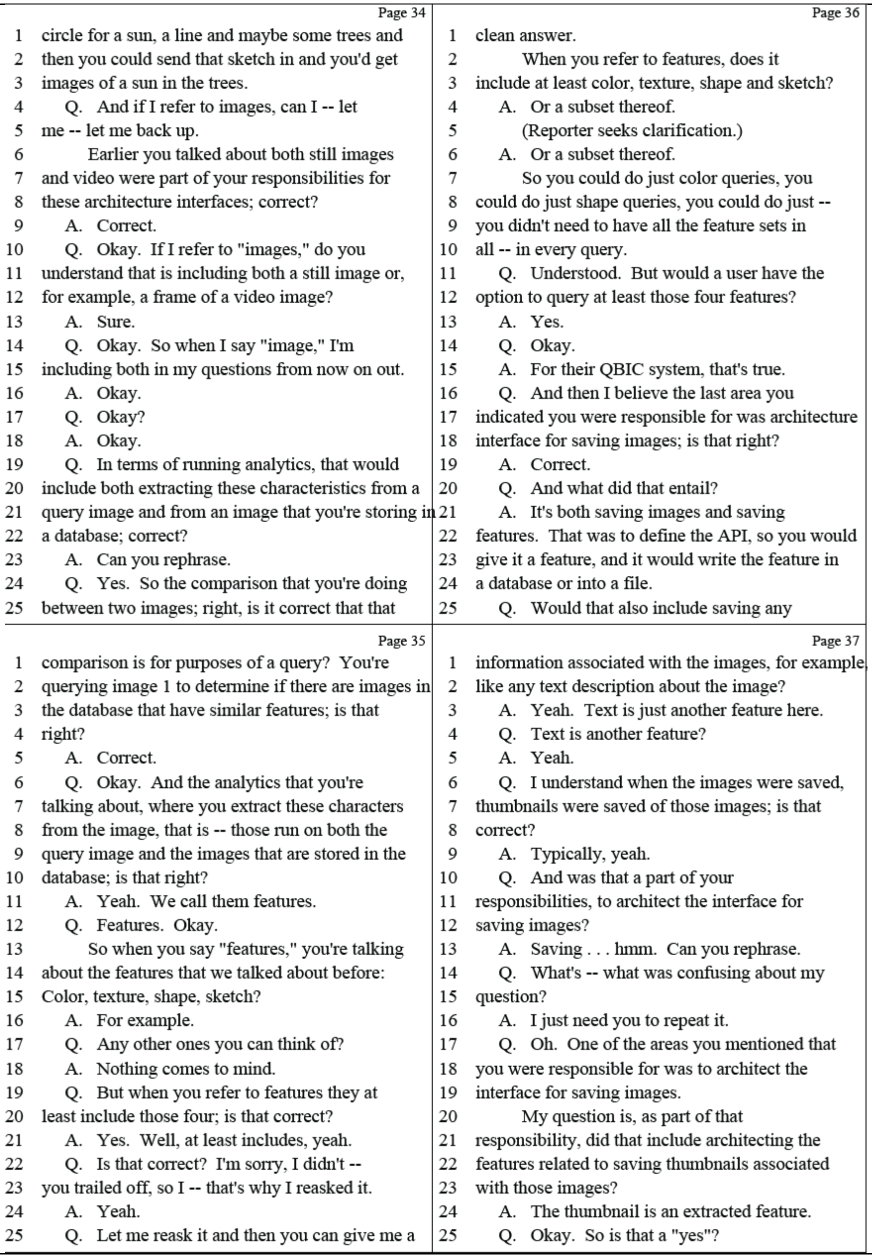

Page 34

circle for a sun, a line and maybe some trees and then you could send that sketch in and you'd get images of a sun in the trees.
Q. And if I refer to images, can I -- let me -- let me back up.
Earlier you talked about both still images and video were part of your responsibilities for these architecture interfaces; correct?
A. Correct.
Q. Okay. If I refer to "images," do you understand that is including both a still image or, for example, a frame of a video image?
A. Sure.
Q. Okay. So when I say "image," I'm including both in my questions from now on out.
A. Okay.
Q. Okay?
A. Okay.
Q. In terms of running analytics, that would include both extracting these characteristics from a query image and from an image that you're storing in a database; correct?
A. Can you rephrase.
Q. Yes. So the comparison that you're doing between two images; right, is it correct that that

Page 35

comparison is for purposes of a query? You're querying image 1 to determine if there are images in the database that have similar features; is that right?
A. Correct.
Q. Okay. And the analytics that you're talking about, where you extract these characters from the image, that is -- those run on both the query image and the images that are stored in the database; is that right?
A. Yeah. We call them features.
Q. Features. Okay.
So when you say "features," you're talking about the features that we talked about before: Color, texture, shape, sketch?
A. For example.
Q. Any other ones you can think of?
A. Nothing comes to mind.
Q. But when you refer to features they at least include those four; is that correct?
A. Yes. Well, at least includes, yeah.
Q. Is that correct? I'm sorry, I didn't -- you trailed off, so I -- that's why I reasked it.
A. Yeah.
Q. Let me reask it and then you can give me a

Page 36

clean answer.
When you refer to features, does it include at least color, texture, shape and sketch?
A. Or a subset thereof.
(Reporter seeks clarification.)
A. Or a subset thereof.
So you could do just color queries, you could do just shape queries, you could do just -- you didn't need to have all the feature sets in all -- in every query.
Q. Understood. But would a user have the option to query at least those four features?
A. Yes.
Q. Okay.
A. For their QBIC system, that's true.
Q. And then I believe the last area you indicated you were responsible for was architecture interface for saving images; is that right?
A. Correct.
Q. And what did that entail?
A. It's both saving images and saving features. That was to define the API, so you would give it a feature, and it would write the feature in a database or into a file.
Q. Would that also include saving any

Page 37

information associated with the images, for example, like any text description about the image?
A. Yeah. Text is just another feature here.
Q. Text is another feature?
A. Yeah.
Q. I understand when the images were saved, thumbnails were saved of those images; is that correct?
A. Typically, yeah.
Q. And was that a part of your responsibilities, to architect the interface for saving images?
A. Saving . . . hmm. Can you rephrase.
Q. What's -- what was confusing about my question?
A. I just need you to repeat it.
Q. Oh. One of the areas you mentioned that you were responsible for was to architect the interface for saving images.
My question is, as part of that responsibility, did that include architecting the features related to saving thumbnails associated with those images?
A. The thumbnail is an extracted feature.
Q. Okay. So is that a "yes"?

Element	'004 Claim Recitation	Exemplary Citations in Reference
		Id. at 52–53 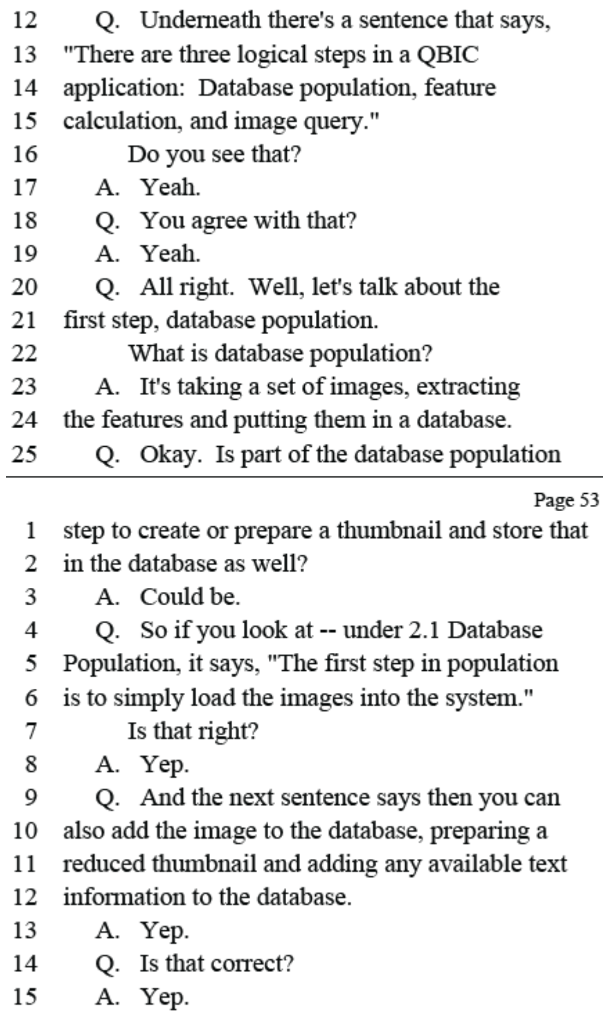 Q. Underneath there's a sentence that says, "There are three logical steps in a QBIC application: Database population, feature calculation, and image query." Do you see that? A. Yeah. Q. You agree with that? A. Yeah. Q. All right. Well, let's talk about the first step, database population. What is database population? A. It's taking a set of images, extracting the features and putting them in a database. Q. Okay. Is part of the database population step to create or prepare a thumbnail and store that in the database as well? Page 53 A. Could be. Q. So if you look at -- under 2.1 Database Population, it says, "The first step in population is to simply load the images into the system." Is that right? A. Yep. Q. And the next sentence says then you can also add the image to the database, preparing a reduced thumbnail and adding any available text information to the database. A. Yep. Q. Is that correct? A. Yep. *Id.* at 77–78

Element	'004 Claim Recitation	Exemplary Citations in Reference
		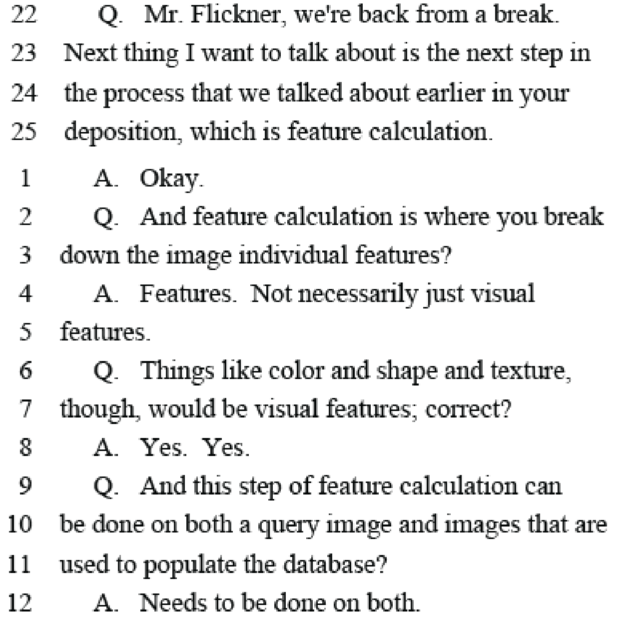 Q. Mr. Flickner, we're back from a break. Next thing I want to talk about is the next step in the process that we talked about earlier in your deposition, which is feature calculation. A. Okay. Q. And feature calculation is where you break down the image individual features? A. Features. Not necessarily just visual features. Q. Things like color and shape and texture, though, would be visual features; correct? A. Yes. Yes. Q. And this step of feature calculation can be done on both a query image and images that are used to populate the database? A. Needs to be done on both. *Id.* at 85–99 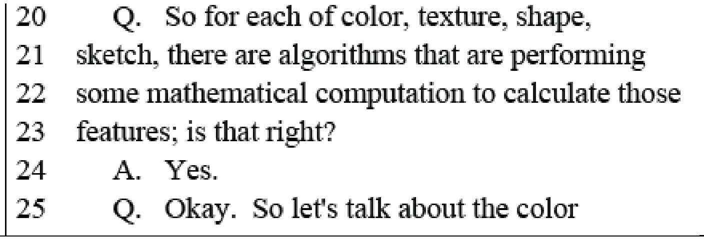 Q. So for each of color, texture, shape, sketch, there are algorithms that are performing some mathematical computation to calculate those features; is that right? A. Yes. Q. Okay. So let's talk about the color

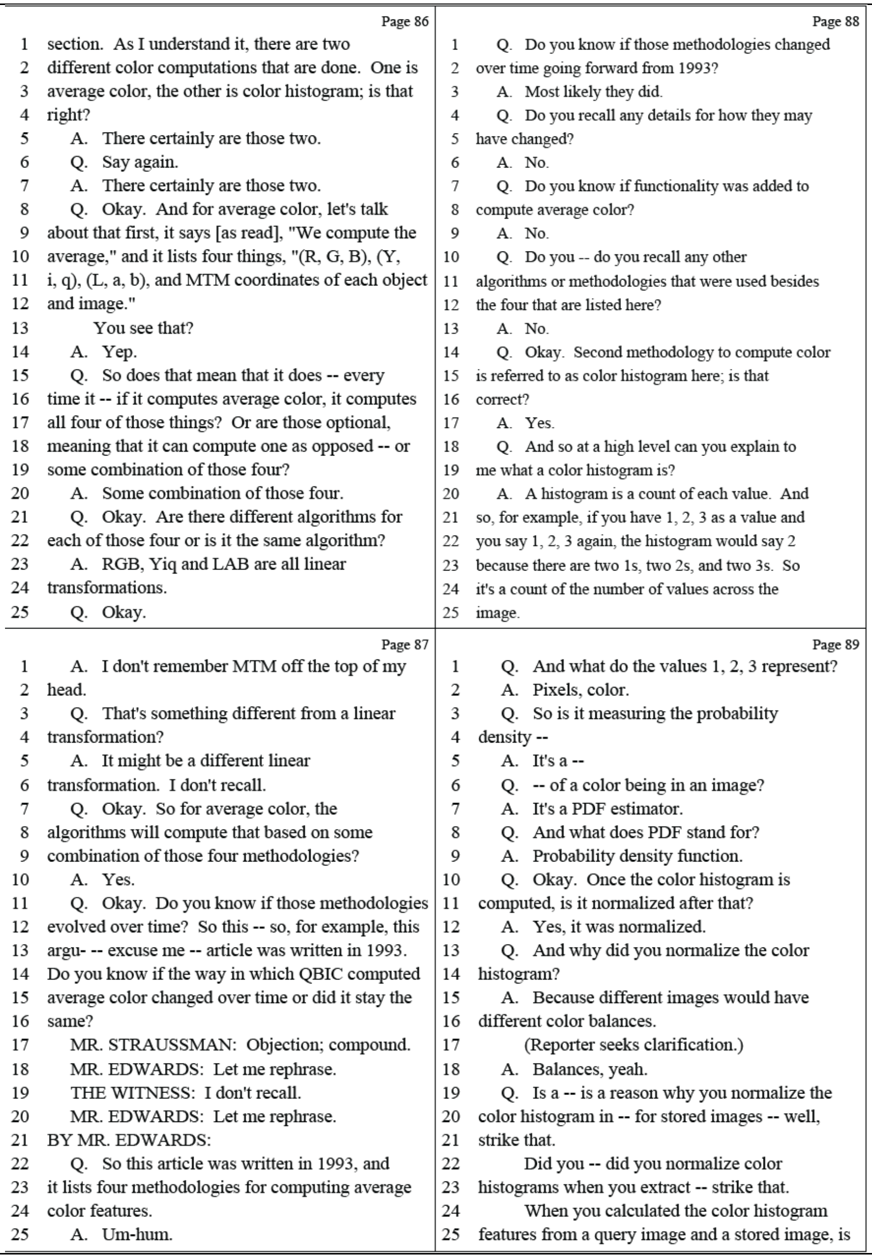

Page 86

section. As I understand it, there are two different color computations that are done. One is average color, the other is color histogram; is that right?
A. There certainly are those two.
Q. Say again.
A. There certainly are those two.
Q. Okay. And for average color, let's talk about that first, it says [as read], "We compute the average," and it lists four things, "(R, G, B), (Y, i, q), (L, a, b), and MTM coordinates of each object and image."
You see that?
A. Yep.
Q. So does that mean that it does -- every time it -- if it computes average color, it computes all four of those things? Or are those optional, meaning that it can compute one as opposed -- or some combination of those four?
A. Some combination of those four.
Q. Okay. Are there different algorithms for each of those four or is it the same algorithm?
A. RGB, Yiq and LAB are all linear transformations.
Q. Okay.

Page 87

A. I don't remember MTM off the top of my head.
Q. That's something different from a linear transformation?
A. It might be a different linear transformation. I don't recall.
Q. Okay. So for average color, the algorithms will compute that based on some combination of those four methodologies?
A. Yes.
Q. Okay. Do you know if those methodologies evolved over time? So this -- so, for example, this argu- -- excuse me -- article was written in 1993. Do you know if the way in which QBIC computed average color changed over time or did it stay the same?
MR. STRAUSSMAN: Objection; compound.
MR. EDWARDS: Let me rephrase.
THE WITNESS: I don't recall.
MR. EDWARDS: Let me rephrase.
BY MR. EDWARDS:
Q. So this article was written in 1993, and it lists four methodologies for computing average color features.
A. Um-hum.

Page 88

Q. Do you know if those methodologies changed over time going forward from 1993?
A. Most likely they did.
Q. Do you recall any details for how they may have changed?
A. No.
Q. Do you know if functionality was added to compute average color?
A. No.
Q. Do you -- do you recall any other algorithms or methodologies that were used besides the four that are listed here?
A. No.
Q. Okay. Second methodology to compute color is referred to as color histogram here; is that correct?
A. Yes.
Q. And so at a high level can you explain to me what a color histogram is?
A. A histogram is a count of each value. And so, for example, if you have 1, 2, 3 as a value and you say 1, 2, 3 again, the histogram would say 2 because there are two 1s, two 2s, and two 3s. So it's a count of the number of values across the image.

Page 89

Q. And what do the values 1, 2, 3 represent?
A. Pixels, color.
Q. So is it measuring the probability density --
A. It's a --
Q. -- of a color being in an image?
A. It's a PDF estimator.
Q. And what does PDF stand for?
A. Probability density function.
Q. Okay. Once the color histogram is computed, is it normalized after that?
A. Yes, it was normalized.
Q. And why did you normalize the color histogram?
A. Because different images would have different color balances.
(Reporter seeks clarification.)
A. Balances, yeah.
Q. Is a -- is a reason why you normalize the color histogram in -- for stored images -- well, strike that.
Did you -- did you normalize color histograms when you extract -- strike that.
When you calculated the color histogram features from a query image and a stored image, is

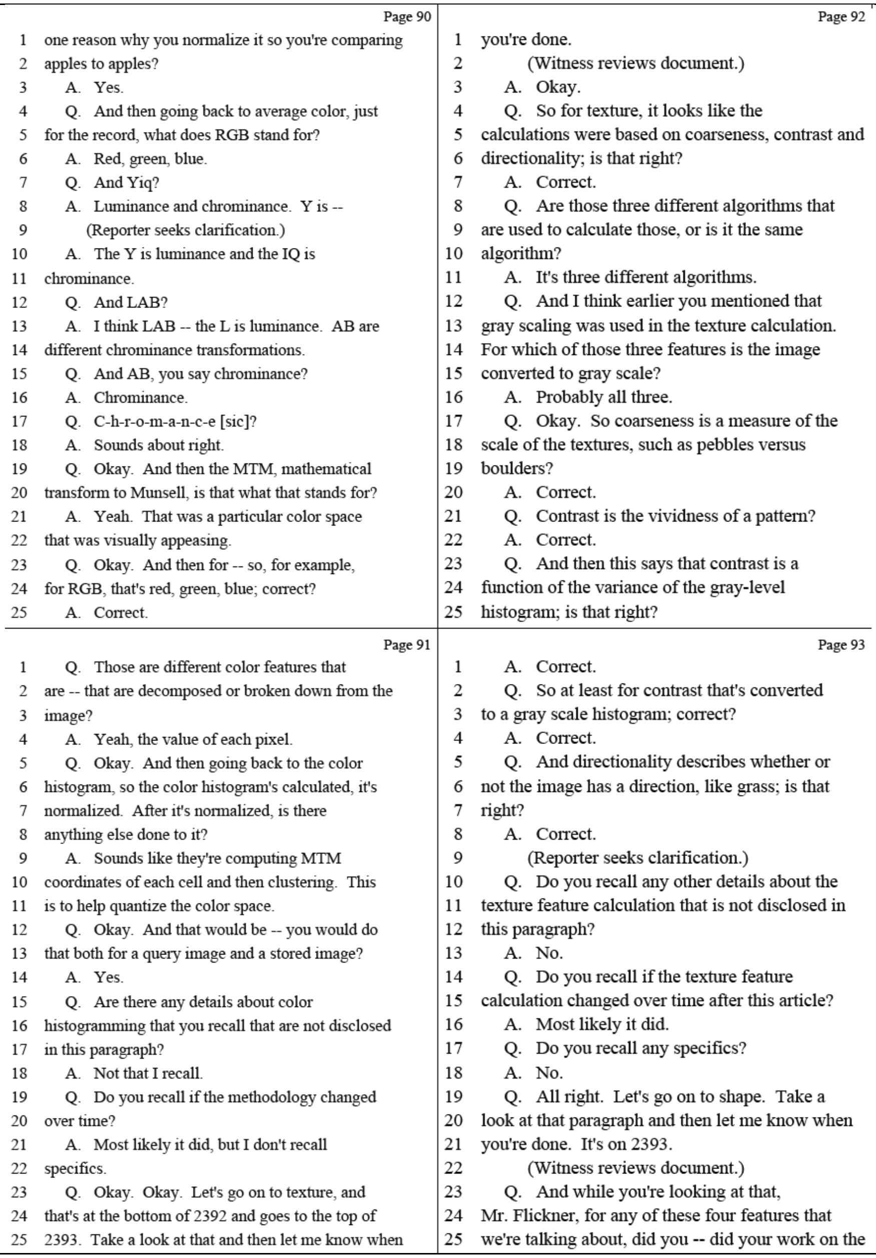

Page 90

one reason why you normalize it so you're comparing apples to apples?
A. Yes.
Q. And then going back to average color, just for the record, what does RGB stand for?
A. Red, green, blue.
Q. And Yiq?
A. Luminance and chrominance. Y is --
(Reporter seeks clarification.)
A. The Y is luminance and the IQ is chrominance.
Q. And LAB?
A. I think LAB -- the L is luminance. AB are different chrominance transformations.
Q. And AB, you say chrominance?
A. Chrominance.
Q. C-h-r-o-m-a-n-c-e [sic]?
A. Sounds about right.
Q. Okay. And then the MTM, mathematical transform to Munsell, is that what that stands for?
A. Yeah. That was a particular color space that was visually appeasing.
Q. Okay. And then for -- so, for example, for RGB, that's red, green, blue; correct?
A. Correct.

Page 91

Q. Those are different color features that are -- that are decomposed or broken down from the image?
A. Yeah, the value of each pixel.
Q. Okay. And then going back to the color histogram, so the color histogram's calculated, it's normalized. After it's normalized, is there anything else done to it?
A. Sounds like they're computing MTM coordinates of each cell and then clustering. This is to help quantize the color space.
Q. Okay. And that would be -- you would do that both for a query image and a stored image?
A. Yes.
Q. Are there any details about color histogramming that you recall that are not disclosed in this paragraph?
A. Not that I recall.
Q. Do you recall if the methodology changed over time?
A. Most likely it did, but I don't recall specifics.
Q. Okay. Okay. Let's go on to texture, and that's at the bottom of 2392 and goes to the top of 2393. Take a look at that and then let me know when

Page 92

you're done.
(Witness reviews document.)
A. Okay.
Q. So for texture, it looks like the calculations were based on coarseness, contrast and directionality; is that right?
A. Correct.
Q. Are those three different algorithms that are used to calculate those, or is it the same algorithm?
A. It's three different algorithms.
Q. And I think earlier you mentioned that gray scaling was used in the texture calculation. For which of those three features is the image converted to gray scale?
A. Probably all three.
Q. Okay. So coarseness is a measure of the scale of the textures, such as pebbles versus boulders?
A. Correct.
Q. Contrast is the vividness of a pattern?
A. Correct.
Q. And then this says that contrast is a function of the variance of the gray-level histogram; is that right?

Page 93

A. Correct.
Q. So at least for contrast that's converted to a gray scale histogram; correct?
A. Correct.
Q. And directionality describes whether or not the image has a direction, like grass; is that right?
A. Correct.
(Reporter seeks clarification.)
Q. Do you recall any other details about the texture feature calculation that is not disclosed in this paragraph?
A. No.
Q. Do you recall if the texture feature calculation changed over time after this article?
A. Most likely it did.
Q. Do you recall any specifics?
A. No.
Q. All right. Let's go on to shape. Take a look at that paragraph and then let me know when you're done. It's on 2393.
(Witness reviews document.)
Q. And while you're looking at that, Mr. Flickner, for any of these four features that we're talking about, did you -- did your work on the

Page 94

QBIC project -- were you more involved in one or more of these features than the other?
A. Yes.
Q. Which one?
A. Shape.
Q. Okay.
(Witness reviews document.)
A. So one thing I can state is originally I couldn't remember the name of the edge operator -- and the name is Canny, the Canny edge operator -- when we were talking about finding edges.
Q. Okay. So this -- the paragraph on IBM 2393 refreshed your memory that shape calculation used Canny edge detection?
A. No, that's sketch.
Q. Oh, I'm sorry.
A. I'm just stating that originally I was asked a question about standard edge operators. I couldn't remember the name. The name was Canny.
Q. Canny. And Canny is used in the sketch calculation?
A. Yeah, it's used in the sketch calculation.
Q. Understood. Okay. Let's go back to shape.
And for shape, is it a first step -- well,

Page 95

back up, strike that.
To calculate the shape feature in a query image or a stored database image, is the first step to convert that image into a binary image?
A. I don't think so.
Q. So if you look at the second paragraph -- and the reason why I ask is -- it's the second sentence in the second paragraph under "Shape features." It says [as read], "All shapes are assumed to be non-occluded planar shapes allowing for each shape to be represented as a binary image."
See that?
A. Yeah.
Q. And then the next sentence says [as read], "The area is computed as the number of pixels set in the binary image."
A. Right.
Q. So that's why I ask is it the first step in calculating shape to convert the image to binary?
A. Yes.
Q. And then once you convert it to binary image, then what do you do next in order to calculate the shape?
A. You're computing something called perround, which is perimeter squared over the area,

Page 96

of the object, of the blob.
Q. Can you say the name of the word again?
A. I think it's called perround.
Q. Per round?
A. Yeah.
Q. Those -- are those two words, p-e-r, r-o-u-n-d?
A. P-e-r -- one word.
Q. Okay. Perround, one word.
A. It's basically the perimeter squared over area.
(Reporter seeks clarification.)
A. Perimeter squared over area.
(Reporter seeks clarification.)
A. Length of the -- over area.
Q. Okay. So let's step back a little bit. So it says -- the second paragraph says [as read], "Our shape features are based on a combination of heuristic" -- h-e-u-r-i-s-t-i-c -- "shape features area, circularity, eccentricity, major access orientation and a set of algebraic moment invariants," i-n-v-a-r-i-a-n-t-s.
Did I read that correctly?
A. Um-hum.
Q. Okay. For circularity, that is computed

Page 97

as perimeter squared over area; correct?
A. Right.
Q. All right. Area is the number of pixels in the binary image?
A. In the binary blob.
Q. In the binary what?
A. Blob.
Q. Blob. And when you say "blob," what does that refer to?
A. The outline of the object that you're doing a shape query on.
Q. And so are each of these, area, circularity, eccentricity, major axis orientation and algebraic moment invariants, are those each different algorithms or is it all the same algorithm?
A. It's different algorithms.
Q. Okay. And since you worked on the shape feature the most, do you recall any other details that are not disclosed in this paragraph?
A. Well, this evolved and we ended up doing an implementation of something that's not in here.
Q. Okay. And what was that implementation?
A. Something called turning angle per round, turning angle, which is -- it looks at the tangent

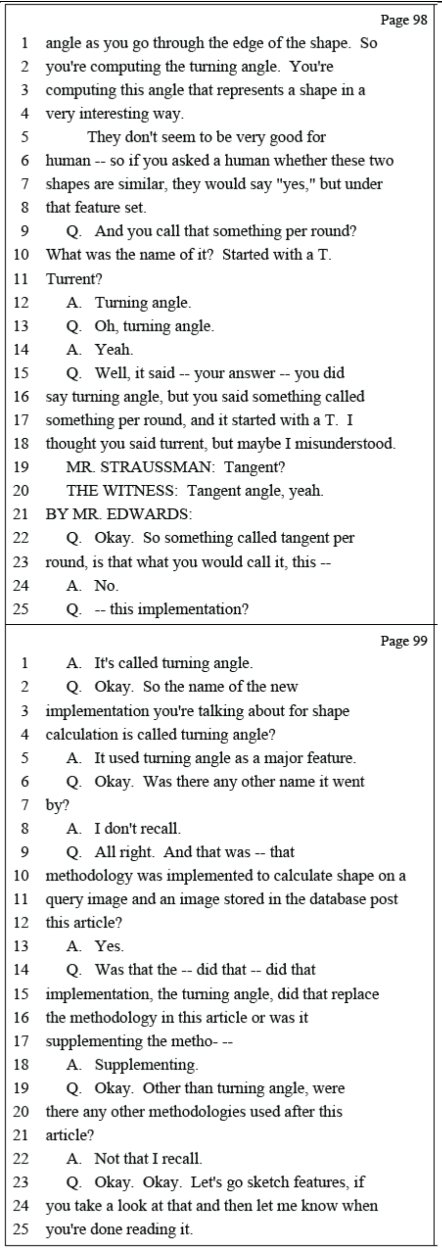

Page 98

angle as you go through the edge of the shape. So you're computing the turning angle. You're computing this angle that represents a shape in a very interesting way.

They don't seem to be very good for human -- so if you asked a human whether these two shapes are similar, they would say "yes," but under that feature set.

Q. And you call that something per round? What was the name of it? Started with a T. Turrent?

A. Turning angle.

Q. Oh, turning angle.

A. Yeah.

Q. Well, it said -- your answer -- you did say turning angle, but you said something called something per round, and it started with a T. I thought you said turrent, but maybe I misunderstood.

MR. STRAUSSMAN: Tangent?

THE WITNESS: Tangent angle, yeah.

BY MR. EDWARDS:

Q. Okay. So something called tangent per round, is that what you would call it, this --

A. No.

Q. -- this implementation?

Page 99

A. It's called turning angle.

Q. Okay. So the name of the new implementation you're talking about for shape calculation is called turning angle?

A. It used turning angle as a major feature.

Q. Okay. Was there any other name it went by?

A. I don't recall.

Q. All right. And that was -- that methodology was implemented to calculate shape on a query image and an image stored in the database post this article?

A. Yes.

Q. Was that the -- did that -- did that implementation, the turning angle, did that replace the methodology in this article or was it supplementing the metho- --

A. Supplementing.

Q. Okay. Other than turning angle, were there any other methodologies used after this article?

A. Not that I recall.

Q. Okay. Okay. Let's go sketch features, if you take a look at that and then let me know when you're done reading it.

Element	'004 Claim Recitation	Exemplary Citations in Reference
		Id. at 119–128: 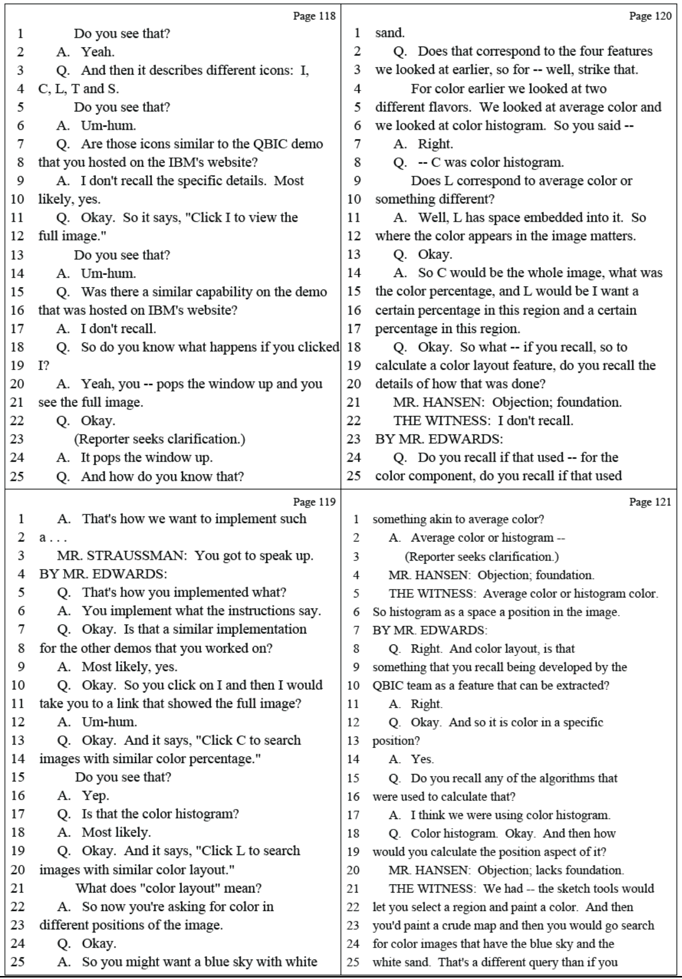

Page 118

Do you see that?
A. Yeah.
Q. And then it describes different icons: I, C, L, T and S.
Do you see that?
A. Um-hum.
Q. Are those icons similar to the QBIC demo that you hosted on the IBM's website?
A. I don't recall the specific details. Most likely, yes.
Q. Okay. So it says, "Click I to view the full image."
Do you see that?
A. Um-hum.
Q. Was there a similar capability on the demo that was hosted on IBM's website?
A. I don't recall.
Q. So do you know what happens if you clicked I?
A. Yeah, you -- pops the window up and you see the full image.
Q. Okay.
(Reporter seeks clarification.)
A. It pops the window up.
Q. And how do you know that?

Page 119

A. That's how we want to implement such a . . .
MR. STRAUSSMAN: You got to speak up.
BY MR. EDWARDS:
Q. That's how you implemented what?
A. You implement what the instructions say.
Q. Okay. Is that a similar implementation for the other demos that you worked on?
A. Most likely, yes.
Q. Okay. So you click on I and then I would take you to a link that showed the full image?
A. Um-hum.
Q. Okay. And it says, "Click C to search images with similar color percentage."
Do you see that?
A. Yep.
Q. Is that the color histogram?
A. Most likely.
Q. Okay. And it says, "Click L to search images with similar color layout."
What does "color layout" mean?
A. So now you're asking for color in different positions of the image.
Q. Okay.
A. So you might want a blue sky with white

Page 120

sand.
Q. Does that correspond to the four features we looked at earlier, so for -- well, strike that.
For color earlier we looked at two different flavors. We looked at average color and we looked at color histogram. So you said --
A. Right.
Q. -- C was color histogram.
Does L correspond to average color or something different?
A. Well, L has space embedded into it. So where the color appears in the image matters.
Q. Okay.
A. So C would be the whole image, what was the color percentage, and L would be I want a certain percentage in this region and a certain percentage in this region.
Q. Okay. So what -- if you recall, so to calculate a color layout feature, do you recall the details of how that was done?
MR. HANSEN: Objection; foundation.
THE WITNESS: I don't recall.
BY MR. EDWARDS:
Q. Do you recall if that used -- for the color component, do you recall if that used

Page 121

something akin to average color?
A. Average color or histogram --
(Reporter seeks clarification.)
MR. HANSEN: Objection; foundation.
THE WITNESS: Average color or histogram color. So histogram as a space a position in the image.
BY MR. EDWARDS:
Q. Right. And color layout, is that something that you recall being developed by the QBIC team as a feature that can be extracted?
A. Right.
Q. Okay. And so it is color in a specific position?
A. Yes.
Q. Do you recall any of the algorithms that were used to calculate that?
A. I think we were using color histogram.
Q. Color histogram. Okay. And then how would you calculate the position aspect of it?
MR. HANSEN: Objection; lacks foundation.
THE WITNESS: We had -- the sketch tools would let you select a region and paint a color. And then you'd paint a crude map and then you would go search for color images that have the blue sky and the white sand. That's a different query than if you

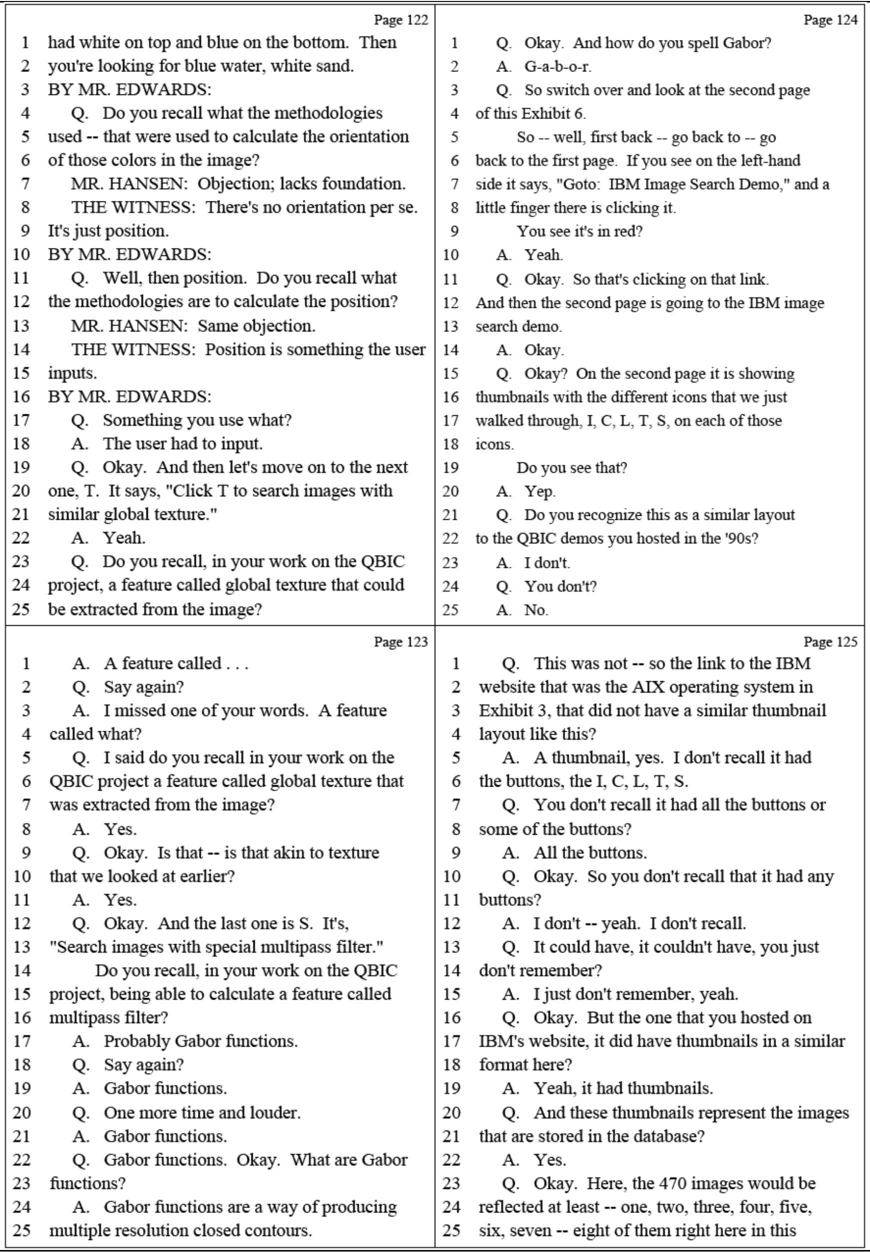

Page 122

had white on top and blue on the bottom. Then you're looking for blue water, white sand.
BY MR. EDWARDS:
Q. Do you recall what the methodologies used -- that were used to calculate the orientation of those colors in the image?
MR. HANSEN: Objection; lacks foundation.
THE WITNESS: There's no orientation per se. It's just position.
BY MR. EDWARDS:
Q. Well, then position. Do you recall what the methodologies are to calculate the position?
MR. HANSEN: Same objection.
THE WITNESS: Position is something the user inputs.
BY MR. EDWARDS:
Q. Something you use what?
A. The user had to input.
Q. Okay. And then let's move on to the next one, T. It says, "Click T to search images with similar global texture."
A. Yeah.
Q. Do you recall, in your work on the QBIC project, a feature called global texture that could be extracted from the image?

Page 123

A. A feature called . . .
Q. Say again?
A. I missed one of your words. A feature called what?
Q. I said do you recall in your work on the QBIC project a feature called global texture that was extracted from the image?
A. Yes.
Q. Okay. Is that -- is that akin to texture that we looked at earlier?
A. Yes.
Q. Okay. And the last one is S. It's, "Search images with special multipass filter."
Do you recall, in your work on the QBIC project, being able to calculate a feature called multipass filter?
A. Probably Gabor functions.
Q. Say again?
A. Gabor functions.
Q. One more time and louder.
A. Gabor functions.
Q. Gabor functions. Okay. What are Gabor functions?
A. Gabor functions are a way of producing multiple resolution closed contours.

Page 124

Q. Okay. And how do you spell Gabor?
A. G-a-b-o-r.
Q. So switch over and look at the second page of this Exhibit 6.
So -- well, first back -- go back to -- go back to the first page. If you see on the left-hand side it says, "Goto: IBM Image Search Demo," and a little finger there is clicking it.
You see it's in red?
A. Yeah.
Q. Okay. So that's clicking on that link. And then the second page is going to the IBM image search demo.
A. Okay.
Q. Okay? On the second page it is showing thumbnails with the different icons that we just walked through, I, C, L, T, S, on each of those icons.
Do you see that?
A. Yep.
Q. Do you recognize this as a similar layout to the QBIC demos you hosted in the '90s?
A. I don't.
Q. You don't?
A. No.

Page 125

Q. This was not -- so the link to the IBM website that was the AIX operating system in Exhibit 3, that did not have a similar thumbnail layout like this?
A. A thumbnail, yes. I don't recall it had the buttons, the I, C, L, T, S.
Q. You don't recall it had all the buttons or some of the buttons?
A. All the buttons.
Q. Okay. So you don't recall that it had any buttons?
A. I don't -- yeah. I don't recall.
Q. It could have, it couldn't have, you just don't remember?
A. I just don't remember, yeah.
Q. Okay. But the one that you hosted on IBM's website, it did have thumbnails in a similar format here?
A. Yeah, it had thumbnails.
Q. And these thumbnails represent the images that are stored in the database?
A. Yes.
Q. Okay. Here, the 470 images would be reflected at least -- one, two, three, four, five, six, seven -- eight of them right here in this

Element	'004 Claim Recitation	Exemplary Citations in Reference
		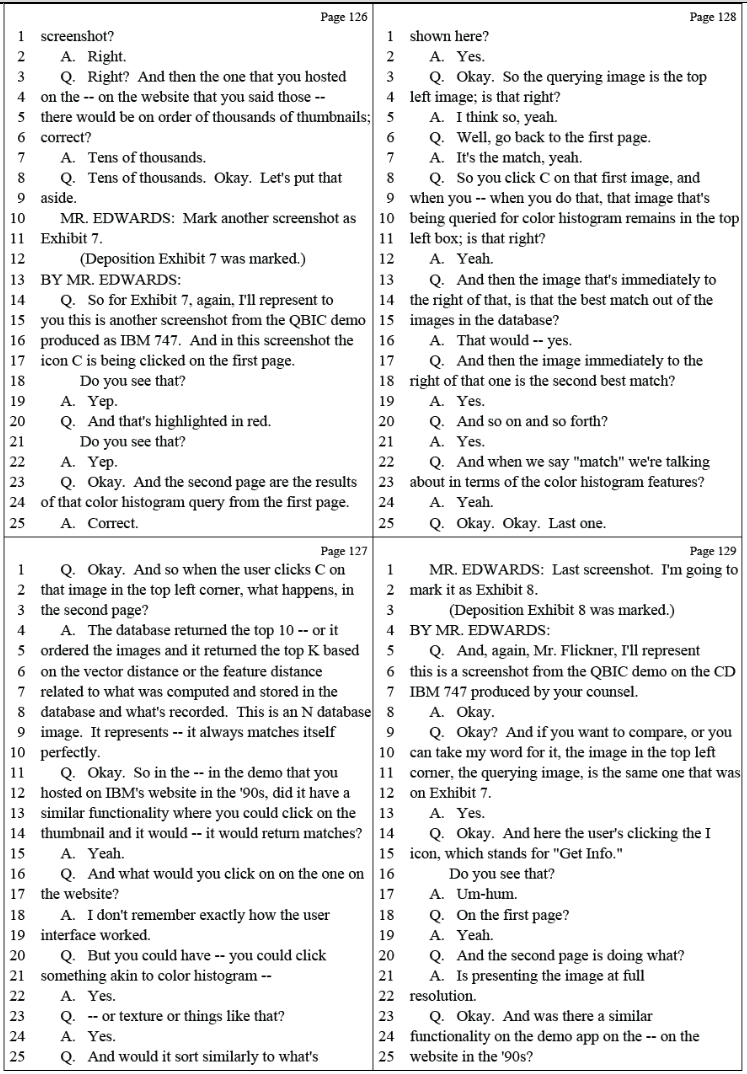 Page 126 screenshot? A. Right. Q. Right? And then the one that you hosted on the -- on the website that you said those -- there would be on order of thousands of thumbnails; correct? A. Tens of thousands. Q. Tens of thousands. Okay. Let's put that aside. MR. EDWARDS: Mark another screenshot as Exhibit 7. (Deposition Exhibit 7 was marked.) BY MR. EDWARDS: Q. So for Exhibit 7, again, I'll represent to you this is another screenshot from the QBIC demo produced as IBM 747. And in this screenshot the icon C is being clicked on the first page. Do you see that? A. Yep. Q. And that's highlighted in red. Do you see that? A. Yep. Q. Okay. And the second page are the results of that color histogram query from the first page. A. Correct. Page 127 Q. Okay. And so when the user clicks C on that image in the top left corner, what happens, in the second page? A. The database returned the top 10 -- or it ordered the images and it returned the top K based on the vector distance or the feature distance related to what was computed and stored in the database and what's recorded. This is an N database image. It represents -- it always matches itself perfectly. Q. Okay. So in the -- in the demo that you hosted on IBM's website in the '90s, did it have a similar functionality where you could click on the thumbnail and it would -- it would return matches? A. Yeah. Q. And what would you click on on the one on the website? A. I don't remember exactly how the user interface worked. Q. But you could have -- you could click something akin to color histogram -- A. Yes. Q. -- or texture or things like that? A. Yes. Q. And would it sort similarly to what's Page 128 shown here? A. Yes. Q. Okay. So the querying image is the top left image; is that right? A. I think so, yeah. Q. Well, go back to the first page. A. It's the match, yeah. Q. So you click C on that first image, and when you -- when you do that, that image that's being queried for color histogram remains in the top left box; is that right? A. Yeah. Q. And then the image that's immediately to the right of that, is that the best match out of the images in the database? A. That would -- yes. Q. And then the image immediately to the right of that one is the second best match? A. Yes. Q. And so on and so forth? A. Yes. Q. And when we say "match" we're talking about in terms of the color histogram features? A. Yeah. Q. Okay. Okay. Last one. Page 129 MR. EDWARDS: Last screenshot. I'm going to mark it as Exhibit 8. (Deposition Exhibit 8 was marked.) BY MR. EDWARDS: Q. And, again, Mr. Flickner, I'll represent this is a screenshot from the QBIC demo on the CD IBM 747 produced by your counsel. A. Okay. Q. Okay? And if you want to compare, or you can take my word for it, the image in the top left corner, the querying image, is the same one that was on Exhibit 7. A. Yes. Q. Okay. And here the user's clicking the I icon, which stands for "Get Info." Do you see that? A. Um-hum. Q. On the first page? A. Yeah. Q. And the second page is doing what? A. Is presenting the image at full resolution. Q. Okay. And was there a similar functionality on the demo app on the -- on the website in the '90s? *Id*. at 135 – 136

Element	'004 Claim Recitation	Exemplary Citations in Reference
		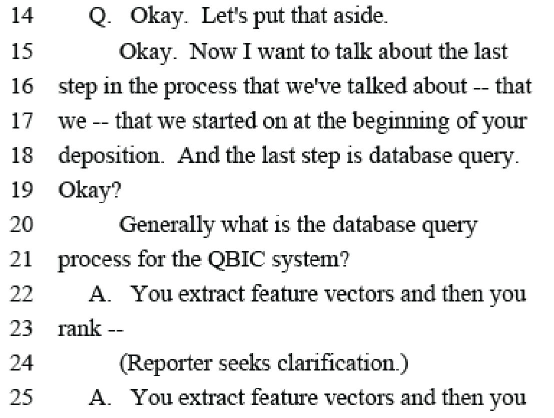 Q. Okay. Let's put that aside. Okay. Now I want to talk about the last step in the process that we've talked about -- that we -- that we started on at the beginning of your deposition. And the last step is database query. Okay? Generally what is the database query process for the QBIC system? A. You extract feature vectors and then you rank -- (Reporter seeks clarification.) A. You extract feature vectors and then you 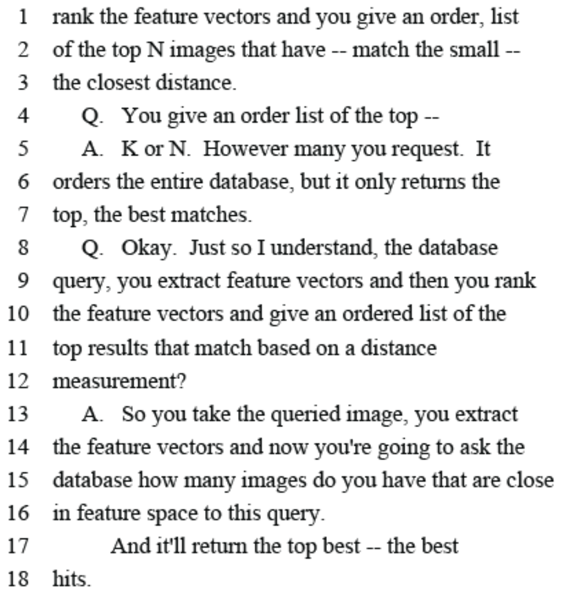 rank the feature vectors and you give an order, list of the top N images that have -- match the small -- the closest distance. Q. You give an order list of the top -- A. K or N. However many you request. It orders the entire database, but it only returns the top, the best matches. Q. Okay. Just so I understand, the database query, you extract feature vectors and then you rank the feature vectors and give an ordered list of the top results that match based on a distance measurement? A. So you take the queried image, you extract the feature vectors and now you're going to ask the database how many images do you have that are close in feature space to this query. And it'll return the top best -- the best hits. 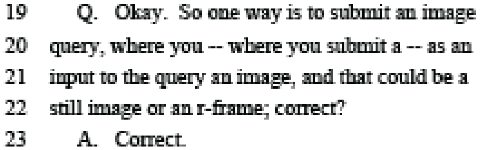 Q. Okay. So one way is to submit an image query, where you -- where you submit a -- as an input to the query an image, and that could be a still image or an r-frame; correct? A. Correct.

Element	'004 Claim Recitation	Exemplary Citations in Reference
		Id. at 139–165 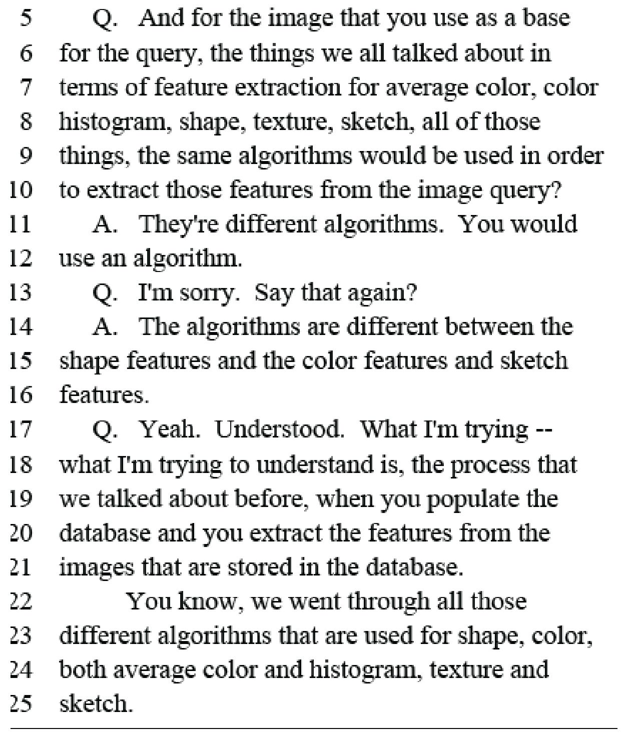 Q. And for the image that you use as a base for the query, the things we all talked about in terms of feature extraction for average color, color histogram, shape, texture, sketch, all of those things, the same algorithms would be used in order to extract those features from the image query? A. They're different algorithms. You would use an algorithm. Q. I'm sorry. Say that again? A. The algorithms are different between the shape features and the color features and sketch features. Q. Yeah. Understood. What I'm trying -- what I'm trying to understand is, the process that we talked about before, when you populate the database and you extract the features from the images that are stored in the database. You know, we went through all those different algorithms that are used for shape, color, both average color and histogram, texture and sketch.

Element	'004 Claim Recitation	Exemplary Citations in Reference
		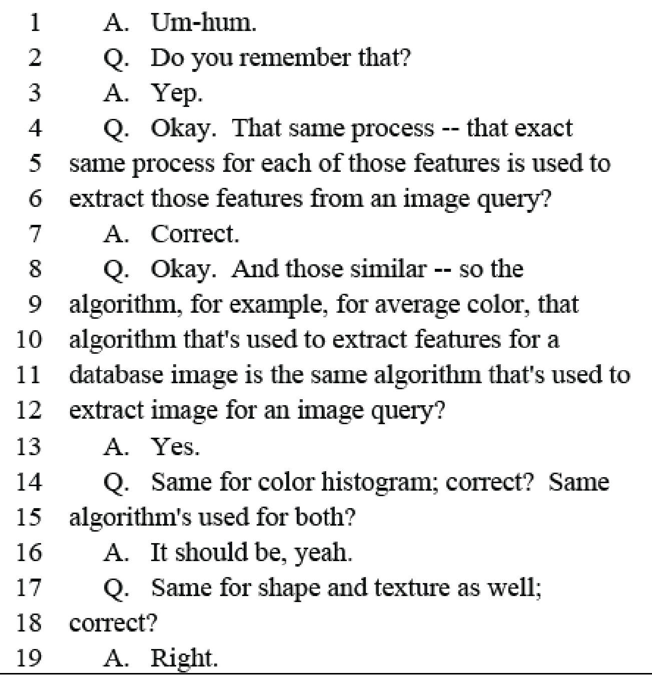 A. Um-hum. Q. Do you remember that? A. Yep. Q. Okay. That same process -- that exact same process for each of those features is used to extract those features from an image query? A. Correct. Q. Okay. And those similar -- so the algorithm, for example, for average color, that algorithm that's used to extract features for a database image is the same algorithm that's used to extract image for an image query? A. Yes. Q. Same for color histogram; correct? Same algorithm's used for both? A. It should be, yeah. Q. Same for shape and texture as well; correct? A. Right.

Element	'004 Claim Recitation	Exemplary Citations in Reference
		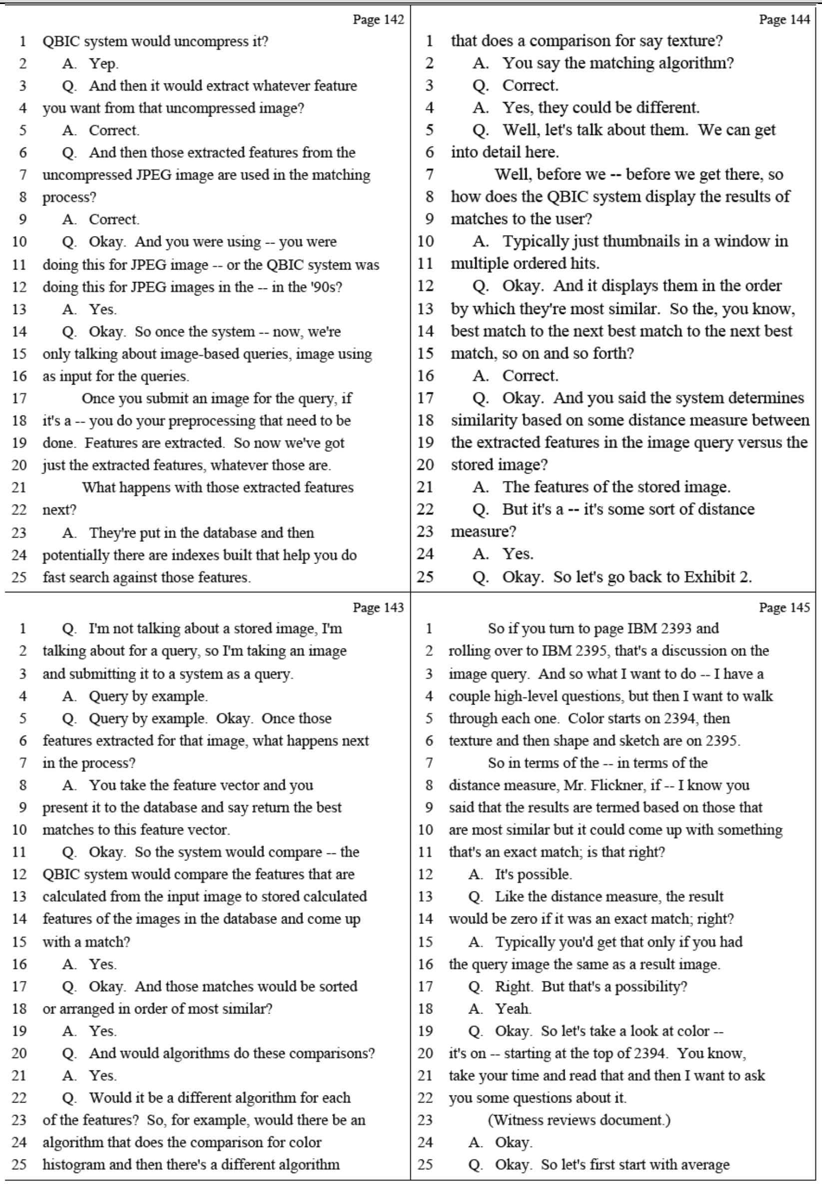 Page 142 QBIC system would uncompress it? A. Yep. Q. And then it would extract whatever feature you want from that uncompressed image? A. Correct. Q. And then those extracted features from the uncompressed JPEG image are used in the matching process? A. Correct. Q. Okay. And you were using -- you were doing this for JPEG image -- or the QBIC system was doing this for JPEG images in the -- in the '90s? A. Yes. Q. Okay. So once the system -- now, we're only talking about image-based queries, image using as input for the queries. Once you submit an image for the query, if it's a -- you do your preprocessing that need to be done. Features are extracted. So now we've got just the extracted features, whatever those are. What happens with those extracted features next? A. They're put in the database and then potentially there are indexes built that help you do fast search against those features. Page 143 Q. I'm not talking about a stored image, I'm talking about for a query, so I'm taking an image and submitting it to a system as a query. A. Query by example. Q. Query by example. Okay. Once those features extracted for that image, what happens next in the process? A. You take the feature vector and you present it to the database and say return the best matches to this feature vector. Q. Okay. So the system would compare -- the QBIC system would compare the features that are calculated from the input image to stored calculated features of the images in the database and come up with a match? A. Yes. Q. Okay. And those matches would be sorted or arranged in order of most similar? A. Yes. Q. And would algorithms do these comparisons? A. Yes. Q. Would it be a different algorithm for each of the features? So, for example, would there be an algorithm that does the comparison for color histogram and then there's a different algorithm Page 144 that does a comparison for say texture? A. You say the matching algorithm? Q. Correct. A. Yes, they could be different. Q. Well, let's talk about them. We can get into detail here. Well, before we -- before we get there, so how does the QBIC system display the results of matches to the user? A. Typically just thumbnails in a window in multiple ordered hits. Q. Okay. And it displays them in the order by which they're most similar. So the, you know, best match to the next best match to the next best match, so on and so forth? A. Correct. Q. Okay. And you said the system determines similarity based on some distance measure between the extracted features in the image query versus the stored image? A. The features of the stored image. Q. But it's a -- it's some sort of distance measure? A. Yes. Q. Okay. So let's go back to Exhibit 2. Page 145 So if you turn to page IBM 2393 and rolling over to IBM 2395, that's a discussion on the image query. And so what I want to do -- I have a couple high-level questions, but then I want to walk through each one. Color starts on 2394, then texture and then shape and sketch are on 2395. So in terms of the -- in terms of the distance measure, Mr. Flickner, if -- I know you said that the results are termed based on those that are most similar but it could come up with something that's an exact match; is that right? A. It's possible. Q. Like the distance measure, the result would be zero if it was an exact match; right? A. Typically you'd get that only if you had the query image the same as a result image. Q. Right. But that's a possibility? A. Yeah. Q. Okay. So let's take a look at color -- it's on -- starting at the top of 2394. You know, take your time and read that and then I want to ask you some questions about it. (Witness reviews document.) A. Okay. Q. Okay. So let's first start with average

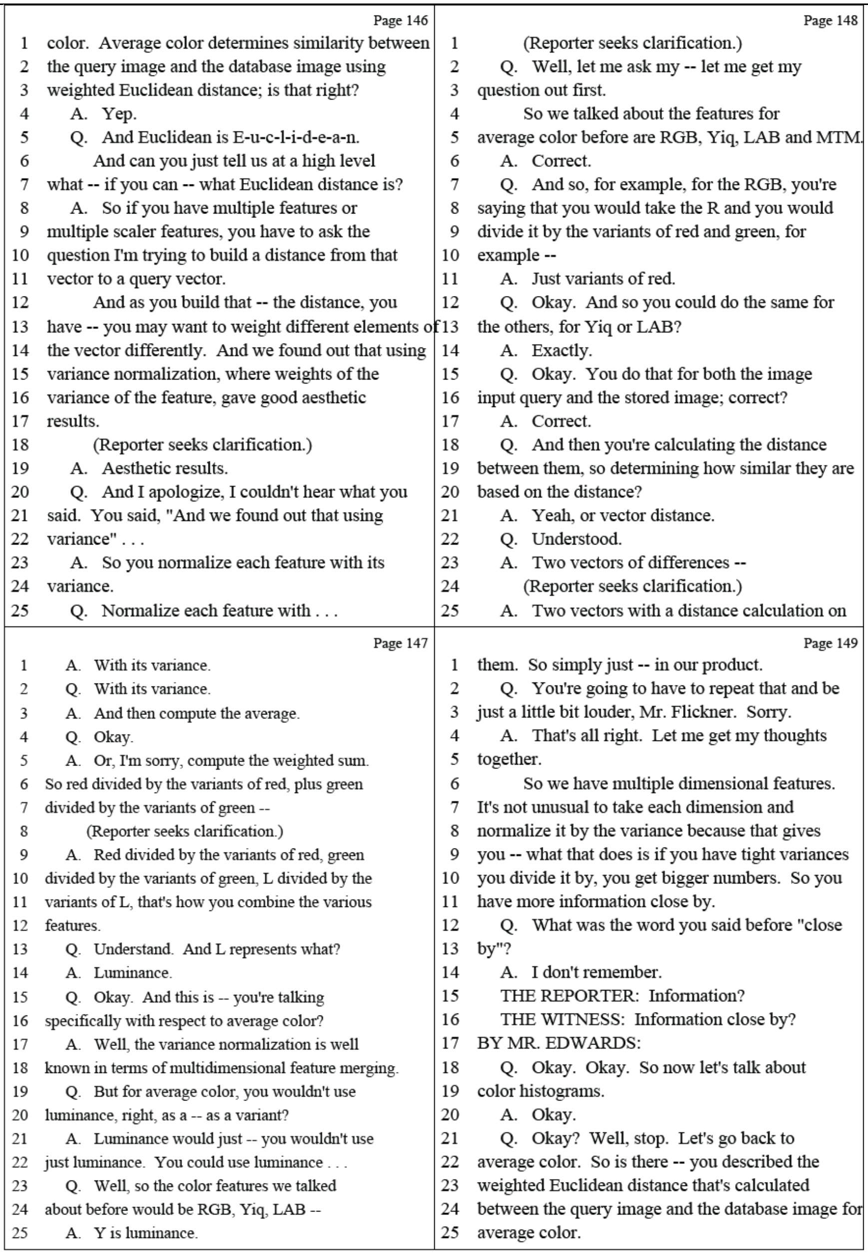

Page 146

color. Average color determines similarity between the query image and the database image using weighted Euclidean distance; is that right?
A. Yep.
Q. And Euclidean is E-u-c-l-i-d-e-a-n.
And can you just tell us at a high level what -- if you can -- what Euclidean distance is?
A. So if you have multiple features or multiple scaler features, you have to ask the question I'm trying to build a distance from that vector to a query vector.
And as you build that -- the distance, you have -- you may want to weight different elements of the vector differently. And we found out that using variance normalization, where weights of the variance of the feature, gave good aesthetic results.
(Reporter seeks clarification.)
A. Aesthetic results.
Q. And I apologize, I couldn't hear what you said. You said, "And we found out that using variance" . . .
A. So you normalize each feature with its variance.
Q. Normalize each feature with . . .

Page 147

A. With its variance.
Q. With its variance.
A. And then compute the average.
Q. Okay.
A. Or, I'm sorry, compute the weighted sum. So red divided by the variants of red, plus green divided by the variants of green --
(Reporter seeks clarification.)
A. Red divided by the variants of red, green divided by the variants of green, L divided by the variants of L, that's how you combine the various features.
Q. Understand. And L represents what?
A. Luminance.
Q. Okay. And this is -- you're talking specifically with respect to average color?
A. Well, the variance normalization is well known in terms of multidimensional feature merging.
Q. But for average color, you wouldn't use luminance, right, as a -- as a variant?
A. Luminance would just -- you wouldn't use just luminance. You could use luminance . . .
Q. Well, so the color features we talked about before would be RGB, Yiq, LAB --
A. Y is luminance.

Page 148

(Reporter seeks clarification.)
Q. Well, let me ask my -- let me get my question out first.
So we talked about the features for average color before are RGB, Yiq, LAB and MTM.
A. Correct.
Q. And so, for example, for the RGB, you're saying that you would take the R and you would divide it by the variants of red and green, for example --
A. Just variants of red.
Q. Okay. And so you could do the same for the others, for Yiq or LAB?
A. Exactly.
Q. Okay. You do that for both the image input query and the stored image; correct?
A. Correct.
Q. And then you're calculating the distance between them, so determining how similar they are based on the distance?
A. Yeah, or vector distance.
Q. Understood.
A. Two vectors of differences --
(Reporter seeks clarification.)
A. Two vectors with a distance calculation on

Page 149

them. So simply just -- in our product.
Q. You're going to have to repeat that and be just a little bit louder, Mr. Flickner. Sorry.
A. That's all right. Let me get my thoughts together.
So we have multiple dimensional features. It's not unusual to take each dimension and normalize it by the variance because that gives you -- what that does is if you have tight variances you divide it by, you get bigger numbers. So you have more information close by.
Q. What was the word you said before "close by"?
A. I don't remember.
THE REPORTER: Information?
THE WITNESS: Information close by?
BY MR. EDWARDS:
Q. Okay. Okay. So now let's talk about color histograms.
A. Okay.
Q. Okay? Well, stop. Let's go back to average color. So is there -- you described the weighted Euclidean distance that's calculated between the query image and the database image for average color.

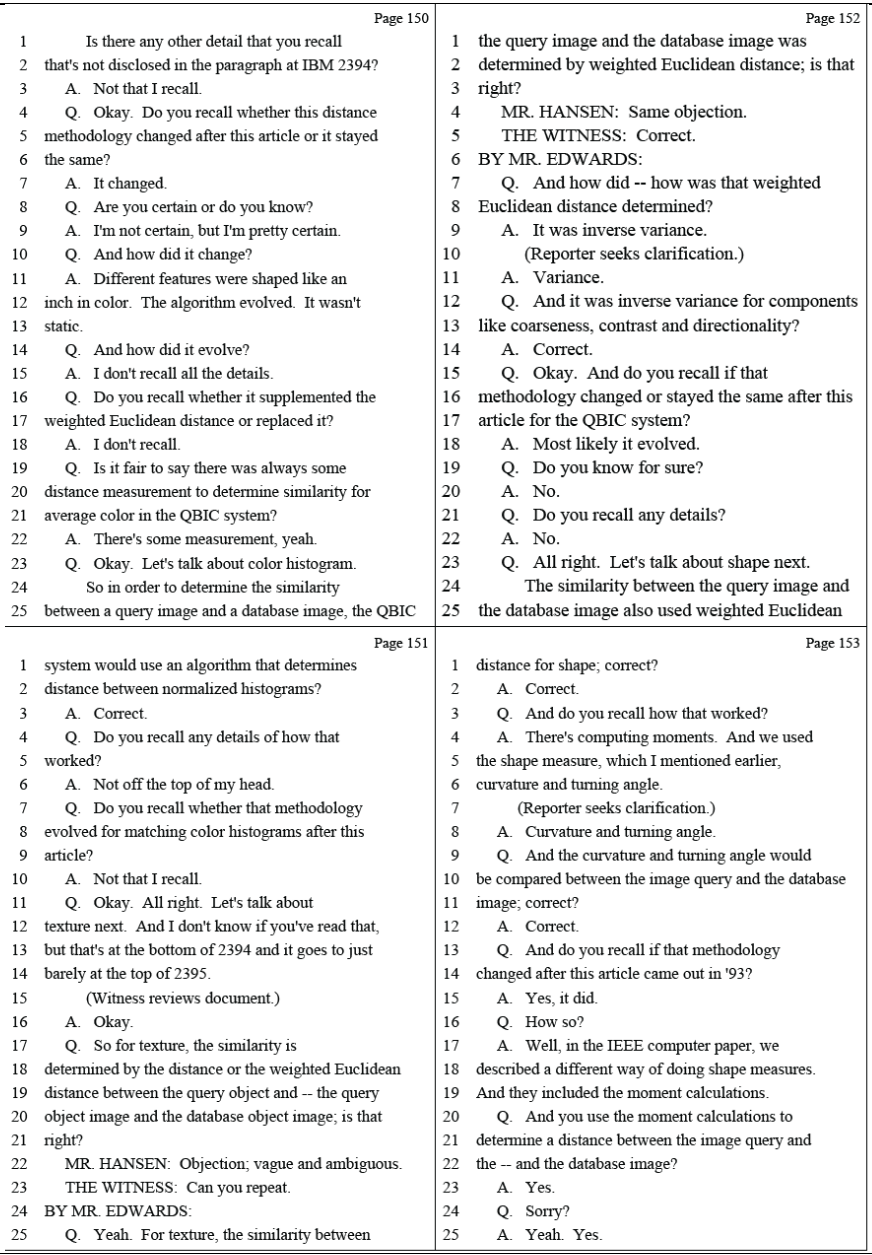

Page 150

Is there any other detail that you recall that's not disclosed in the paragraph at IBM 2394?

A. Not that I recall.

Q. Okay. Do you recall whether this distance methodology changed after this article or it stayed the same?

A. It changed.

Q. Are you certain or do you know?

A. I'm not certain, but I'm pretty certain.

Q. And how did it change?

A. Different features were shaped like an inch in color. The algorithm evolved. It wasn't static.

Q. And how did it evolve?

A. I don't recall all the details.

Q. Do you recall whether it supplemented the weighted Euclidean distance or replaced it?

A. I don't recall.

Q. Is it fair to say there was always some distance measurement to determine similarity for average color in the QBIC system?

A. There's some measurement, yeah.

Q. Okay. Let's talk about color histogram.

So in order to determine the similarity between a query image and a database image, the QBIC

Page 151

system would use an algorithm that determines distance between normalized histograms?

A. Correct.

Q. Do you recall any details of how that worked?

A. Not off the top of my head.

Q. Do you recall whether that methodology evolved for matching color histograms after this article?

A. Not that I recall.

Q. Okay. All right. Let's talk about texture next. And I don't know if you've read that, but that's at the bottom of 2394 and it goes to just barely at the top of 2395.

(Witness reviews document.)

A. Okay.

Q. So for texture, the similarity is determined by the distance or the weighted Euclidean distance between the query object and -- the query object image and the database object image; is that right?

MR. HANSEN: Objection; vague and ambiguous.

THE WITNESS: Can you repeat.

BY MR. EDWARDS:

Q. Yeah. For texture, the similarity between

Page 152

the query image and the database image was determined by weighted Euclidean distance; is that right?

MR. HANSEN: Same objection.

THE WITNESS: Correct.

BY MR. EDWARDS:

Q. And how did -- how was that weighted Euclidean distance determined?

A. It was inverse variance.

(Reporter seeks clarification.)

A. Variance.

Q. And it was inverse variance for components like coarseness, contrast and directionality?

A. Correct.

Q. Okay. And do you recall if that methodology changed or stayed the same after this article for the QBIC system?

A. Most likely it evolved.

Q. Do you know for sure?

A. No.

Q. Do you recall any details?

A. No.

Q. All right. Let's talk about shape next.

The similarity between the query image and the database image also used weighted Euclidean

Page 153

distance for shape; correct?

A. Correct.

Q. And do you recall how that worked?

A. There's computing moments. And we used the shape measure, which I mentioned earlier, curvature and turning angle.

(Reporter seeks clarification.)

A. Curvature and turning angle.

Q. And the curvature and turning angle would be compared between the image query and the database image; correct?

A. Correct.

Q. And do you recall if that methodology changed after this article came out in '93?

A. Yes, it did.

Q. How so?

A. Well, in the IEEE computer paper, we described a different way of doing shape measures. And they included the moment calculations.

Q. And you use the moment calculations to determine a distance between the image query and the -- and the database image?

A. Yes.

Q. Sorry?

A. Yeah. Yes.

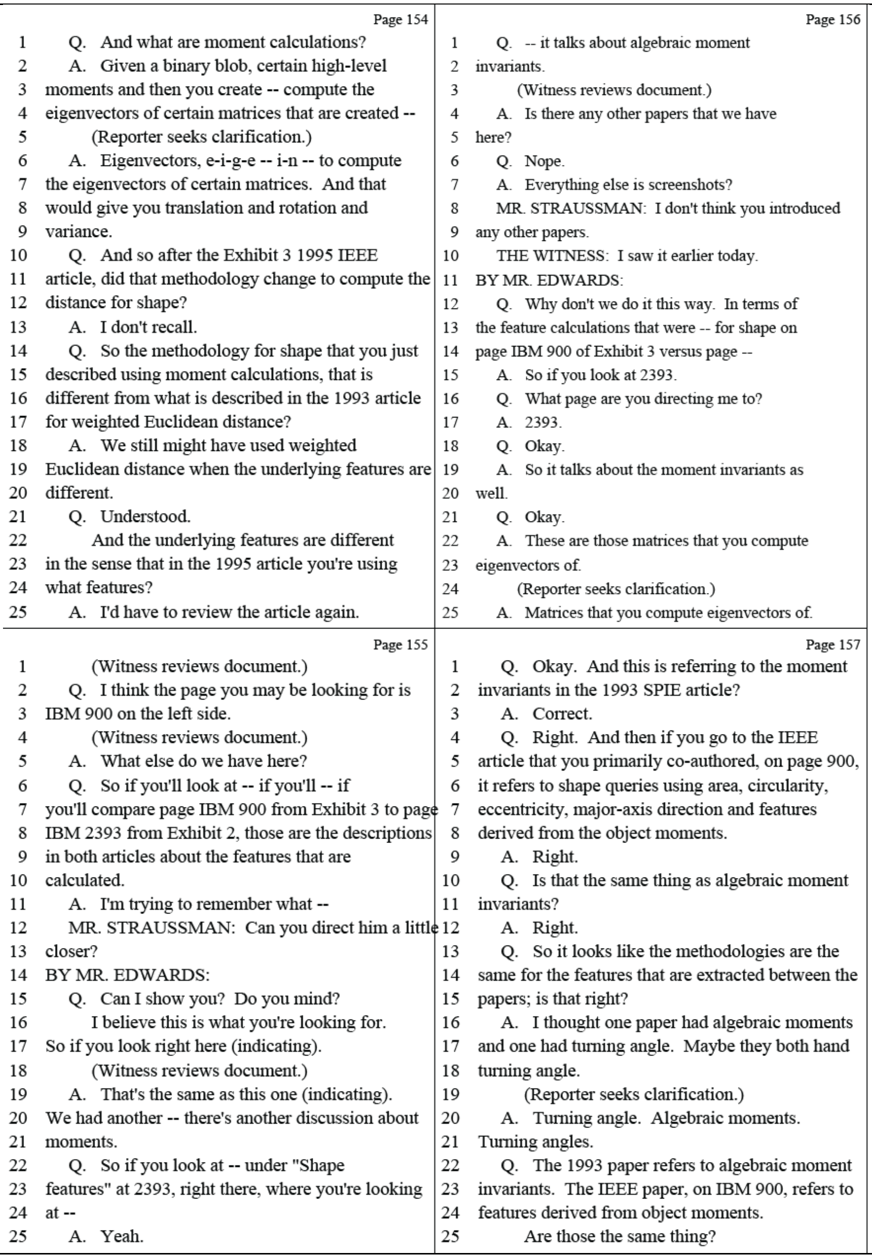

Page 154

Q. And what are moment calculations?
A. Given a binary blob, certain high-level moments and then you create -- compute the eigenvectors of certain matrices that are created --
(Reporter seeks clarification.)
A. Eigenvectors, e-i-g-e -- i-n -- to compute the eigenvectors of certain matrices. And that would give you translation and rotation and variance.
Q. And so after the Exhibit 3 1995 IEEE article, did that methodology change to compute the distance for shape?
A. I don't recall.
Q. So the methodology for shape that you just described using moment calculations, that is different from what is described in the 1993 article for weighted Euclidean distance?
A. We still might have used weighted Euclidean distance when the underlying features are different.
Q. Understood.
And the underlying features are different in the sense that in the 1995 article you're using what features?
A. I'd have to review the article again.

Page 155

(Witness reviews document.)
Q. I think the page you may be looking for is IBM 900 on the left side.
(Witness reviews document.)
A. What else do we have here?
Q. So if you'll look at -- if you'll -- if you'll compare page IBM 900 from Exhibit 3 to page IBM 2393 from Exhibit 2, those are the descriptions in both articles about the features that are calculated.
A. I'm trying to remember what --
MR. STRAUSSMAN: Can you direct him a little closer?
BY MR. EDWARDS:
Q. Can I show you? Do you mind?
I believe this is what you're looking for. So if you look right here (indicating).
(Witness reviews document.)
A. That's the same as this one (indicating). We had another -- there's another discussion about moments.
Q. So if you look at -- under "Shape features" at 2393, right there, where you're looking at --
A. Yeah.

Page 156

Q. -- it talks about algebraic moment invariants.
(Witness reviews document.)
A. Is there any other papers that we have here?
Q. Nope.
A. Everything else is screenshots?
MR. STRAUSSMAN: I don't think you introduced any other papers.
THE WITNESS: I saw it earlier today.
BY MR. EDWARDS:
Q. Why don't we do it this way. In terms of the feature calculations that were -- for shape on page IBM 900 of Exhibit 3 versus page --
A. So if you look at 2393.
Q. What page are you directing me to?
A. 2393.
Q. Okay.
A. So it talks about the moment invariants as well.
Q. Okay.
A. These are those matrices that you compute eigenvectors of.
(Reporter seeks clarification.)
A. Matrices that you compute eigenvectors of.

Page 157

Q. Okay. And this is referring to the moment invariants in the 1993 SPIE article?
A. Correct.
Q. Right. And then if you go to the IEEE article that you primarily co-authored, on page 900, it refers to shape queries using area, circularity, eccentricity, major-axis direction and features derived from the object moments.
A. Right.
Q. Is that the same thing as algebraic moment invariants?
A. Right.
Q. So it looks like the methodologies are the same for the features that are extracted between the papers; is that right?
A. I thought one paper had algebraic moments and one had turning angle. Maybe they both hand turning angle.
(Reporter seeks clarification.)
A. Turning angle. Algebraic moments. Turning angles.
Q. The 1993 paper refers to algebraic moment invariants. The IEEE paper, on IBM 900, refers to features derived from object moments.
Are those the same thing?

		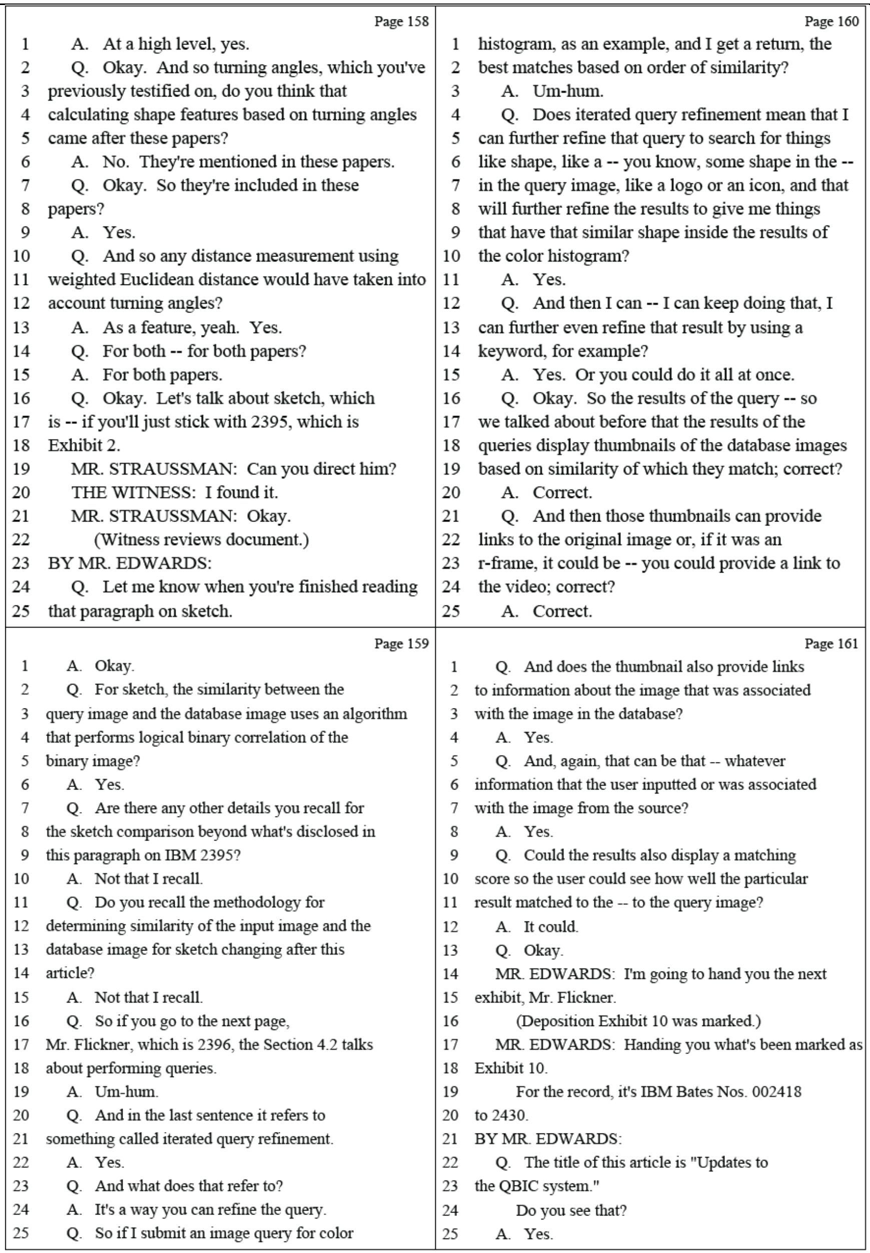

Page 158

A. At a high level, yes.

Q. Okay. And so turning angles, which you've previously testified on, do you think that calculating shape features based on turning angles came after these papers?

A. No. They're mentioned in these papers.

Q. Okay. So they're included in these papers?

A. Yes.

Q. And so any distance measurement using weighted Euclidean distance would have taken into account turning angles?

A. As a feature, yeah. Yes.

Q. For both -- for both papers?

A. For both papers.

Q. Okay. Let's talk about sketch, which is -- if you'll just stick with 2395, which is Exhibit 2.

MR. STRAUSSMAN: Can you direct him?

THE WITNESS: I found it.

MR. STRAUSSMAN: Okay.

(Witness reviews document.)

BY MR. EDWARDS:

Q. Let me know when you're finished reading that paragraph on sketch.

Page 159

A. Okay.

Q. For sketch, the similarity between the query image and the database image uses an algorithm that performs logical binary correlation of the binary image?

A. Yes.

Q. Are there any other details you recall for the sketch comparison beyond what's disclosed in this paragraph on IBM 2395?

A. Not that I recall.

Q. Do you recall the methodology for determining similarity of the input image and the database image for sketch changing after this article?

A. Not that I recall.

Q. So if you go to the next page, Mr. Flickner, which is 2396, the Section 4.2 talks about performing queries.

A. Um-hum.

Q. And in the last sentence it refers to something called iterated query refinement.

A. Yes.

Q. And what does that refer to?

A. It's a way you can refine the query.

Q. So if I submit an image query for color

Page 160

histogram, as an example, and I get a return, the best matches based on order of similarity?

A. Um-hum.

Q. Does iterated query refinement mean that I can further refine that query to search for things like shape, like a -- you know, some shape in the -- in the query image, like a logo or an icon, and that will further refine the results to give me things that have that similar shape inside the results of the color histogram?

A. Yes.

Q. And then I can -- I can keep doing that, I can further even refine that result by using a keyword, for example?

A. Yes. Or you could do it all at once.

Q. Okay. So the results of the query -- so we talked about before that the results of the queries display thumbnails of the database images based on similarity of which they match; correct?

A. Correct.

Q. And then those thumbnails can provide links to the original image or, if it was an r-frame, it could be -- you could provide a link to the video; correct?

A. Correct.

Page 161

Q. And does the thumbnail also provide links to information about the image that was associated with the image in the database?

A. Yes.

Q. And, again, that can be that -- whatever information that the user inputted or was associated with the image from the source?

A. Yes.

Q. Could the results also display a matching score so the user could see how well the particular result matched to the -- to the query image?

A. It could.

Q. Okay.

MR. EDWARDS: I'm going to hand you the next exhibit, Mr. Flickner.

(Deposition Exhibit 10 was marked.)

MR. EDWARDS: Handing you what's been marked as Exhibit 10.

For the record, it's IBM Bates Nos. 002418 to 2430.

BY MR. EDWARDS:

Q. The title of this article is "Updates to the QBIC system."

Do you see that?

A. Yes.

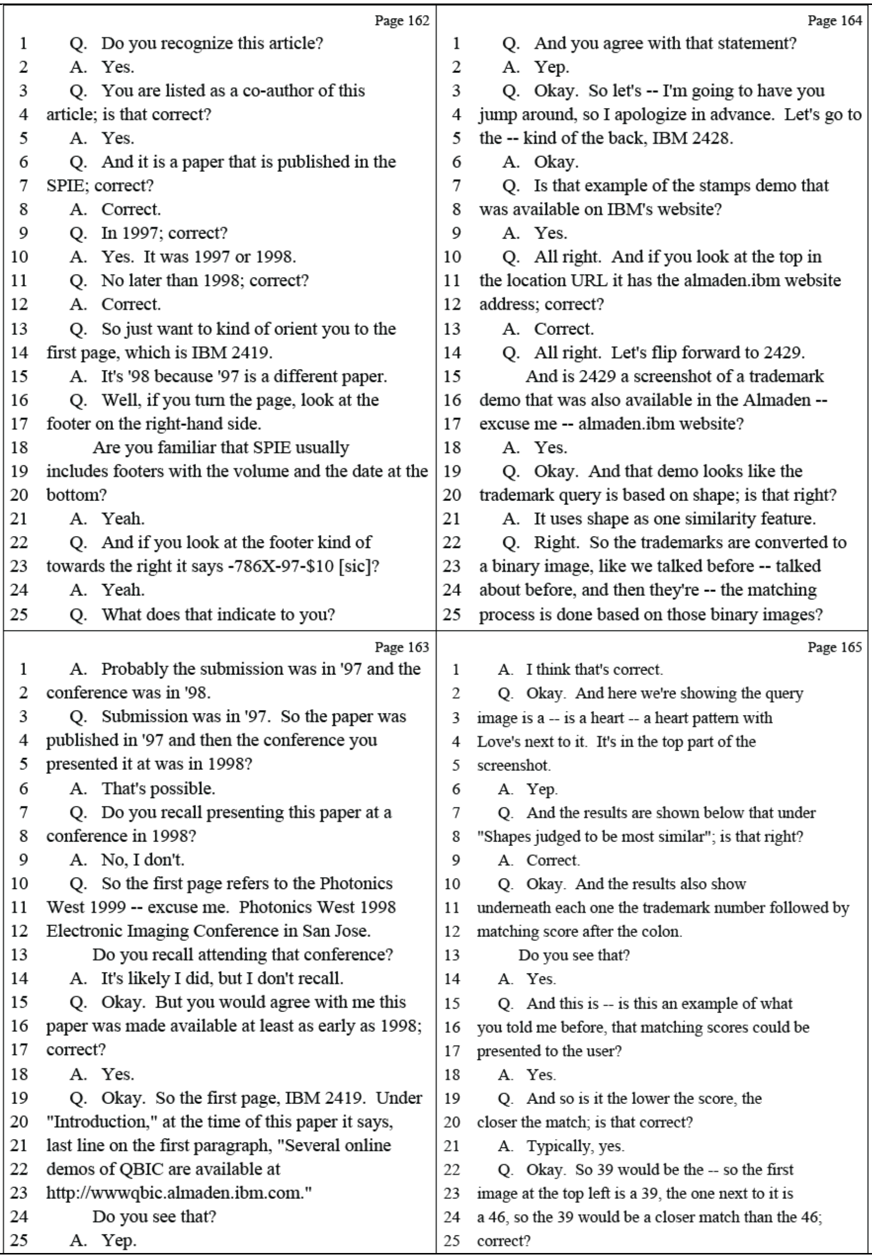

Page 162

Q. Do you recognize this article?
A. Yes.
Q. You are listed as a co-author of this article; is that correct?
A. Yes.
Q. And it is a paper that is published in the SPIE; correct?
A. Correct.
Q. In 1997; correct?
A. Yes. It was 1997 or 1998.
Q. No later than 1998; correct?
A. Correct.
Q. So just want to kind of orient you to the first page, which is IBM 2419.
A. It's '98 because '97 is a different paper.
Q. Well, if you turn the page, look at the footer on the right-hand side.
Are you familiar that SPIE usually includes footers with the volume and the date at the bottom?
A. Yeah.
Q. And if you look at the footer kind of towards the right it says -786X-97-$10 [sic]?
A. Yeah.
Q. What does that indicate to you?

Page 163

A. Probably the submission was in '97 and the conference was in '98.
Q. Submission was in '97. So the paper was published in '97 and then the conference you presented it at was in 1998?
A. That's possible.
Q. Do you recall presenting this paper at a conference in 1998?
A. No, I don't.
Q. So the first page refers to the Photonics West 1999 -- excuse me. Photonics West 1998 Electronic Imaging Conference in San Jose.
Do you recall attending that conference?
A. It's likely I did, but I don't recall.
Q. Okay. But you would agree with me this paper was made available at least as early as 1998; correct?
A. Yes.
Q. Okay. So the first page, IBM 2419. Under "Introduction," at the time of this paper it says, last line on the first paragraph, "Several online demos of QBIC are available at http://wwwqbic.almaden.ibm.com."
Do you see that?
A. Yep.

Page 164

Q. And you agree with that statement?
A. Yep.
Q. Okay. So let's -- I'm going to have you jump around, so I apologize in advance. Let's go to the -- kind of the back, IBM 2428.
A. Okay.
Q. Is that example of the stamps demo that was available on IBM's website?
A. Yes.
Q. All right. And if you look at the top in the location URL it has the almaden.ibm website address; correct?
A. Correct.
Q. All right. Let's flip forward to 2429.
And is 2429 a screenshot of a trademark demo that was also available in the Almaden -- excuse me -- almaden.ibm website?
A. Yes.
Q. Okay. And that demo looks like the trademark query is based on shape; is that right?
A. It uses shape as one similarity feature.
Q. Right. So the trademarks are converted to a binary image, like we talked before -- talked about before, and then they're -- the matching process is done based on those binary images?

Page 165

A. I think that's correct.
Q. Okay. And here we're showing the query image is a -- is a heart -- a heart pattern with Love's next to it. It's in the top part of the screenshot.
A. Yep.
Q. And the results are shown below that under "Shapes judged to be most similar"; is that right?
A. Correct.
Q. Okay. And the results also show underneath each one the trademark number followed by matching score after the colon.
Do you see that?
A. Yes.
Q. And this is -- is this an example of what you told me before, that matching scores could be presented to the user?
A. Yes.
Q. And so is it the lower the score, the closer the match; is that correct?
A. Typically, yes.
Q. Okay. So 39 would be the -- so the first image at the top left is a 39, the one next to it is a 46, so the 39 would be a closer match than the 46; correct?

Element	'004 Claim Recitation	Exemplary Citations in Reference
		Id. at 174–182 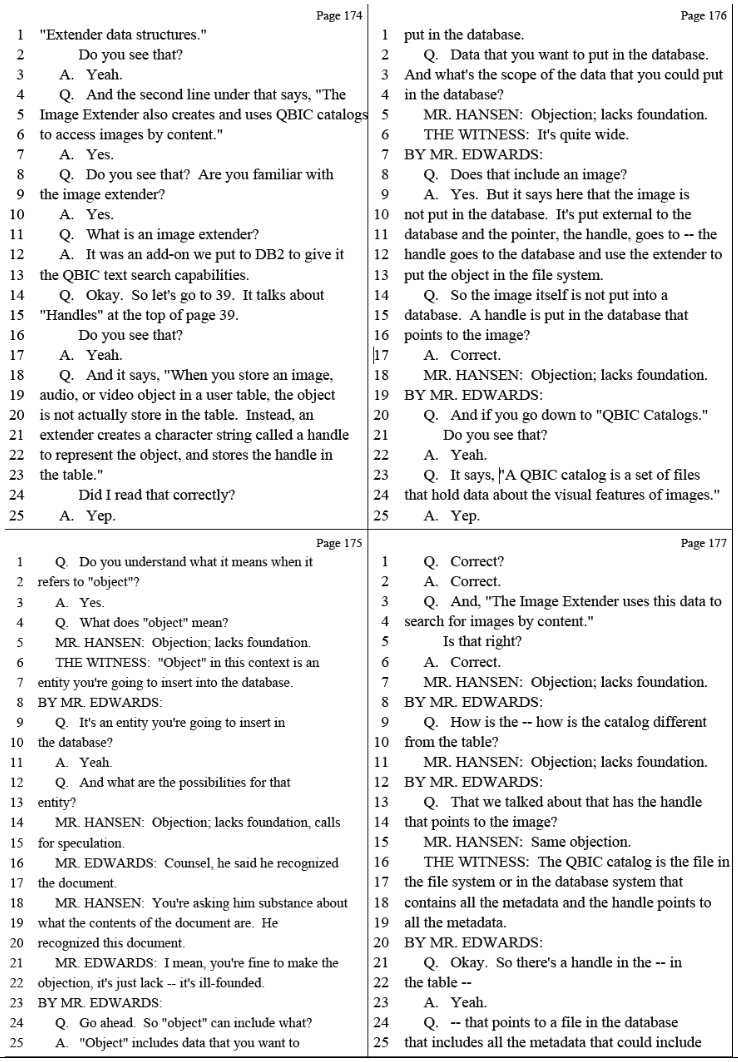

Page 174

"Extender data structures."
Do you see that?
A. Yeah.
Q. And the second line under that says, "The Image Extender also creates and uses QBIC catalogs to access images by content."
A. Yes.
Q. Do you see that? Are you familiar with the image extender?
A. Yes.
Q. What is an image extender?
A. It was an add-on we put to DB2 to give it the QBIC text search capabilities.
Q. Okay. So let's go to 39. It talks about "Handles" at the top of page 39.
Do you see that?
A. Yeah.
Q. And it says, "When you store an image, audio, or video object in a user table, the object is not actually store in the table. Instead, an extender creates a character string called a handle to represent the object, and stores the handle in the table."
Did I read that correctly?
A. Yep.

Page 175

Q. Do you understand what it means when it refers to "object"?
A. Yes.
Q. What does "object" mean?
MR. HANSEN: Objection; lacks foundation.
THE WITNESS: "Object" in this context is an entity you're going to insert into the database.
BY MR. EDWARDS:
Q. It's an entity you're going to insert in the database?
A. Yeah.
Q. And what are the possibilities for that entity?
MR. HANSEN: Objection; lacks foundation, calls for speculation.
MR. EDWARDS: Counsel, he said he recognized the document.
MR. HANSEN: You're asking him substance about what the contents of the document are. He recognized this document.
MR. EDWARDS: I mean, you're fine to make the objection, it's just lack -- it's ill-founded.
BY MR. EDWARDS:
Q. Go ahead. So "object" can include what?
A. "Object" includes data that you want to

Page 176

put in the database.
Q. Data that you want to put in the database. And what's the scope of the data that you could put in the database?
MR. HANSEN: Objection; lacks foundation.
THE WITNESS: It's quite wide.
BY MR. EDWARDS:
Q. Does that include an image?
A. Yes. But it says here that the image is not put in the database. It's put external to the database and the pointer, the handle, goes to -- the handle goes to the database and use the extender to put the object in the file system.
Q. So the image itself is not put into a database. A handle is put in the database that points to the image?
A. Correct.
MR. HANSEN: Objection; lacks foundation.
BY MR. EDWARDS:
Q. And if you go down to "QBIC Catalogs."
Do you see that?
A. Yeah.
Q. It says, "A QBIC catalog is a set of files that hold data about the visual features of images."
A. Yep.

Page 177

Q. Correct?
A. Correct.
Q. And, "The Image Extender uses this data to search for images by content."
Is that right?
A. Correct.
MR. HANSEN: Objection; lacks foundation.
BY MR. EDWARDS:
Q. How is the -- how is the catalog different from the table?
MR. HANSEN: Objection; lacks foundation.
BY MR. EDWARDS:
Q. That we talked about that has the handle that points to the image?
MR. HANSEN: Same objection.
THE WITNESS: The QBIC catalog is the file in the file system or in the database system that contains all the metadata and the handle points to all the metadata.
BY MR. EDWARDS:
Q. Okay. So there's a handle in the -- in the table --
A. Yeah.
Q. -- that points to a file in the database that includes all the metadata that could include

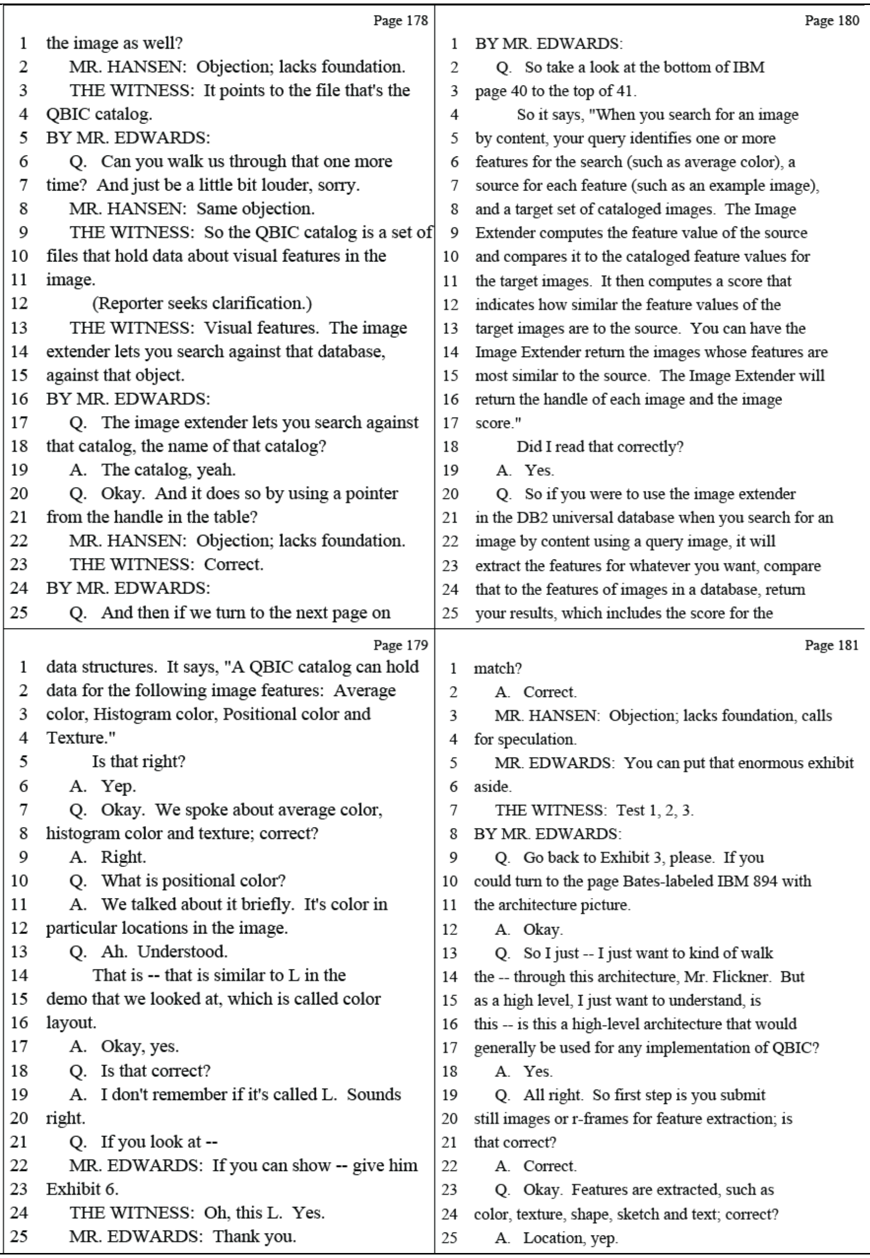

Page 178

the image as well?

MR. HANSEN: Objection; lacks foundation.

THE WITNESS: It points to the file that's the QBIC catalog.

BY MR. EDWARDS:

Q. Can you walk us through that one more time? And just be a little bit louder, sorry.

MR. HANSEN: Same objection.

THE WITNESS: So the QBIC catalog is a set of files that hold data about visual features in the image.

(Reporter seeks clarification.)

THE WITNESS: Visual features. The image extender lets you search against that database, against that object.

BY MR. EDWARDS:

Q. The image extender lets you search against that catalog, the name of that catalog?

A. The catalog, yeah.

Q. Okay. And it does so by using a pointer from the handle in the table?

MR. HANSEN: Objection; lacks foundation.

THE WITNESS: Correct.

BY MR. EDWARDS:

Q. And then if we turn to the next page on

Page 179

data structures. It says, "A QBIC catalog can hold data for the following image features: Average color, Histogram color, Positional color and Texture."

Is that right?

A. Yep.

Q. Okay. We spoke about average color, histogram color and texture; correct?

A. Right.

Q. What is positional color?

A. We talked about it briefly. It's color in particular locations in the image.

Q. Ah. Understood.

That is -- that is similar to L in the demo that we looked at, which is called color layout.

A. Okay, yes.

Q. Is that correct?

A. I don't remember if it's called L. Sounds right.

Q. If you look at --

MR. EDWARDS: If you can show -- give him Exhibit 6.

THE WITNESS: Oh, this L. Yes.

MR. EDWARDS: Thank you.

Page 180

BY MR. EDWARDS:

Q. So take a look at the bottom of IBM page 40 to the top of 41.

So it says, "When you search for an image by content, your query identifies one or more features for the search (such as average color), a source for each feature (such as an example image), and a target set of cataloged images. The Image Extender computes the feature value of the source and compares it to the cataloged feature values for the target images. It then computes a score that indicates how similar the feature values of the target images are to the source. You can have the Image Extender return the images whose features are most similar to the source. The Image Extender will return the handle of each image and the image score."

Did I read that correctly?

A. Yes.

Q. So if you were to use the image extender in the DB2 universal database when you search for an image by content using a query image, it will extract the features for whatever you want, compare that to the features of images in a database, return your results, which includes the score for the

Page 181

match?

A. Correct.

MR. HANSEN: Objection; lacks foundation, calls for speculation.

MR. EDWARDS: You can put that enormous exhibit aside.

THE WITNESS: Test 1, 2, 3.

BY MR. EDWARDS:

Q. Go back to Exhibit 3, please. If you could turn to the page Bates-labeled IBM 894 with the architecture picture.

A. Okay.

Q. So I just -- I just want to kind of walk the -- through this architecture, Mr. Flickner. But as a high level, I just want to understand, is this -- is this a high-level architecture that would generally be used for any implementation of QBIC?

A. Yes.

Q. All right. So first step is you submit still images or r-frames for feature extraction; is that correct?

A. Correct.

Q. Okay. Features are extracted, such as color, texture, shape, sketch and text; correct?

A. Location, yep.

Element	'004 Claim Recitation	Exemplary Citations in Reference
		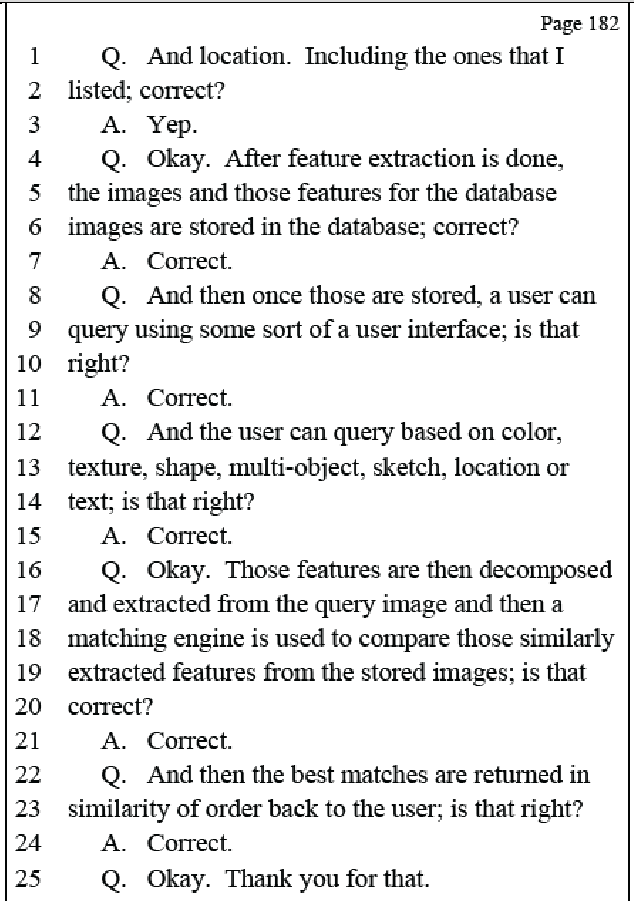 Page 182 Q. And location. Including the ones that I listed; correct? A. Yep. Q. Okay. After feature extraction is done, the images and those features for the database images are stored in the database; correct? A. Correct. Q. And then once those are stored, a user can query using some sort of a user interface; is that right? A. Correct. Q. And the user can query based on color, texture, shape, multi-object, sketch, location or text; is that right? A. Correct. Q. Okay. Those features are then decomposed and extracted from the query image and then a matching engine is used to compare those similarly extracted features from the stored images; is that correct? A. Correct. Q. And then the best matches are returned in similarity of order back to the user; is that right? A. Correct. Q. Okay. Thank you for that. *Id.* at 190–191 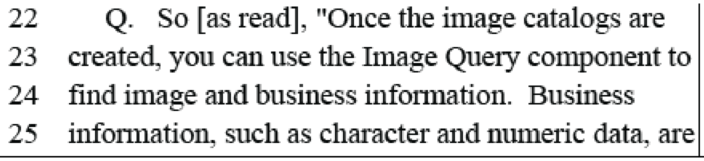 Q. So [as read], "Once the image catalogs are created, you can use the Image Query component to find image and business information. Business information, such as character and numeric data, are

		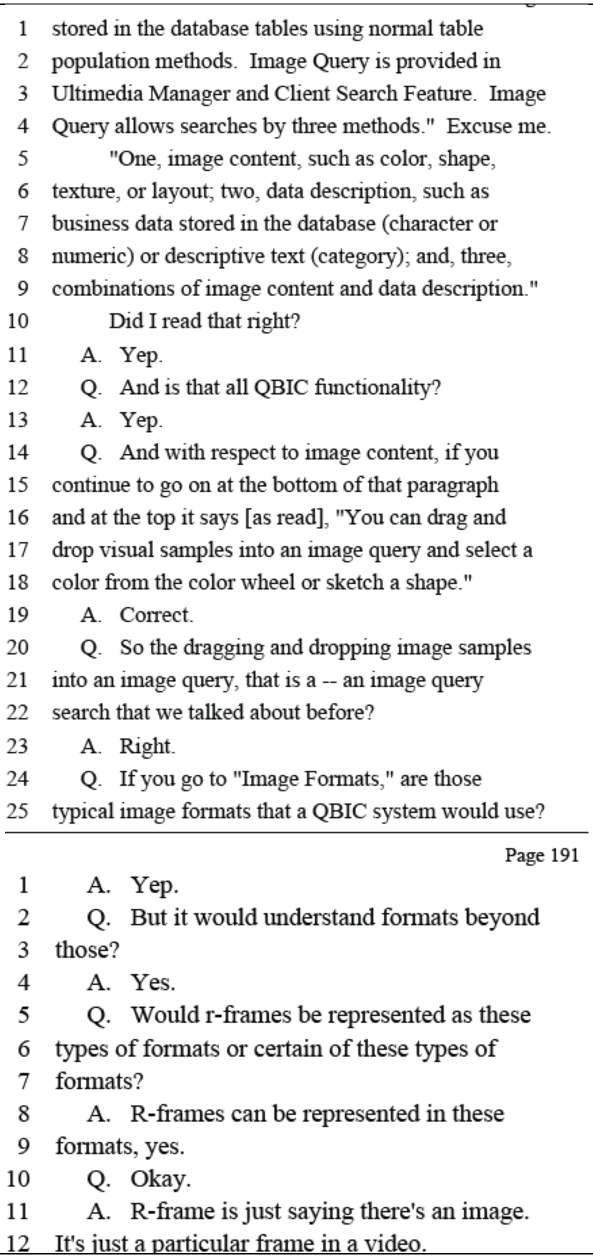 stored in the database tables using normal table population methods. Image Query is provided in Ultimedia Manager and Client Search Feature. Image Query allows searches by three methods." Excuse me. "One, image content, such as color, shape, texture, or layout; two, data description, such as business data stored in the database (character or numeric) or descriptive text (category); and, three, combinations of image content and data description." Did I read that right? A. Yep. Q. And is that all QBIC functionality? A. Yep. Q. And with respect to image content, if you continue to go on at the bottom of that paragraph and at the top it says [as read], "You can drag and drop visual samples into an image query and select a color from the color wheel or sketch a shape." A. Correct. Q. So the dragging and dropping image samples into an image query, that is a -- an image query search that we talked about before? A. Right. Q. If you go to "Image Formats," are those typical image formats that a QBIC system would use? Page 191 A. Yep. Q. But it would understand formats beyond those? A. Yes. Q. Would r-frames be represented as these types of formats or certain of these types of formats? A. R-frames can be represented in these formats, yes. Q. Okay. A. R-frame is just saying there's an image. It's just a particular frame in a video.

Element	'004 Claim Recitation	Exemplary Citations in Reference
		Id. at 198 – 199 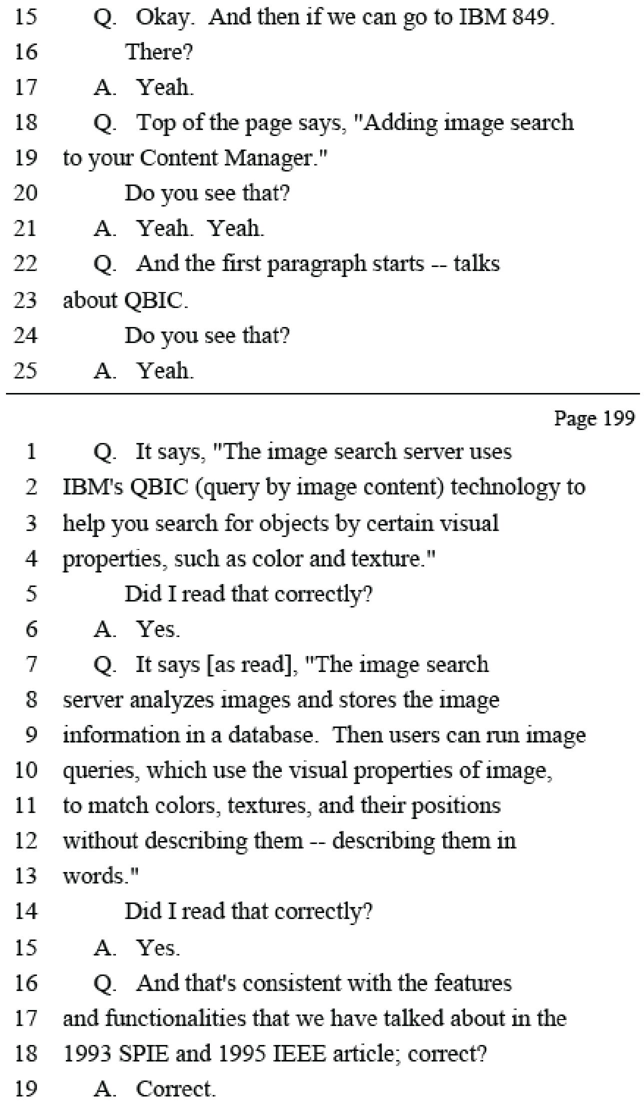 Q. Okay. And then if we can go to IBM 849. There? A. Yeah. Q. Top of the page says, "Adding image search to your Content Manager." Do you see that? A. Yeah. Yeah. Q. And the first paragraph starts -- talks about QBIC. Do you see that? A. Yeah. Page 199 Q. It says, "The image search server uses IBM's QBIC (query by image content) technology to help you search for objects by certain visual properties, such as color and texture." Did I read that correctly? A. Yes. Q. It says [as read], "The image search server analyzes images and stores the image information in a database. Then users can run image queries, which use the visual properties of image, to match colors, textures, and their positions without describing them -- describing them in words." Did I read that correctly? A. Yes. Q. And that's consistent with the features and functionalities that we have talked about in the 1993 SPIE and 1995 IEEE article; correct? A. Correct. *Id.* at 312 – 313

Element	'004 Claim Recitation	Exemplary Citations in Reference
		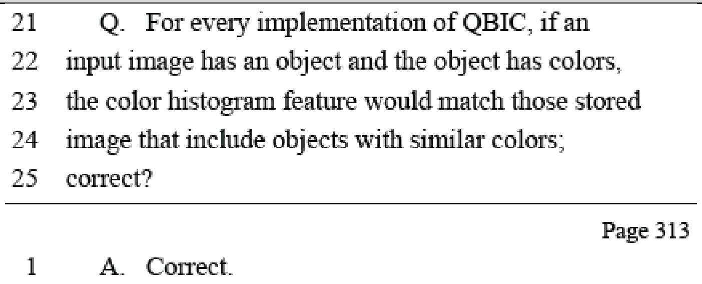 Q. For every implementation of QBIC, if an input image has an object and the object has colors, the color histogram feature would match those stored image that include objects with similar colors; correct? Page 313 A. Correct. *Id.* at 317–322 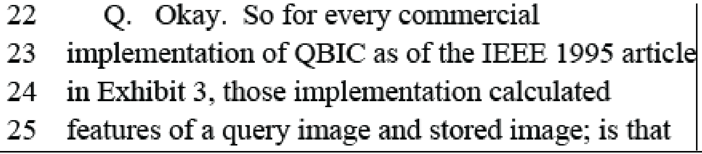 Q. Okay. So for every commercial implementation of QBIC as of the IEEE 1995 article in Exhibit 3, those implementation calculated features of a query image and stored image; is that

		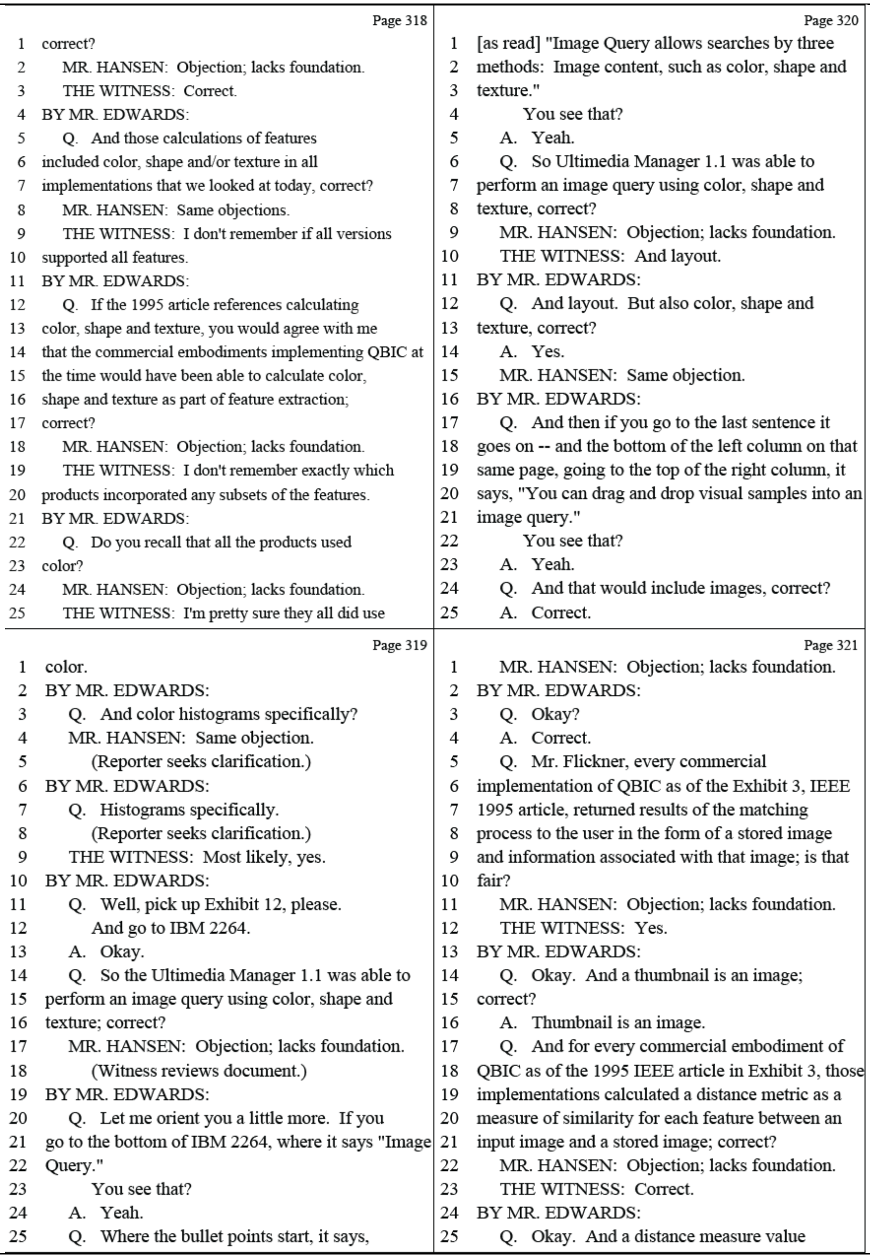

Page 318

correct?

MR. HANSEN: Objection; lacks foundation.

THE WITNESS: Correct.

BY MR. EDWARDS:

Q. And those calculations of features included color, shape and/or texture in all implementations that we looked at today, correct?

MR. HANSEN: Same objections.

THE WITNESS: I don't remember if all versions supported all features.

BY MR. EDWARDS:

Q. If the 1995 article references calculating color, shape and texture, you would agree with me that the commercial embodiments implementing QBIC at the time would have been able to calculate color, shape and texture as part of feature extraction; correct?

MR. HANSEN: Objection; lacks foundation.

THE WITNESS: I don't remember exactly which products incorporated any subsets of the features.

BY MR. EDWARDS:

Q. Do you recall that all the products used color?

MR. HANSEN: Objection; lacks foundation.

THE WITNESS: I'm pretty sure they all did use

Page 319

color.

BY MR. EDWARDS:

Q. And color histograms specifically?

MR. HANSEN: Same objection.

(Reporter seeks clarification.)

BY MR. EDWARDS:

Q. Histograms specifically.

(Reporter seeks clarification.)

THE WITNESS: Most likely, yes.

BY MR. EDWARDS:

Q. Well, pick up Exhibit 12, please.

And go to IBM 2264.

A. Okay.

Q. So the Ultimedia Manager 1.1 was able to perform an image query using color, shape and texture; correct?

MR. HANSEN: Objection; lacks foundation.

(Witness reviews document.)

BY MR. EDWARDS:

Q. Let me orient you a little more. If you go to the bottom of IBM 2264, where it says "Image Query."

You see that?

A. Yeah.

Q. Where the bullet points start, it says,

Page 320

[as read] "Image Query allows searches by three methods: Image content, such as color, shape and texture."

You see that?

A. Yeah.

Q. So Ultimedia Manager 1.1 was able to perform an image query using color, shape and texture, correct?

MR. HANSEN: Objection; lacks foundation.

THE WITNESS: And layout.

BY MR. EDWARDS:

Q. And layout. But also color, shape and texture, correct?

A. Yes.

MR. HANSEN: Same objection.

BY MR. EDWARDS:

Q. And then if you go to the last sentence it goes on -- and the bottom of the left column on that same page, going to the top of the right column, it says, "You can drag and drop visual samples into an image query."

You see that?

A. Yeah.

Q. And that would include images, correct?

A. Correct.

Page 321

MR. HANSEN: Objection; lacks foundation.

BY MR. EDWARDS:

Q. Okay?

A. Correct.

Q. Mr. Flickner, every commercial implementation of QBIC as of the Exhibit 3, IEEE 1995 article, returned results of the matching process to the user in the form of a stored image and information associated with that image; is that fair?

MR. HANSEN: Objection; lacks foundation.

THE WITNESS: Yes.

BY MR. EDWARDS:

Q. Okay. And a thumbnail is an image; correct?

A. Thumbnail is an image.

Q. And for every commercial embodiment of QBIC as of the 1995 IEEE article in Exhibit 3, those implementations calculated a distance metric as a measure of similarity for each feature between an input image and a stored image; correct?

MR. HANSEN: Objection; lacks foundation.

THE WITNESS: Correct.

BY MR. EDWARDS:

Q. Okay. And a distance measure value

Element	'004 Claim Recitation	Exemplary Citations in Reference
		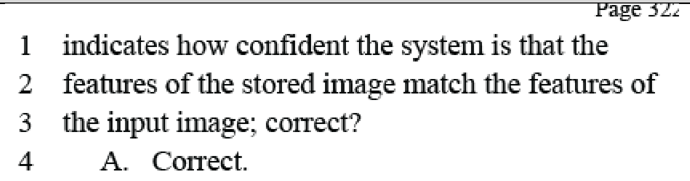 Page 322 indicates how confident the system is that the features of the stored image match the features of the input image; correct? A. Correct. IBM0002263 at 2264:

Description

Ultimedia Manager 1.1 opens new dimensions in visual queries, helping you manage your image database more efficiently. With its easy-to-use interface, even entry-level users can perform sophisticated searches. You can search by image content, data description, or a combination of both. Once the images are found, you can view the images in full resolution, use the images in another application, or read the associated textual comments.

You can use Ultimedia Manager 1.1 either as a stand-alone product or in a LAN environment. As a stand-alone product, Ultimedia Manager 1.1 employs full image classification and image query capabilities. Within the LAN environment, the Ultimedia Manager image catalogs can be sorted on a DB2 for OS/2 or DB2 for AIX database server. The image catalogs can be accessed by Windows or OS/2 clients using the Client Search Feature or the Visualizer Ultimedia Query.

Ultimedia Manager runs as an OS/2 application to a DB2 for OS/2 or DB2 for AIX database server in a LAN, or as a stand-alone application with a local DB2 for OS/2 database. The Client Search Feature runs on DOS or Windows, and accesses files in dBASE flat file format. A database manager is not required to access catalogs in dBASE file format.

Image Classification

With the Image Classifier, you can create visual indexes based on color, shape, category, and layout elements, which are areas of the image containing the patterns you want to use for searches. You can store these elements in catalogs in DB2 for OS/2 databases. The catalogs can be used locally or, through the use of utility programs, be transferred to DB2 for OS/2 or AIX servers. In addition, you can convert the catalog data into dBASE flat file format for access by Windows client applications.

The Image Classifier and Analyzer allows automatic or interactive classification of images for later image queries, flexibility in storage of catalog information, and access to multiple users. You can organize large image catalogs so that later searches for images can be performed easily. You can classify automatically using a batch program, or interactively by specifying areas of interest in an image and indicating the visual cue (color, shape, or texture) to be saved for later queries.

Image Query

Once the image catalogs are created, you can use the Image Query component to find image and business information. Business information, such as character and numeric data, are stored in the database tables using normal table population methods. Image Query is provided in Ultimedia Manager and Client Search Feature.

Image Query allows searches by three methods:

- Image content, such as color, shape, texture, or layout
- Data description, such as business data stored in the database (character or numeric) or descriptive text (category)
- Combinations of image content and data description.

Image Content: To search by image content, specify the attributes (color, shape, texture, or layout). You can drag and drop visual samples into an image query, select a color from the color wheel, or sketch a shape.

Data Description: To search by data description, enter character and numeric data values in the database tables using the simple fill-in-the-blanks query panel. You can enter data such as name, number, date, or a range of dates. You can integrate character or numeric data as part of the image query, enabling a more effective search for better results. The simple fill-in-the-blanks method for specifying desired conditions is so easy even non-programmers can use database table information as part of their searches.

Combinations: Image queries using either image content or data description, or combinations of both, allow you to search through thousands of images more quickly and effectively than the traditional text and keyword search methods.

Image Formats

Ultimedia Manager 1.1 understands these image formats:

- Audio Visual Connection (AVC) images
- OS/2 bitmap
- Windows 3.x bitmap
- PCX
- Intel and Motorola TIFF
- TARGA TGA
- TIFF (FAX)
- GIF
- Kodak Photo-CD

Product Positioning

The advanced image search technology in Ultimedia Manager, query by image content, was developed in conjunction with the IBM Almaden Research Center. It provides function not available with conventional, text-only image management systems.

Ultimedia Manager fulfills image search and retrieval needs of business users and creative professionals through use of online image libraries or collections. It is intended for users who have large image libraries and need fast, flexible access to their image data, users who may also want to add character or numeric data to their image queries, and users who feel limited by text-based systems.

Ultimedia Manager 1.1 works with the DB2 Family of products and the Visualizer Ultimedia Query product. Catalogs created by Utlimedia Manager can be stored in DB2 for OS/2 database tables. Image catalogs also can be stored in DB2 for OS/2 or DB2 for AIX database servers in client/server environments.

OS/2 and Windows clients in a LAN environment using the Client Search Feature and the Visualizer Ultimedia Query can access catalogs created with Ultimedia Manager 1.1. Visualizer Ultimedia Query provides users with the capability to:

- Add multimedia columns (such as image, video, audio, and formatted text) to DB2 for OS/2 database tables
- Query and display images
- View or listen to video and audio information
- Display formatted text

Users of Ultimedia Manager 1.1 and Client Search Feature are focused on image management and query applications. Visualizer Ultimedia Query users are

Element	'004 Claim Recitation	Exemplary Citations in Reference
		IBM000001 at 37-39 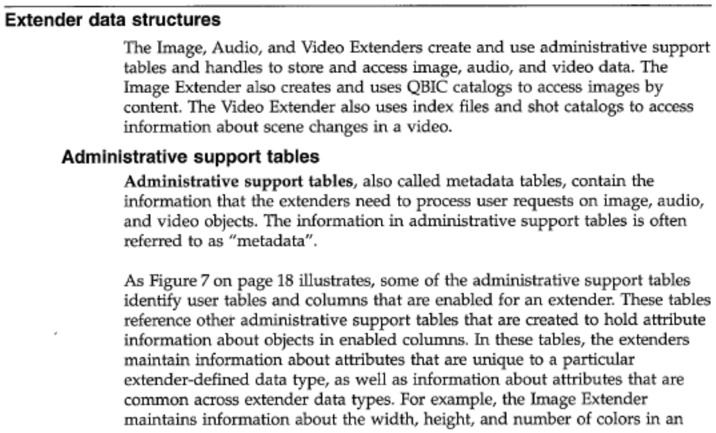 **Extender data structures** The Image, Audio, and Video Extenders create and use administrative support tables and handles to store and access image, audio, and video data. The Image Extender also creates and uses QBIC catalogs to access images by content. The Video Extender also uses index files and shot catalogs to access information about scene changes in a video. **Administrative support tables** **Administrative support tables**, also called metadata tables, contain the information that the extenders need to process user requests on image, audio, and video objects. The information in administrative support tables is often referred to as "metadata". As Figure 7 on page 18 illustrates, some of the administrative support tables identify user tables and columns that are enabled for an extender. These tables reference other administrative support tables that are created to hold attribute information about objects in enabled columns. In these tables, the extenders maintain information about attributes that are unique to a particular extender-defined data type, as well as information about attributes that are common across extender data types. For example, the Image Extender maintains information about the width, height, and number of colors in an

Element	'004 Claim Recitation	Exemplary Citations in Reference
		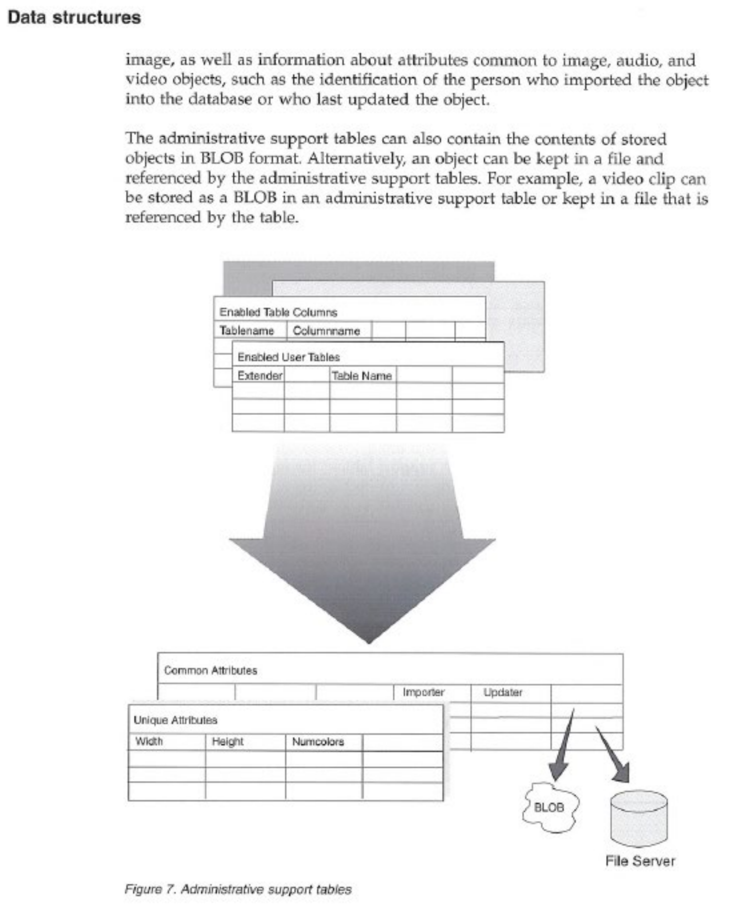 **Data structures** image, as well as information about attributes common to image, audio, and video objects, such as the identification of the person who imported the object into the database or who last updated the object. The administrative support tables can also contain the contents of stored objects in BLOB format. Alternatively, an object can be kept in a file and referenced by the administrative support tables. For example, a video clip can be stored as a BLOB in an administrative support table or kept in a file that is referenced by the table. *Figure 7. Administrative support tables* IBM000001 at 40-41

Element	'004 Claim Recitation	Exemplary Citations in Reference
		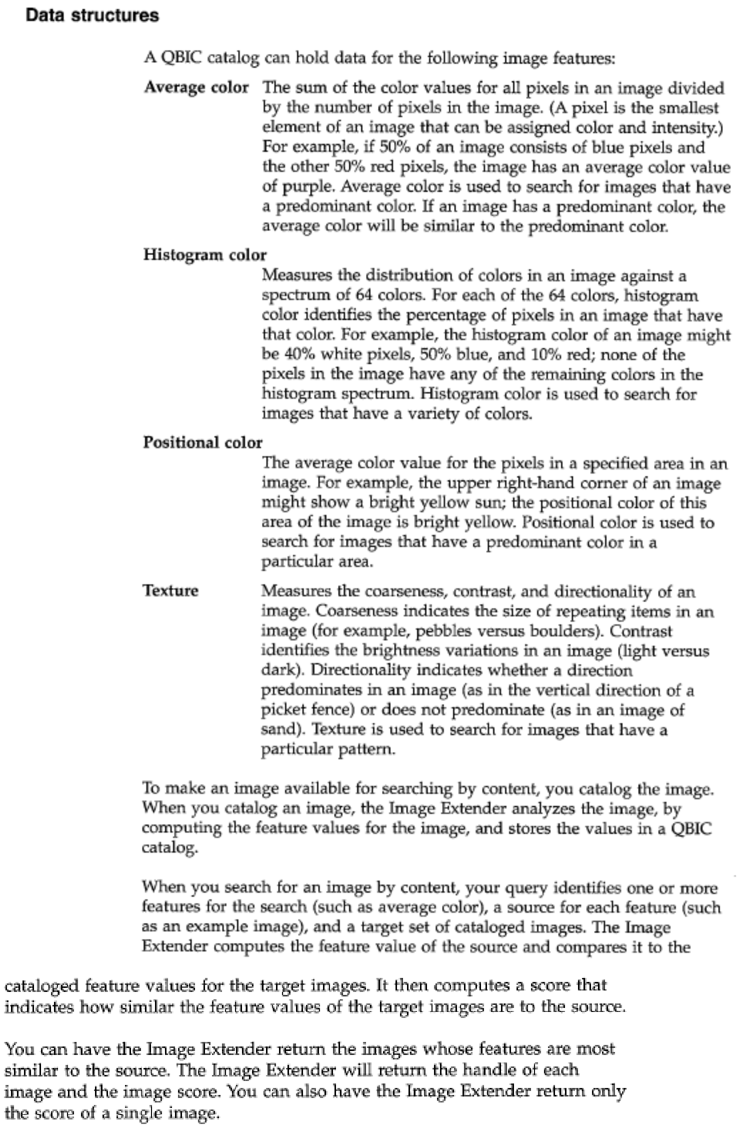 **Data structures** A QBIC catalog can hold data for the following image features: **Average color** The sum of the color values for all pixels in an image divided by the number of pixels in the image. (A pixel is the smallest element of an image that can be assigned color and intensity.) For example, if 50% of an image consists of blue pixels and the other 50% red pixels, the image has an average color value of purple. Average color is used to search for images that have a predominant color. If an image has a predominant color, the average color will be similar to the predominant color. **Histogram color** Measures the distribution of colors in an image against a spectrum of 64 colors. For each of the 64 colors, histogram color identifies the percentage of pixels in an image that have that color. For example, the histogram color of an image might be 40% white pixels, 50% blue, and 10% red; none of the pixels in the image have any of the remaining colors in the histogram spectrum. Histogram color is used to search for images that have a variety of colors. **Positional color** The average color value for the pixels in a specified area in an image. For example, the upper right-hand corner of an image might show a bright yellow sun; the positional color of this area of the image is bright yellow. Positional color is used to search for images that have a predominant color in a particular area. **Texture** Measures the coarseness, contrast, and directionality of an image. Coarseness indicates the size of repeating items in an image (for example, pebbles versus boulders). Contrast identifies the brightness variations in an image (light versus dark). Directionality indicates whether a direction predominates in an image (as in the vertical direction of a picket fence) or does not predominate (as in an image of sand). Texture is used to search for images that have a particular pattern. To make an image available for searching by content, you catalog the image. When you catalog an image, the Image Extender analyzes the image, by computing the feature values for the image, and stores the values in a QBIC catalog. When you search for an image by content, your query identifies one or more features for the search (such as average color), a source for each feature (such as an example image), and a target set of cataloged images. The Image Extender computes the feature value of the source and compares it to the cataloged feature values for the target images. It then computes a score that indicates how similar the feature values of the target images are to the source. You can have the Image Extender return the images whose features are most similar to the source. The Image Extender will return the handle of each image and the image score. You can also have the Image Extender return only the score of a single image. Flickner 8/29/23 Depo Ex. 7 at 1

Element	'004 Claim Recitation	Exemplary Citations in Reference
		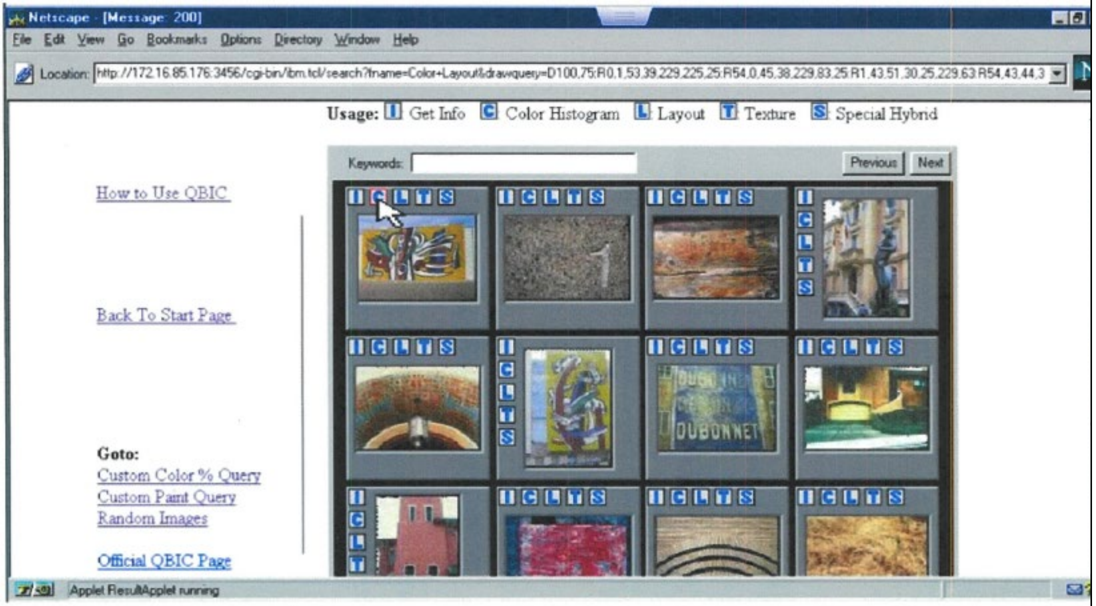 Location: http://172.16.85.176:3456/cgi-bin/ibm.tcl/search?fname=Color+Layout&drawquery=D100,75:R0,1,53,39,229,225,25:R54,0,45,38,229,83,25:R1,43,51,30,25,229,63:R54,43,44,3 *Id.* at 2

Element	'004 Claim Recitation	Exemplary Citations in Reference
		 *See also* Flickner Depo at 25-26, 38-41, 40-42, 87, 201, 209-210, 214-215. IBM000001 at 39 – 41

<table>
<tr><th>Element</th><th>'004 Claim Recitation</th><th>Exemplary Citations in Reference</th></tr>
<tr><td></td><td></td><td>
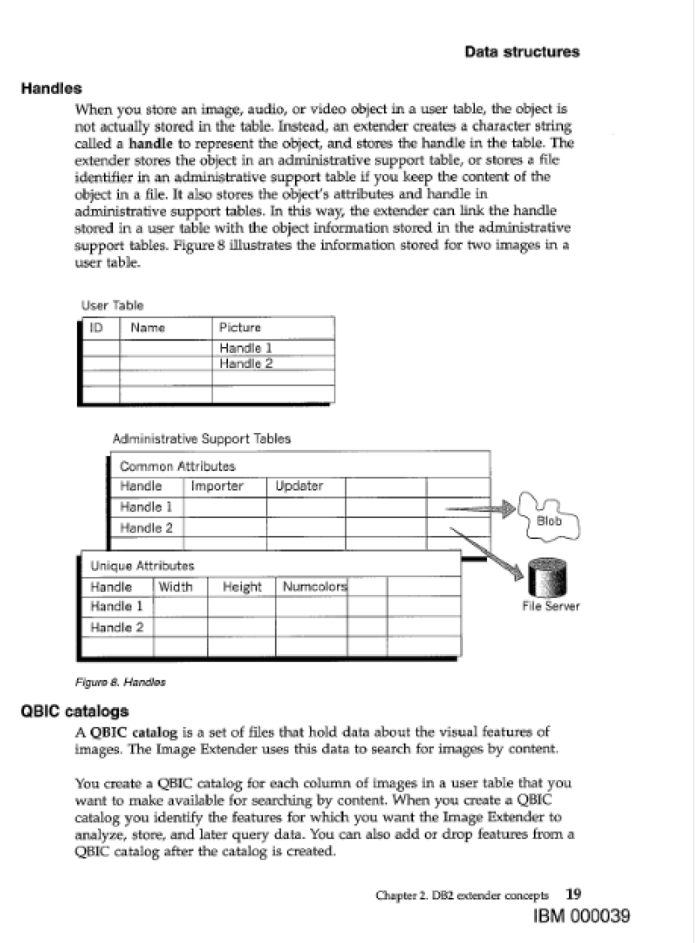

Data structures

Handles

When you store an image, audio, or video object in a user table, the object is not actually stored in the table. Instead, an extender creates a character string called a **handle** to represent the object, and stores the handle in the table. The extender stores the object in an administrative support table, or stores a file identifier in an administrative support table if you keep the content of the object in a file. It also stores the object's attributes and handle in administrative support tables. In this way, the extender can link the handle stored in a user table with the object information stored in the administrative support tables. Figure 8 illustrates the information stored for two images in a user table.

Figure 8. Handles

QBIC catalogs

A **QBIC catalog** is a set of files that hold data about the visual features of images. The Image Extender uses this data to search for images by content.

You create a QBIC catalog for each column of images in a user table that you want to make available for searching by content. When you create a QBIC catalog you identify the features for which you want the Image Extender to analyze, store, and later query data. You can also add or drop features from a QBIC catalog after the catalog is created.

Chapter 2. DB2 extender concepts 19

IBM 000039
</td></tr>
</table>

Element	'004 Claim Recitation	Exemplary Citations in Reference
		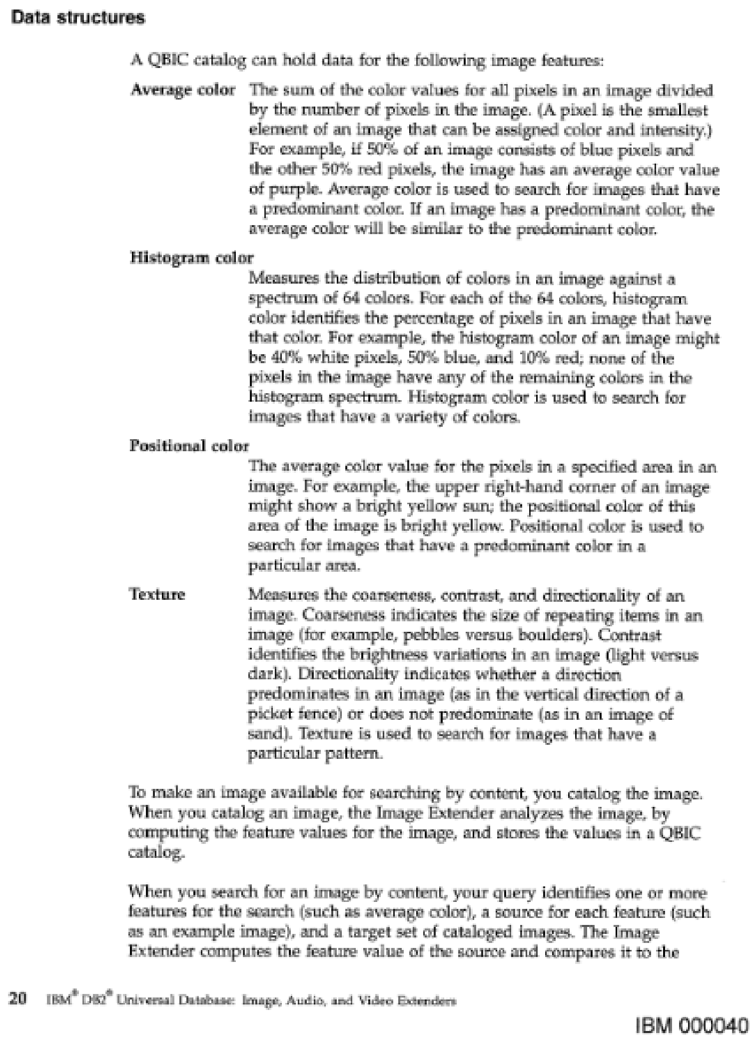 **Data structures** A QBIC catalog can hold data for the following image features: **Average color** The sum of the color values for all pixels in an image divided by the number of pixels in the image. (A pixel is the smallest element of an image that can be assigned color and intensity.) For example, if 50% of an image consists of blue pixels and the other 50% red pixels, the image has an average color value of purple. Average color is used to search for images that have a predominant color. If an image has a predominant color, the average color will be similar to the predominant color. **Histogram color** Measures the distribution of colors in an image against a spectrum of 64 colors. For each of the 64 colors, histogram color identifies the percentage of pixels in an image that have that color. For example, the histogram color of an image might be 40% white pixels, 50% blue, and 10% red; none of the pixels in the image have any of the remaining colors in the histogram spectrum. Histogram color is used to search for images that have a variety of colors. **Positional color** The average color value for the pixels in a specified area in an image. For example, the upper right-hand corner of an image might show a bright yellow sun; the positional color of this area of the image is bright yellow. Positional color is used to search for images that have a predominant color in a particular area. **Texture** Measures the coarseness, contrast, and directionality of an image. Coarseness indicates the size of repeating items in an image (for example, pebbles versus boulders). Contrast identifies the brightness variations in an image (light versus dark). Directionality indicates whether a direction predominates in an image (as in the vertical direction of a picket fence) or does not predominate (as in an image of sand). Texture is used to search for images that have a particular pattern. To make an image available for searching by content, you catalog the image. When you catalog an image, the Image Extender analyzes the image, by computing the feature values for the image, and stores the values in a QBIC catalog. When you search for an image by content, your query identifies one or more features for the search (such as average color), a source for each feature (such as an example image), and a target set of cataloged images. The Image Extender computes the feature value of the source and compares it to the 20 IBM® DB2® Universal Database: Image, Audio, and Video Extenders IBM 000040

Element	'004 Claim Recitation	Exemplary Citations in Reference
		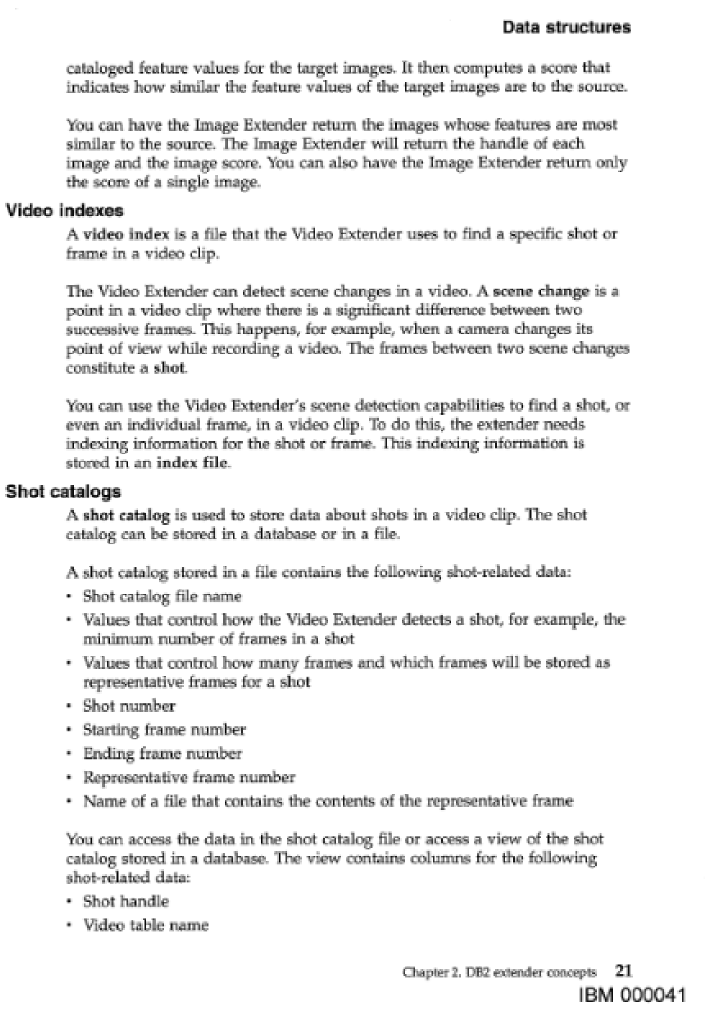 Data structures cataloged feature values for the target images. It then computes a score that indicates how similar the feature values of the target images are to the source. You can have the Image Extender return the images whose features are most similar to the source. The Image Extender will return the handle of each image and the image score. You can also have the Image Extender return only the score of a single image. **Video indexes** A **video index** is a file that the Video Extender uses to find a specific shot or frame in a video clip. The Video Extender can detect scene changes in a video. A **scene change** is a point in a video clip where there is a significant difference between two successive frames. This happens, for example, when a camera changes its point of view while recording a video. The frames between two scene changes constitute a **shot**. You can use the Video Extender's scene detection capabilities to find a shot, or even an individual frame, in a video clip. To do this, the extender needs indexing information for the shot or frame. This indexing information is stored in an **index file**. **Shot catalogs** A **shot catalog** is used to store data about shots in a video clip. The shot catalog can be stored in a database or in a file. A shot catalog stored in a file contains the following shot-related data: - Shot catalog file name - Values that control how the Video Extender detects a shot, for example, the minimum number of frames in a shot - Values that control how many frames and which frames will be stored as representative frames for a shot - Shot number - Starting frame number - Ending frame number - Representative frame number - Name of a file that contains the contents of the representative frame You can access the data in the shot catalog file or access a view of the shot catalog stored in a database. The view contains columns for the following shot-related data: - Shot handle - Video table name Chapter 2. DB2 extender concepts 21 IBM 000041 IBM0002263 at 226

Element	'004 Claim Recitation	Exemplary Citations in Reference
		 IBM000777 at 794

Description

Ultimedia Manager 1.1 opens new dimensions in visual queries, helping you manage your image database more efficiently. With its easy-to-use interface, even entry-level users can perform sophisticated searches. You can search by image content, data description, or a combination of both. Once the images are found, you can view the images in full resolution, use the images in another application, or read the associated textual comments.

You can use Ultimedia Manager 1.1 either as a stand-alone product or in a LAN environment. As a stand-alone product, Ultimedia Manager 1.1 employs full image classification and image query capabilities. Within the LAN environment, the Ultimedia Manager image catalogs can be sorted on a DB2 for OS/2 or DB2 for AIX database server. The image catalogs can be accessed by Windows or OS/2 clients using the Client Search Feature or the Visualizer Ultimedia Query.

Ultimedia Manager runs as an OS/2 application to a DB2 for OS/2 or DB2 for AIX database server in a LAN, or as a stand-alone application with a local DB2 for OS/2 database. The Client Search Feature runs on DOS or Windows, and accesses files in dBASE flat file format. A database manager is not required to access catalogs in dBASE file format.

Image Classification

With the Image Classifier, you can create visual indexes based on color, shape, category, and layout elements, which are areas of the image containing the patterns you want to use for searches. You can store these elements in catalogs in DB2 for OS/2 databases. The catalogs can be used locally or, through the use of utility programs, be transferred to DB2 for OS/2 or AIX servers. In addition, you can convert the catalog data into dBASE flat file format for access by Windows client applications.

The Image Classifier and Analyzer allows automatic or interactive classification of images for later image queries, flexibility in storage of catalog information, and access to multiple users. You can organize large image catalogs so that later searches for images can be performed easily. You can classify automatically using a batch program, or interactively by specifying areas of interest in an image and indicating the visual cue (color, shape, or texture) to be saved for later queries.

Image Query

Once the image catalogs are created, you can use the Image Query component to find image and business information. Business information, such as character and numeric data, are stored in the database tables using normal table population methods. Image Query is provided in Ultimedia Manager and Client Search Feature.

Image Query allows searches by three methods:

- Image content, such as color, shape, texture, or layout
- Data description, such as business data stored in the database (character or numeric) or descriptive text (category)
- Combinations of image content and data description.

Image Content: To search by image content, specify the attributes (color, shape, texture, or layout). You can drag and drop visual samples into an image query, select a color from the color wheel, or sketch a shape.

Data Description: To search by data description, enter character and numeric data values in the database tables using the simple fill-in-the-blanks query panel. You can enter data such as name, number, date, or a range of dates. You can integrate character or numeric data as part of the image query, enabling a more effective search for better results. The simple fill-in-the-blanks method for specifying desired conditions is so easy even non-programmers can use database table information as part of their searches.

Combinations: Image queries using either image content or data description, or combinations of both, allow you to search through thousands of images more quickly and effectively than the traditional text and keyword search methods.

Image Formats

Ultimedia Manager 1.1 understands these image formats:

- Audio Visual Connection (AVC) images
- OS/2 bitmap
- Windows 3.x bitmap
- PCX
- Intel and Motorola TIFF
- TARGA TGA
- TIFF (FAX)
- GIF
- Kodak Photo-CD

Product Positioning

The advanced image search technology in Ultimedia Manager, query by image content, was developed in conjunction with the IBM Almaden Research Center. It provides function not available with conventional, text-only image management systems.

Ultimedia Manager fulfills image search and retrieval needs of business users and creative professionals through use of online image libraries or collections. It is intended for users who have large image libraries and need fast, flexible access to their image data, users who may also want to add character or numeric data to their image queries, and users who feel limited by text-based systems.

Ultimedia Manager 1.1 works with the DB2 Family of products and the Visualizer Ultimedia Query product. Catalogs created by Ultimedia Manager can be stored in DB2 for OS/2 database tables. Image catalogs also can be stored in DB2 for OS/2 or DB2 for AIX database servers in client/server environments.

OS/2 and Windows clients in a LAN environment using the Client Search Feature and the Visualizer Ultimedia Query can access catalogs created with Ultimedia Manager 1.1. Visualizer Ultimedia Query provides users with the capability to:

- Add multimedia columns (such as image, video, audio, and formatted text) to DB2 for OS/2 database tables
- Query and display images
- View or listen to video and audio information
- Display formatted text

Users of Ultimedia Manager 1.1 and Client Search Feature are focused on image management and query applications. Visualizer Ultimedia Query users are

298-042 -2-

IBM 0002264

Element	'004 Claim Recitation	Exemplary Citations in Reference
		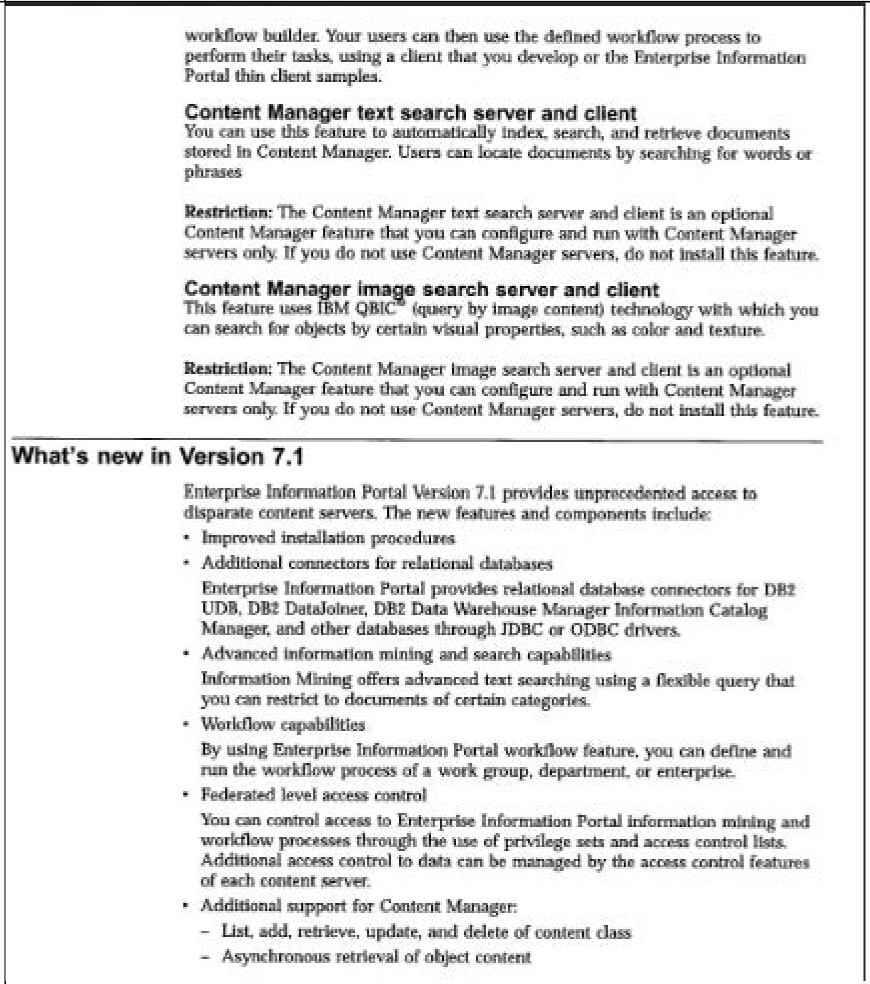 workflow builder. Your users can then use the defined workflow process to perform their tasks, using a client that you develop or the Enterprise Information Portal thin client samples. **Content Manager text search server and client** You can use this feature to automatically index, search, and retrieve documents stored in Content Manager. Users can locate documents by searching for words or phrases **Restriction:** The Content Manager text search server and client is an optional Content Manager feature that you can configure and run with Content Manager servers only. If you do not use Content Manager servers, do not install this feature. **Content Manager image search server and client** This feature uses IBM QBIC® (query by image content) technology with which you can search for objects by certain visual properties, such as color and texture. **Restriction:** The Content Manager image search server and client is an optional Content Manager feature that you can configure and run with Content Manager servers only. If you do not use Content Manager servers, do not install this feature. **What's new in Version 7.1** Enterprise Information Portal Version 7.1 provides unprecedented access to disparate content servers. The new features and components include: • Improved installation procedures • Additional connectors for relational databases Enterprise Information Portal provides relational database connectors for DB2 UDB, DB2 DataJoiner, DB2 Data Warehouse Manager Information Catalog Manager, and other databases through JDBC or ODBC drivers. • Advanced information mining and search capabilities Information Mining offers advanced text searching using a flexible query that you can restrict to documents of certain categories. • Workflow capabilities By using Enterprise Information Portal workflow feature, you can define and run the workflow process of a work group, department, or enterprise. • Federated level access control You can control access to Enterprise Information Portal information mining and workflow processes through the use of privilege sets and access control lists. Additional access control to data can be managed by the access control features of each content server. • Additional support for Content Manager: – List, add, retrieve, update, and delete of content class – Asynchronous retrieval of object content To the extent it is found that QBIC does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that QBIC does not anticipate

Element	'004 Claim Recitation	Exemplary Citations in Reference
		this claim, QBIC renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine QBIC with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1D]	providing an indication upon recognizing the at least one object as the target object; and	QBIC discloses the limitation. The QBIC Project: Querying Images by Content Using Color, Texture, and Shape (IBM 0002390–404) at 2396-98 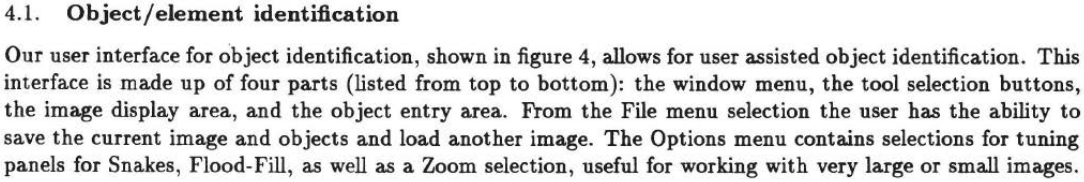 4.1. **Object/element identification** Our user interface for object identification, shown in figure 4, allows for user assisted object identification. This interface is made up of four parts (listed from top to bottom): the window menu, the tool selection buttons, the image display area, and the object entry area. From the File menu selection the user has the ability to save the current image and objects and load another image. The Options menu contains selections for tuning panels for Snakes, Flood-Fill, as well as a Zoom selection, useful for working with very large or small images. The QBIC Project: Querying Images by Content Using Color, Texture, and Shape (IBM 0002390–404) at Fig. 4

Element	'004 Claim Recitation	Exemplary Citations in Reference
		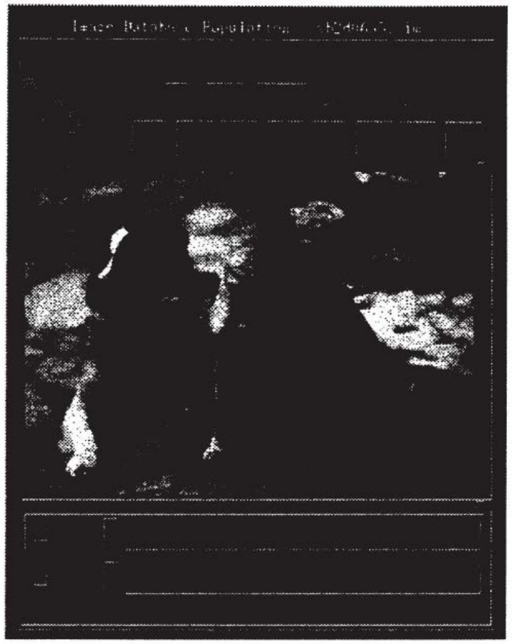 Figure 4: Object Identification Tool. Note for example, the right dog's ear, which has been outlined. The QBIC Project: Querying Images by Content Using Color, Texture, and Shape (IBM 0002390–404) at 2396

Element	'004 Claim Recitation	Exemplary Citations in Reference
		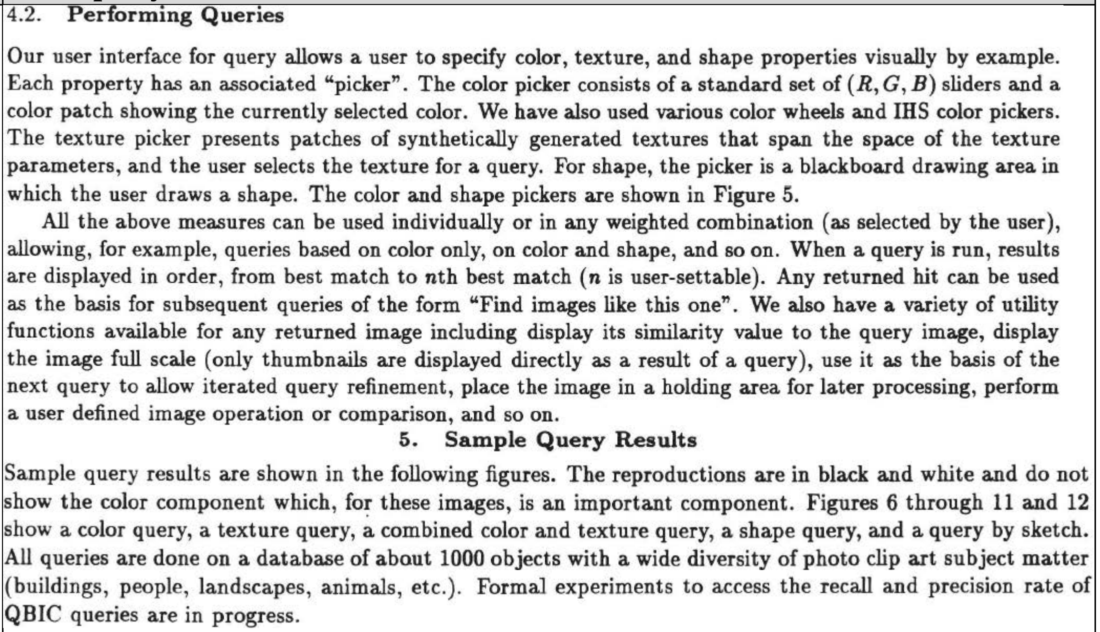 4.2. **Performing Queries** Our user interface for query allows a user to specify color, texture, and shape properties visually by example. Each property has an associated "picker". The color picker consists of a standard set of (R, G, B) sliders and a color patch showing the currently selected color. We have also used various color wheels and IHS color pickers. The texture picker presents patches of synthetically generated textures that span the space of the texture parameters, and the user selects the texture for a query. For shape, the picker is a blackboard drawing area in which the user draws a shape. The color and shape pickers are shown in Figure 5. All the above measures can be used individually or in any weighted combination (as selected by the user), allowing, for example, queries based on color only, on color and shape, and so on. When a query is run, results are displayed in order, from best match to nth best match (n is user-settable). Any returned hit can be used as the basis for subsequent queries of the form "Find images like this one". We also have a variety of utility functions available for any returned image including display its similarity value to the query image, display the image full scale (only thumbnails are displayed directly as a result of a query), use it as the basis of the next query to allow iterated query refinement, place the image in a holding area for later processing, perform a user defined image operation or comparison, and so on. 5. **Sample Query Results** Sample query results are shown in the following figures. The reproductions are in black and white and do not show the color component which, for these images, is an important component. Figures 6 through 11 and 12 show a color query, a texture query, a combined color and texture query, a shape query, and a query by sketch. All queries are done on a database of about 1000 objects with a wide diversity of photo clip art subject matter (buildings, people, landscapes, animals, etc.). Formal experiments to access the recall and precision rate of QBIC queries are in progress.

		The QBIC Project: Querying Images by Content Using Color, Texture, and Shape (IBM 0002390–404) at Figs. 7–9 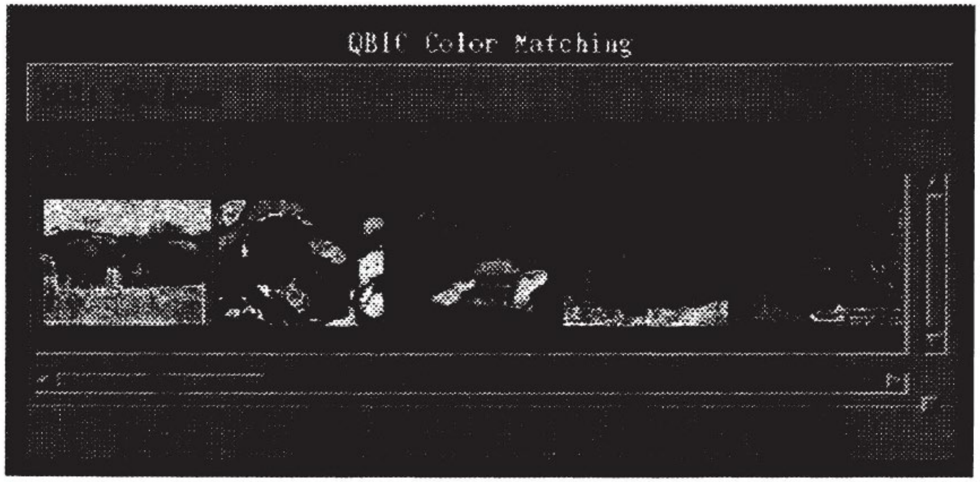 Figure 7: Result of query by example using color only.

Element	'004 Claim Recitation	Exemplary Citations in Reference
		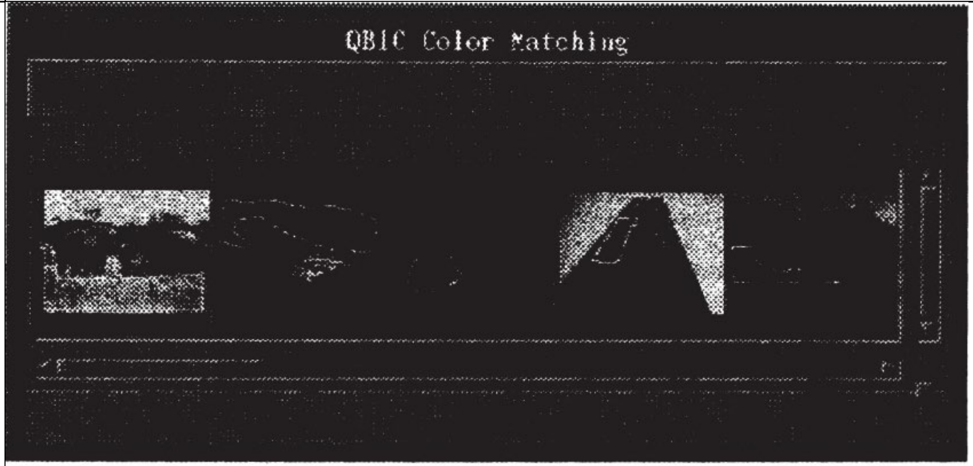 Figure 8: Result of query by example using texture only. 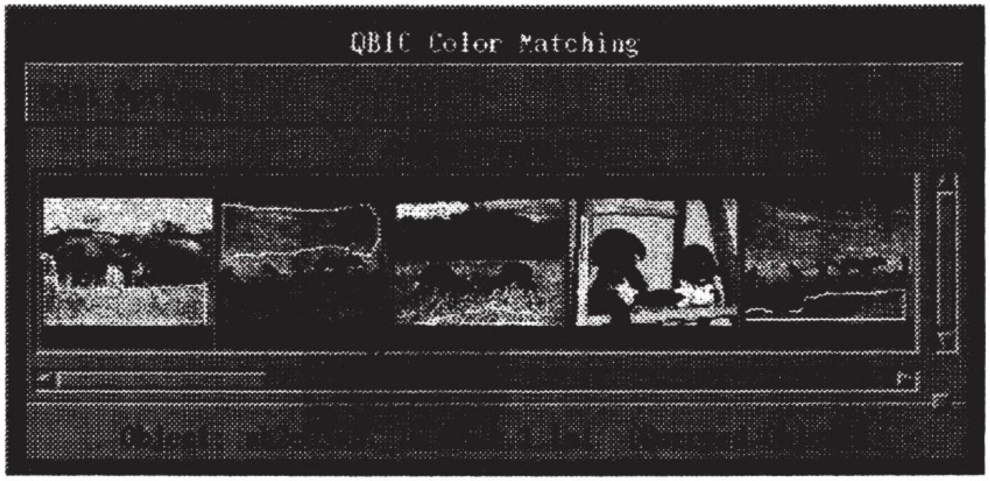 Figure 9: Result of query by example using color and texture.

Element	'004 Claim Recitation	Exemplary Citations in Reference
		Query by Image and Video Content: The QBIC System (IBM 0000893–902): "For both population and query, the QBIC data model has • still images or scenes (full images) that contain objects (subsets of an image), and • video shots that consist of sets of contiguous frames and contain motion objects. For still images, the QBIC data model distinguishes between 'scenes' (or images) and 'objects.' A scene is an image or single representative frame of video. An object is a part of a scene-for example, the fox in Figure 3-or a moving entity in a video. For still image database population, features are extracted from images and objects and stored in a database as shown in the top left part of Figure 2. Videos are broken into clips called shots. Representative frames, or r-frames, are generated for each extracted shot. R-frames are treated as still images, and features are extracted and stored in the database. Further processing of shots generates motion objects-for example, a car moving across the screen. Queries are allowed on objects ('Find images with a red, round object'), scenes ('Find images that have approximately 30-percent red and IS-percent blue colors'), shots ('Find all shots panning from left to right'), or any combination ('Find images that have 30 percent red and contain a blue textured object')." *Id*. at p. 3. Image, Audio, and Video Extenders Administration and Programming (IBM 000001–617): "To make an image available for searching by content, you catalog the image. When you catalog an image, the Image Extender analyzes the image, by computing the feature values for the image, and stores the values in a QBIC catalog. When you search for an image by content, your query identifies one or more features for the search (such as average color), a source for each feature (such as an example image), and a target set of cataloged images. The Image Extender computes the feature value of the source and compares it to the cataloged feature values for the target images. It then computes a score that indicates how

Element	'004 Claim Recitation	Exemplary Citations in Reference
		similar the feature values of the target images are to the source. You can have the Image Extender return the images whose features are most similar to the source. The Image Extender will return the handle of each image and the image score. You can also have the Image Extender return only the score of a single image." *Id.* at pp. 20–21. QBIC Demo (IBM000747) 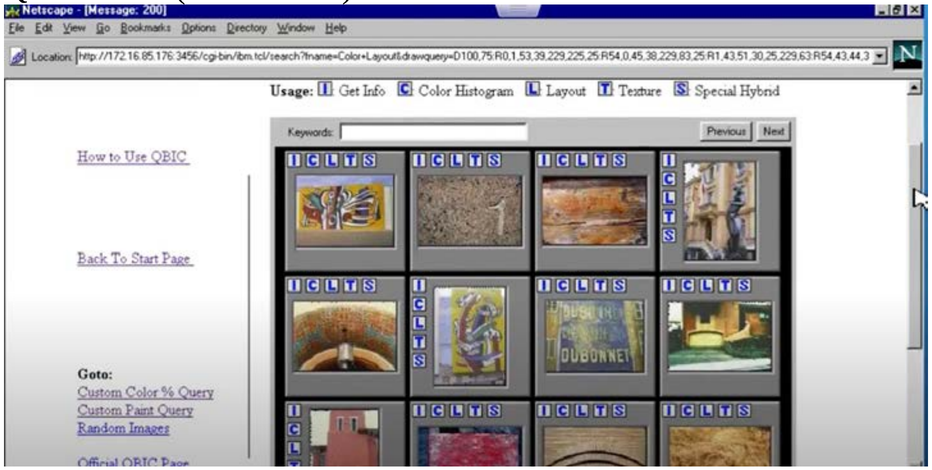 Location: http://172.16.85.176:3456/cgi-bin/ibm.tcl/search?fname=Color+Layout&drawquery=D100,75:R0,1,53,39,229,225,25:R54,0,45,38,229,83,25:R1,43,51,30,25,229,63:R54,43,44,3

Element	'004 Claim Recitation	Exemplary Citations in Reference
		QBIC Demo (IBM000747)  http://172.16.85.176:3456/cgi-bin/ibm.tcl/info?france50%2Ejpg Image france50.jpg Flickner Depo at 201-202 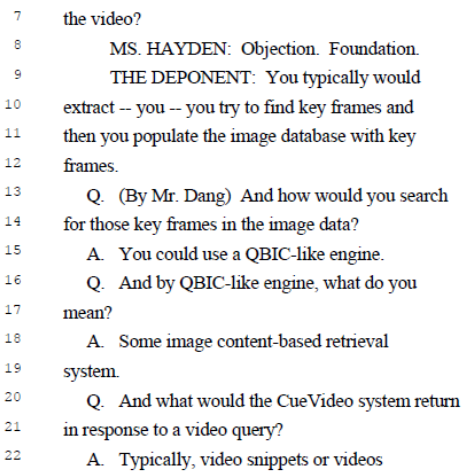 Q. Now, to search a video by content using CueVideo, how would the features be extracted from the video? MS. HAYDEN: Objection. Foundation. THE DEPONENT: You typically would extract -- you -- you try to find key frames and then you populate the image database with key frames. Q. (By Mr. Dang) And how would you search for those key frames in the image data? A. You could use a QBIC-like engine. Q. And by QBIC-like engine, what do you mean? A. Some image content-based retrieval system. Q. And what would the CueVideo system return in response to a video query? A. Typically, video snippets or videos

Element	'004 Claim Recitation	Exemplary Citations in Reference
		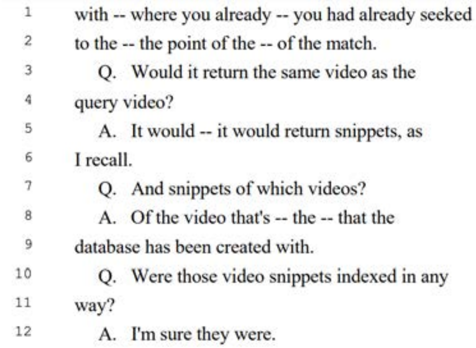 with -- where you already -- you had already seeked to the -- the point of the -- of the match. Q. Would it return the same video as the query video? A. It would -- it would return snippets, as I recall. Q. And snippets of which videos? A. Of the video that's -- the -- that the database has been created with. Q. Were those video snippets indexed in any way? A. I'm sure they were. Flickner Depo at 41-42 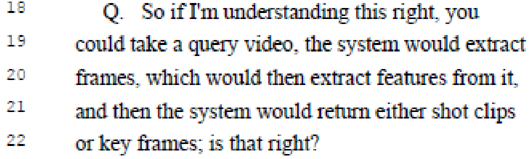 Q. So if I'm understanding this right, you could take a query video, the system would extract frames, which would then extract features from it, and then the system would return either shot clips or key frames; is that right? 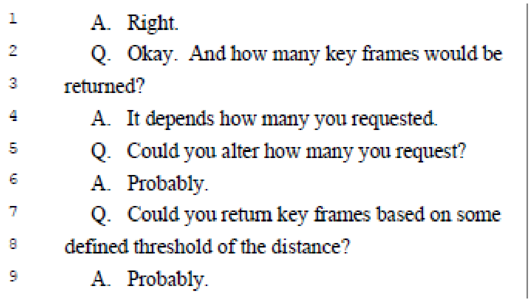 A. Right. Q. Okay. And how many key frames would be returned? A. It depends how many you requested. Q. Could you alter how many you request? A. Probably. Q. Could you return key frames based on some defined threshold of the distance? A. Probably.

Element	'004 Claim Recitation	Exemplary Citations in Reference
		Flickner Depo at 50 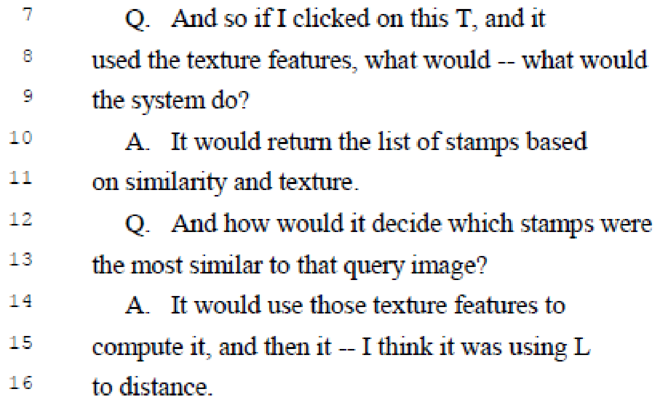 Q. And so if I clicked on this T, and it used the texture features, what would -- what would the system do? A. It would return the list of stamps based on similarity and texture. Q. And how would it decide which stamps were the most similar to that query image? A. It would use those texture features to compute it, and then it -- I think it was using L to distance. Flickner Depo at 162-163 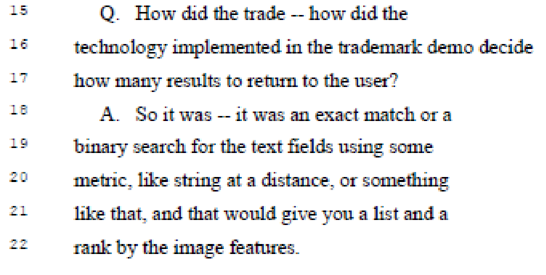 Q. How did the trade -- how did the technology implemented in the trademark demo decide how many results to return to the user? A. So it was -- it was an exact match or a binary search for the text fields using some metric, like string at a distance, or something like that, and that would give you a list and a rank by the image features.

Element	'004 Claim Recitation	Exemplary Citations in Reference
		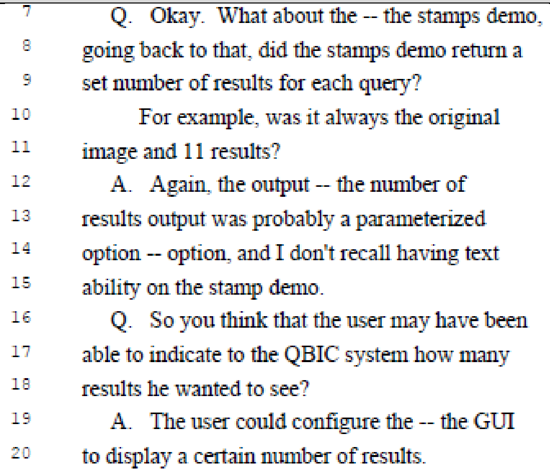Q. Okay. What about the -- the stamps demo, going back to that, did the stamps demo return a set number of results for each query? For example, was it always the original image and 11 results? A. Again, the output -- the number of results output was probably a parameterized option -- option, and I don't recall having text ability on the stamp demo. Q. So you think that the user may have been able to indicate to the QBIC system how many results he wanted to see? A. The user could configure the -- the GUI to display a certain number of results. Flickner Depo at 131-132 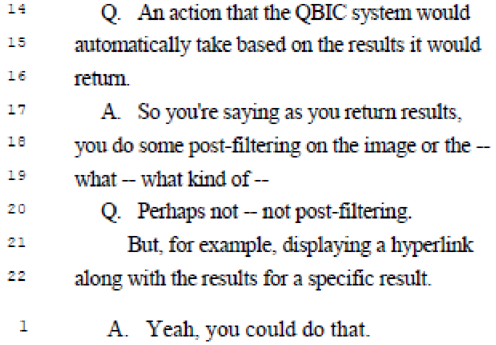Q. An action that the QBIC system would automatically take based on the results it would return. A. So you're saying as you return results, you do some post-filtering on the image or the -- what -- what kind of -- Q. Perhaps not -- not post-filtering. But, for example, displaying a hyperlink along with the results for a specific result. A. Yeah, you could do that. Flickner 8/29/23 Depo at 119–128

Element	’004 Claim Recitation	Exemplary Citations in Reference
		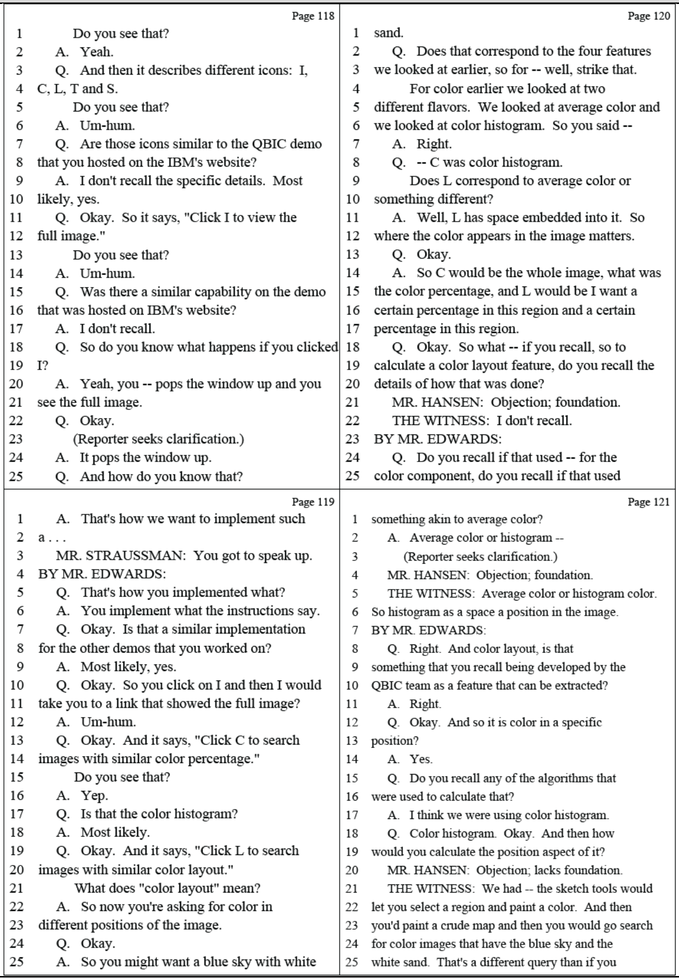

Page 118

Do you see that?
A. Yeah.
Q. And then it describes different icons: I, C, L, T and S.
Do you see that?
A. Um-hum.
Q. Are those icons similar to the QBIC demo that you hosted on the IBM's website?
A. I don't recall the specific details. Most likely, yes.
Q. Okay. So it says, "Click I to view the full image."
Do you see that?
A. Um-hum.
Q. Was there a similar capability on the demo that was hosted on IBM's website?
A. I don't recall.
Q. So do you know what happens if you clicked I?
A. Yeah, you -- pops the window up and you see the full image.
Q. Okay.
(Reporter seeks clarification.)
A. It pops the window up.
Q. And how do you know that?

Page 119

A. That's how we want to implement such a . . .
MR. STRAUSSMAN: You got to speak up.
BY MR. EDWARDS:
Q. That's how you implemented what?
A. You implement what the instructions say.
Q. Okay. Is that a similar implementation for the other demos that you worked on?
A. Most likely, yes.
Q. Okay. So you click on I and then I would take you to a link that showed the full image?
A. Um-hum.
Q. Okay. And it says, "Click C to search images with similar color percentage."
Do you see that?
A. Yep.
Q. Is that the color histogram?
A. Most likely.
Q. Okay. And it says, "Click L to search images with similar color layout."
What does "color layout" mean?
A. So now you're asking for color in different positions of the image.
Q. Okay.
A. So you might want a blue sky with white

Page 120

sand.
Q. Does that correspond to the four features we looked at earlier, so for -- well, strike that.
For color earlier we looked at two different flavors. We looked at average color and we looked at color histogram. So you said --
A. Right.
Q. -- C was color histogram.
Does L correspond to average color or something different?
A. Well, L has space embedded into it. So where the color appears in the image matters.
Q. Okay.
A. So C would be the whole image, what was the color percentage, and L would be I want a certain percentage in this region and a certain percentage in this region.
Q. Okay. So what -- if you recall, so to calculate a color layout feature, do you recall the details of how that was done?
MR. HANSEN: Objection; foundation.
THE WITNESS: I don't recall.
BY MR. EDWARDS:
Q. Do you recall if that used -- for the color component, do you recall if that used

Page 121

something akin to average color?
A. Average color or histogram --
(Reporter seeks clarification.)
MR. HANSEN: Objection; foundation.
THE WITNESS: Average color or histogram color. So histogram as a space a position in the image.
BY MR. EDWARDS:
Q. Right. And color layout, is that something that you recall being developed by the QBIC team as a feature that can be extracted?
A. Right.
Q. Okay. And so it is color in a specific position?
A. Yes.
Q. Do you recall any of the algorithms that were used to calculate that?
A. I think we were using color histogram.
Q. Color histogram. Okay. And then how would you calculate the position aspect of it?
MR. HANSEN: Objection; lacks foundation.
THE WITNESS: We had -- the sketch tools would let you select a region and paint a color. And then you'd paint a crude map and then you would go search for color images that have the blue sky and the white sand. That's a different query than if you

			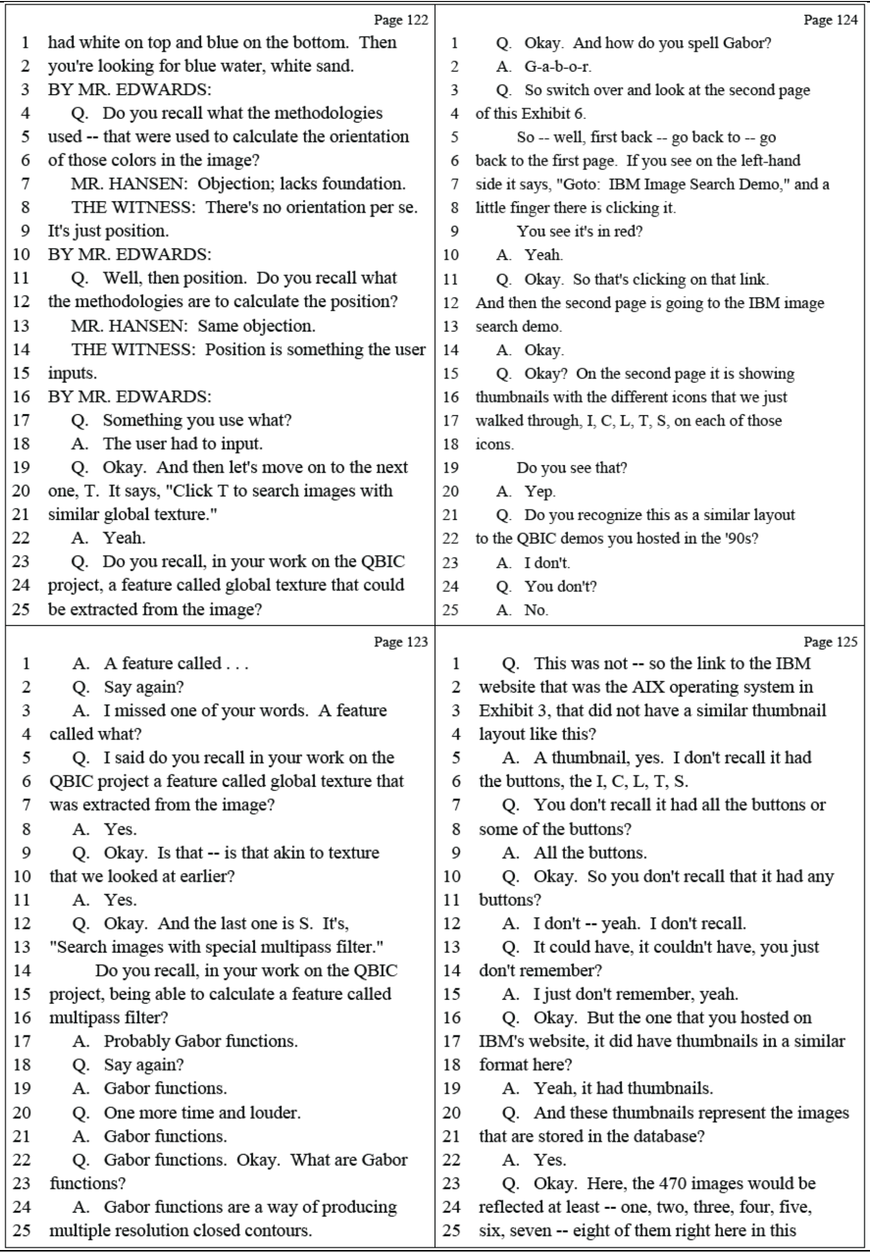 Page 122 had white on top and blue on the bottom. Then you're looking for blue water, white sand. BY MR. EDWARDS: Q. Do you recall what the methodologies used -- that were used to calculate the orientation of those colors in the image? MR. HANSEN: Objection; lacks foundation. THE WITNESS: There's no orientation per se. It's just position. BY MR. EDWARDS: Q. Well, then position. Do you recall what the methodologies are to calculate the position? MR. HANSEN: Same objection. THE WITNESS: Position is something the user inputs. BY MR. EDWARDS: Q. Something you use what? A. The user had to input. Q. Okay. And then let's move on to the next one, T. It says, "Click T to search images with similar global texture." A. Yeah. Q. Do you recall, in your work on the QBIC project, a feature called global texture that could be extracted from the image? Page 123 A. A feature called . . . Q. Say again? A. I missed one of your words. A feature called what? Q. I said do you recall in your work on the QBIC project a feature called global texture that was extracted from the image? A. Yes. Q. Okay. Is that -- is that akin to texture that we looked at earlier? A. Yes. Q. Okay. And the last one is S. It's, "Search images with special multipass filter." Do you recall, in your work on the QBIC project, being able to calculate a feature called multipass filter? A. Probably Gabor functions. Q. Say again? A. Gabor functions. Q. One more time and louder. A. Gabor functions. Q. Gabor functions. Okay. What are Gabor functions? A. Gabor functions are a way of producing multiple resolution closed contours. Page 124 Q. Okay. And how do you spell Gabor? A. G-a-b-o-r. Q. So switch over and look at the second page of this Exhibit 6. So -- well, first back -- go back to -- go back to the first page. If you see on the left-hand side it says, "Goto: IBM Image Search Demo," and a little finger there is clicking it. You see it's in red? A. Yeah. Q. Okay. So that's clicking on that link. And then the second page is going to the IBM image search demo. A. Okay. Q. Okay? On the second page it is showing thumbnails with the different icons that we just walked through, I, C, L, T, S, on each of those icons. Do you see that? A. Yep. Q. Do you recognize this as a similar layout to the QBIC demos you hosted in the '90s? A. I don't. Q. You don't? A. No. Page 125 Q. This was not -- so the link to the IBM website that was the AIX operating system in Exhibit 3, that did not have a similar thumbnail layout like this? A. A thumbnail, yes. I don't recall it had the buttons, the I, C, L, T, S. Q. You don't recall it had all the buttons or some of the buttons? A. All the buttons. Q. Okay. So you don't recall that it had any buttons? A. I don't -- yeah. I don't recall. Q. It could have, it couldn't have, you just don't remember? A. I just don't remember, yeah. Q. Okay. But the one that you hosted on IBM's website, it did have thumbnails in a similar format here? A. Yeah, it had thumbnails. Q. And these thumbnails represent the images that are stored in the database? A. Yes. Q. Okay. Here, the 470 images would be reflected at least -- one, two, three, four, five, six, seven -- eight of them right here in this

Element	’004 Claim Recitation	Exemplary Citations in Reference
		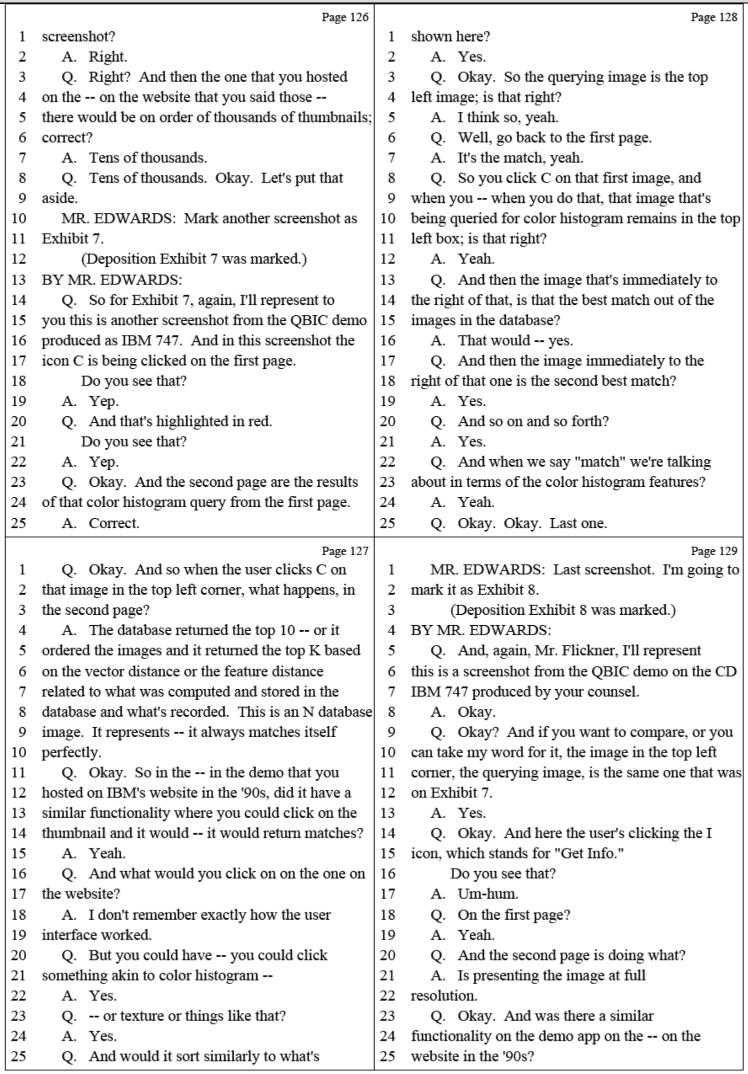 Page 126 screenshot? A. Right. Q. Right? And then the one that you hosted on the -- on the website that you said those -- there would be on order of thousands of thumbnails; correct? A. Tens of thousands. Q. Tens of thousands. Okay. Let's put that aside. MR. EDWARDS: Mark another screenshot as Exhibit 7. (Deposition Exhibit 7 was marked.) BY MR. EDWARDS: Q. So for Exhibit 7, again, I'll represent to you this is another screenshot from the QBIC demo produced as IBM 747. And in this screenshot the icon C is being clicked on the first page. Do you see that? A. Yep. Q. And that's highlighted in red. Do you see that? A. Yep. Q. Okay. And the second page are the results of that color histogram query from the first page. A. Correct. Page 127 Q. Okay. And so when the user clicks C on that image in the top left corner, what happens, in the second page? A. The database returned the top 10 -- or it ordered the images and it returned the top K based on the vector distance or the feature distance related to what was computed and stored in the database and what's recorded. This is an N database image. It represents -- it always matches itself perfectly. Q. Okay. So in the -- in the demo that you hosted on IBM's website in the '90s, did it have a similar functionality where you could click on the thumbnail and it would -- it would return matches? A. Yeah. Q. And what would you click on on the one on the website? A. I don't remember exactly how the user interface worked. Q. But you could have -- you could click something akin to color histogram -- A. Yes. Q. -- or texture or things like that? A. Yes. Q. And would it sort similarly to what's Page 128 shown here? A. Yes. Q. Okay. So the querying image is the top left image; is that right? A. I think so, yeah. Q. Well, go back to the first page. A. It's the match, yeah. Q. So you click C on that first image, and when you -- when you do that, that image that's being queried for color histogram remains in the top left box; is that right? A. Yeah. Q. And then the image that's immediately to the right of that, is that the best match out of the images in the database? A. That would -- yes. Q. And then the image immediately to the right of that one is the second best match? A. Yes. Q. And so on and so forth? A. Yes. Q. And when we say "match" we're talking about in terms of the color histogram features? A. Yeah. Q. Okay. Okay. Last one. Page 129 MR. EDWARDS: Last screenshot. I'm going to mark it as Exhibit 8. (Deposition Exhibit 8 was marked.) BY MR. EDWARDS: Q. And, again, Mr. Flickner, I'll represent this is a screenshot from the QBIC demo on the CD IBM 747 produced by your counsel. A. Okay. Q. Okay? And if you want to compare, or you can take my word for it, the image in the top left corner, the querying image, is the same one that was on Exhibit 7. A. Yes. Q. Okay. And here the user's clicking the I icon, which stands for "Get Info." Do you see that? A. Um-hum. Q. On the first page? A. Yeah. Q. And the second page is doing what? A. Is presenting the image at full resolution. Q. Okay. And was there a similar functionality on the demo app on the -- on the website in the '90s? *Id.* at 144

Element	'004 Claim Recitation	Exemplary Citations in Reference
		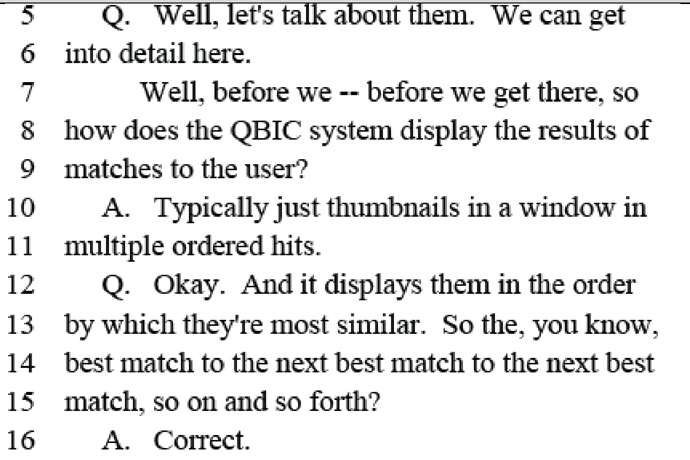 Q. Well, let's talk about them. We can get into detail here. Well, before we -- before we get there, so how does the QBIC system display the results of matches to the user? A. Typically just thumbnails in a window in multiple ordered hits. Q. Okay. And it displays them in the order by which they're most similar. So the, you know, best match to the next best match to the next best match, so on and so forth? A. Correct. *Id.* at 160 – 161 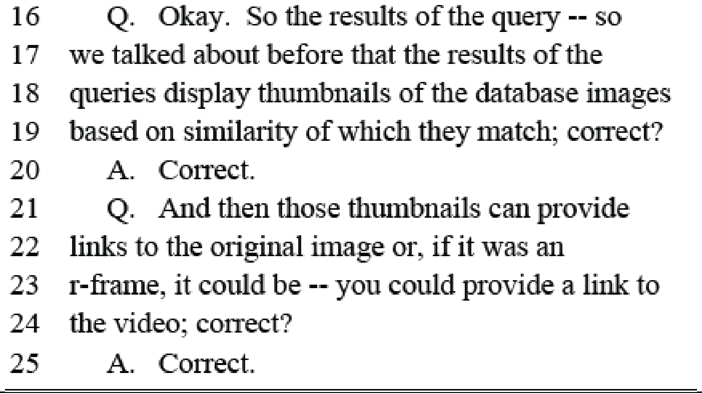 Q. Okay. So the results of the query -- so we talked about before that the results of the queries display thumbnails of the database images based on similarity of which they match; correct? A. Correct. Q. And then those thumbnails can provide links to the original image or, if it was an r-frame, it could be -- you could provide a link to the video; correct? A. Correct.

Element	'004 Claim Recitation	Exemplary Citations in Reference
		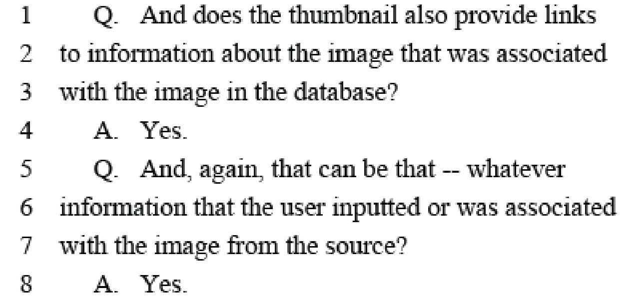 Q. And does the thumbnail also provide links to information about the image that was associated with the image in the database? A. Yes. Q. And, again, that can be that -- whatever information that the user inputted or was associated with the image from the source? A. Yes. Flickner 8/29/23 Depo Ex. 7 at 1 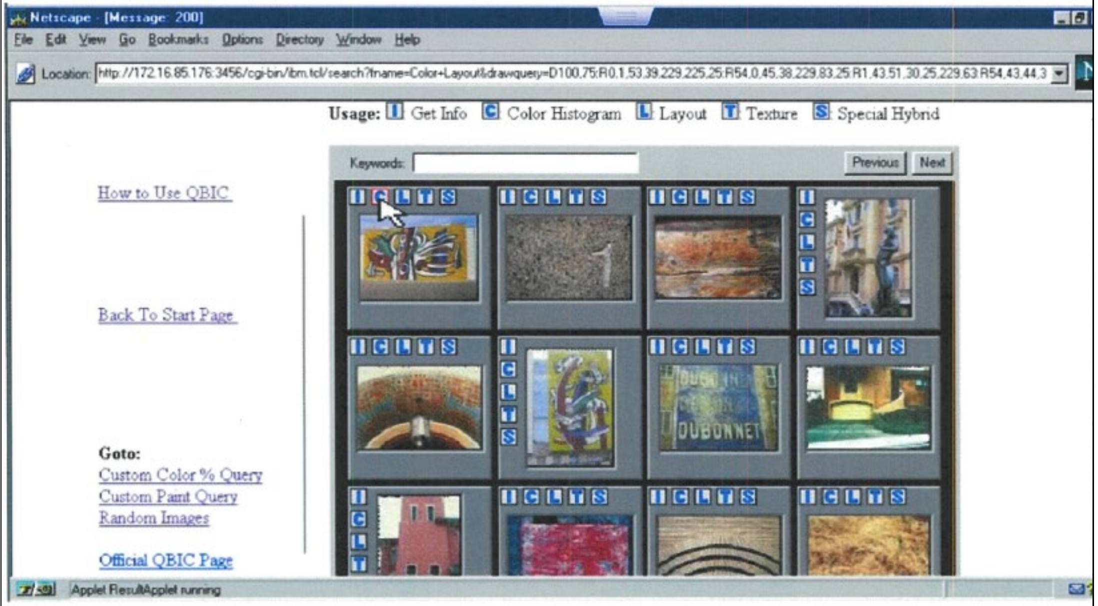 Location: http://172.16.85.176:3456/cgi-bin/ibm.tcl/search?fname=Color+Layout&drawquery=D100,75:R0,1,53,39,229,225,25:R54,0,45,38,229,83,25:R1,43,51,30,25,229,63:R54,43,44,3 *Id.* at 2

Element	'004 Claim Recitation	Exemplary Citations in Reference
		 Location: http://172.16.85.176:3456/cgi-bin/ibm.tcl/search?fname=Color+Histogram&france50%2Ejpg.x=0 Flickner 8/29/23 Depo at Ex. 8

Element	'004 Claim Recitation	Exemplary Citations in Reference
		 Location: http://172.16.85.176:3456/cgi-bin/ibm.tcl/search?fname=Color+Histogram&france50%2Ejpg.x=0 Applet ResultApplet running *Id.* at 2

Element	'004 Claim Recitation	Exemplary Citations in Reference
		 http://172.16.85.176:3456/cgi-bin/ibm.tcl/info?france50%2Ejpg Image france50.jpg *Id.* at 3

Element	'004 Claim Recitation	Exemplary Citations in Reference
		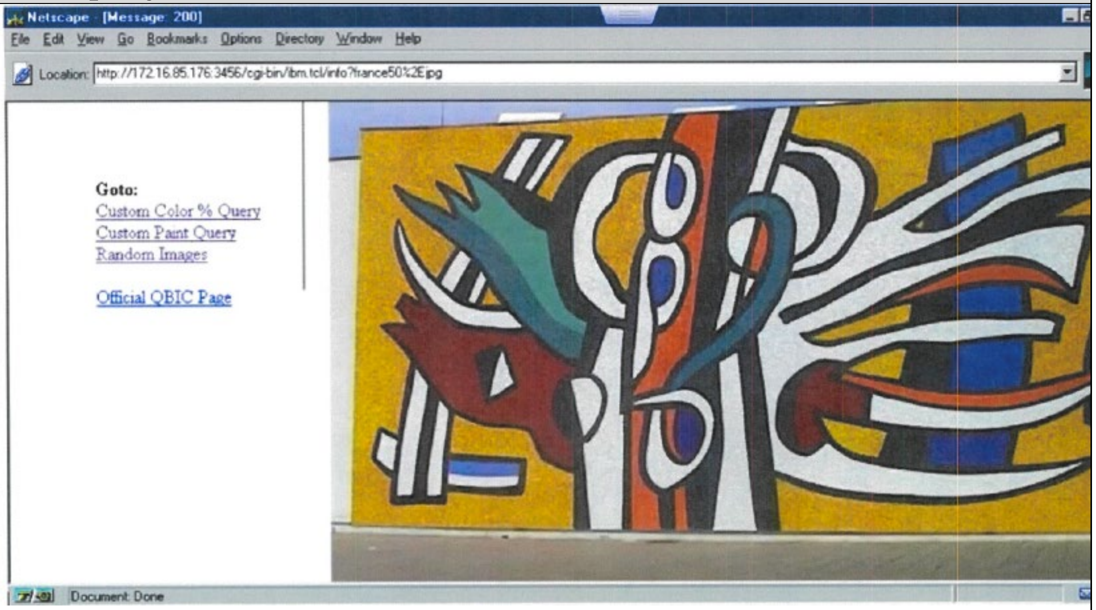 Location: http://172.16.85.176:3456/cgi-bin/ibm.tcl/info?france50%2Ejpg Goto: Custom Color % Query Custom Paint Query Random Images Official QBIC Page Flickner 8/29/23 Depo at Ex. 9

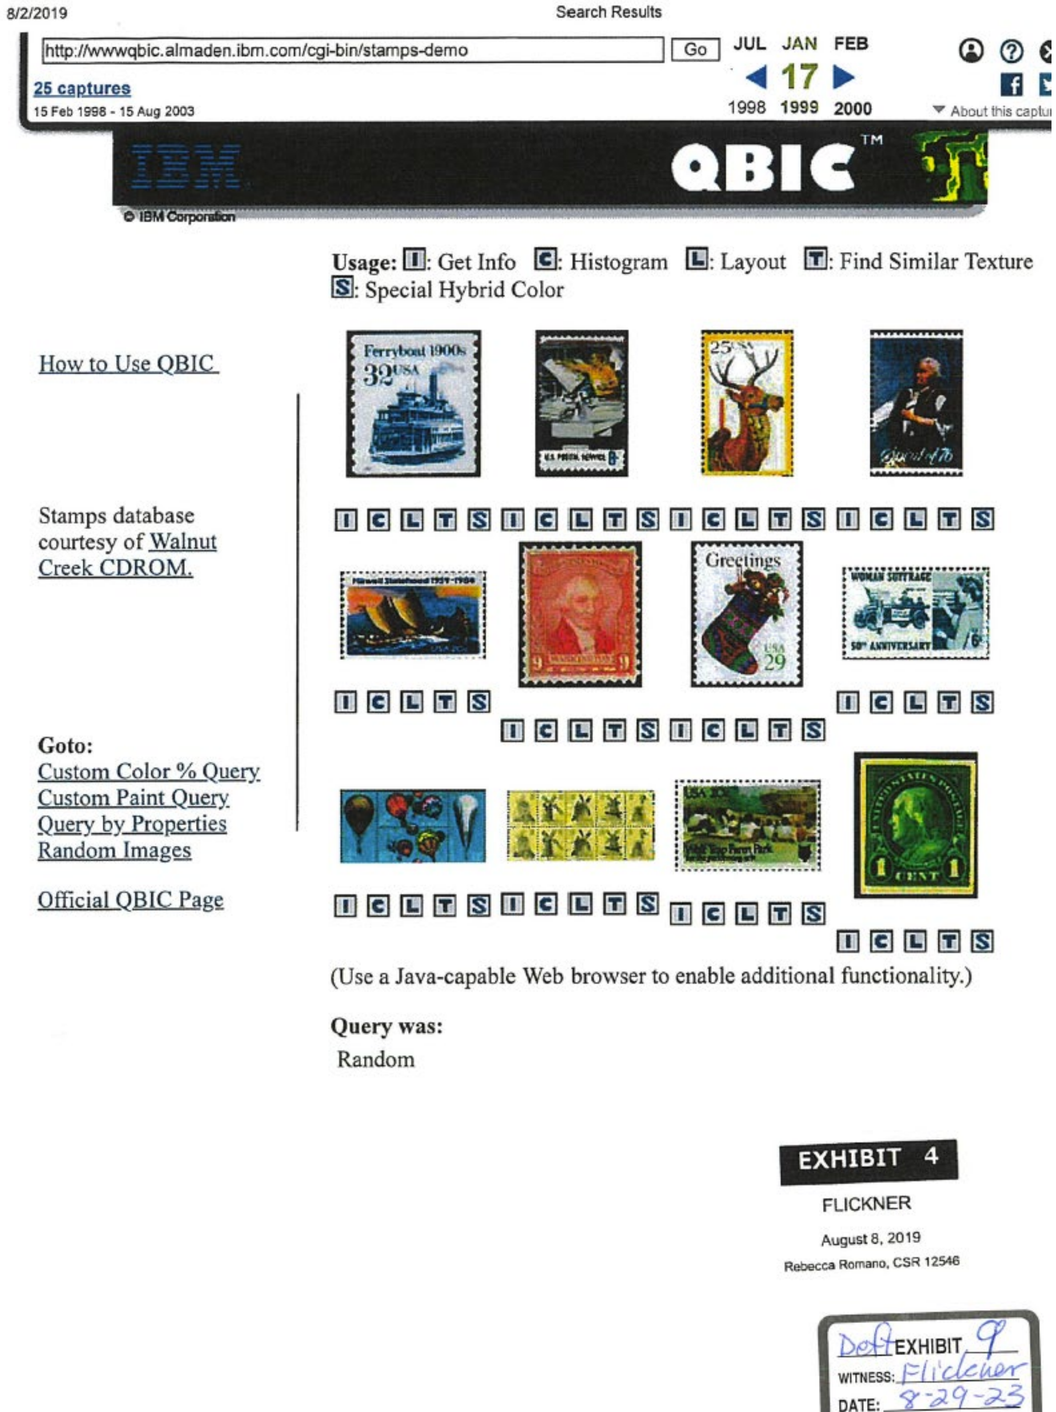

8/2/2019 Search Results

http://wwwqbic.almaden.ibm.com/cgi-bin/stamps-demo

Usage: I: Get Info C: Histogram L: Layout T: Find Similar Texture S: Special Hybrid Color

How to Use QBIC

Stamps database courtesy of Walnut Creek CDROM.

Goto:
Custom Color % Query
Custom Paint Query
Query by Properties
Random Images

Official QBIC Page

(Use a Java-capable Web browser to enable additional functionality.)

Query was:
Random

EXHIBIT 4
FLICKNER
August 8, 2019
Rebecca Romano, CSR 12546

Deft EXHIBIT 9
WITNESS: Flickner
DATE: 8-29-23

Element	'004 Claim Recitation	Exemplary Citations in Reference
		See also Flickner Depo at 22-25, 25-26, 38-41, 209-210 To the extent it is found that QBIC does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that QBIC does not anticipate this claim, QBIC renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine QBIC with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1E]	presenting to a user, via an address, content information associated with the target object.	QBIC discloses the limitation. The QBIC Project: Querying Images by Content Using Color, Texture, and Shape (IBM 0002390–404) at 2396

Element	'004 Claim Recitation	Exemplary Citations in Reference
		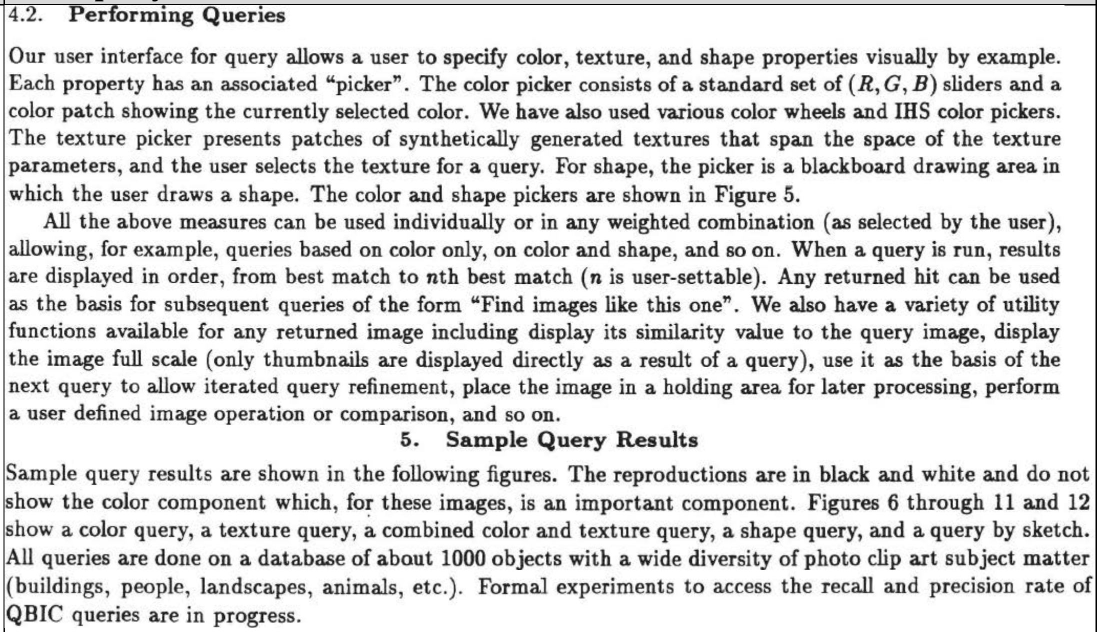 **4.2. Performing Queries** Our user interface for query allows a user to specify color, texture, and shape properties visually by example. Each property has an associated "picker". The color picker consists of a standard set of (R, G, B) sliders and a color patch showing the currently selected color. We have also used various color wheels and IHS color pickers. The texture picker presents patches of synthetically generated textures that span the space of the texture parameters, and the user selects the texture for a query. For shape, the picker is a blackboard drawing area in which the user draws a shape. The color and shape pickers are shown in Figure 5. All the above measures can be used individually or in any weighted combination (as selected by the user), allowing, for example, queries based on color only, on color and shape, and so on. When a query is run, results are displayed in order, from best match to nth best match (n is user-settable). Any returned hit can be used as the basis for subsequent queries of the form "Find images like this one". We also have a variety of utility functions available for any returned image including display its similarity value to the query image, display the image full scale (only thumbnails are displayed directly as a result of a query), use it as the basis of the next query to allow iterated query refinement, place the image in a holding area for later processing, perform a user defined image operation or comparison, and so on. **5. Sample Query Results** Sample query results are shown in the following figures. The reproductions are in black and white and do not show the color component which, for these images, is an important component. Figures 6 through 11 and 12 show a color query, a texture query, a combined color and texture query, a shape query, and a query by sketch. All queries are done on a database of about 1000 objects with a wide diversity of photo clip art subject matter (buildings, people, landscapes, animals, etc.). Formal experiments to access the recall and precision rate of QBIC queries are in progress. IBM0002390 at 2391-92 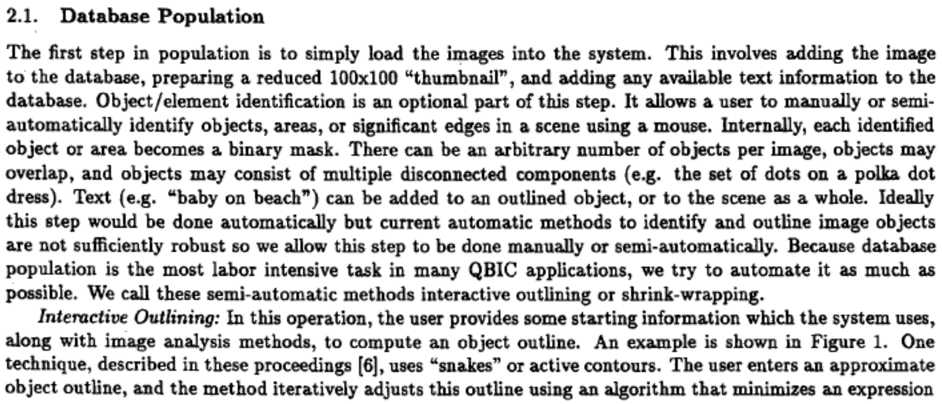 **2.1. Database Population** The first step in population is to simply load the images into the system. This involves adding the image to the database, preparing a reduced 100x100 "thumbnail", and adding any available text information to the database. Object/element identification is an optional part of this step. It allows a user to manually or semi-automatically identify objects, areas, or significant edges in a scene using a mouse. Internally, each identified object or area becomes a binary mask. There can be an arbitrary number of objects per image, objects may overlap, and objects may consist of multiple disconnected components (e.g. the set of dots on a polka dot dress). Text (e.g. "baby on beach") can be added to an outlined object, or to the scene as a whole. Ideally this step would be done automatically but current automatic methods to identify and outline image objects are not sufficiently robust so we allow this step to be done manually or semi-automatically. Because database population is the most labor intensive task in many QBIC applications, we try to automate it as much as possible. We call these semi-automatic methods interactive outlining or shrink-wrapping. *Interactive Outlining:* In this operation, the user provides some starting information which the system uses, along with image analysis methods, to compute an object outline. An example is shown in Figure 1. One technique, described in these proceedings [6], uses "snakes" or active contours. The user enters an approximate object outline, and the method iteratively adjusts this outline using an algorithm that minimizes an expression

Element	'004 Claim Recitation	Exemplary Citations in Reference
		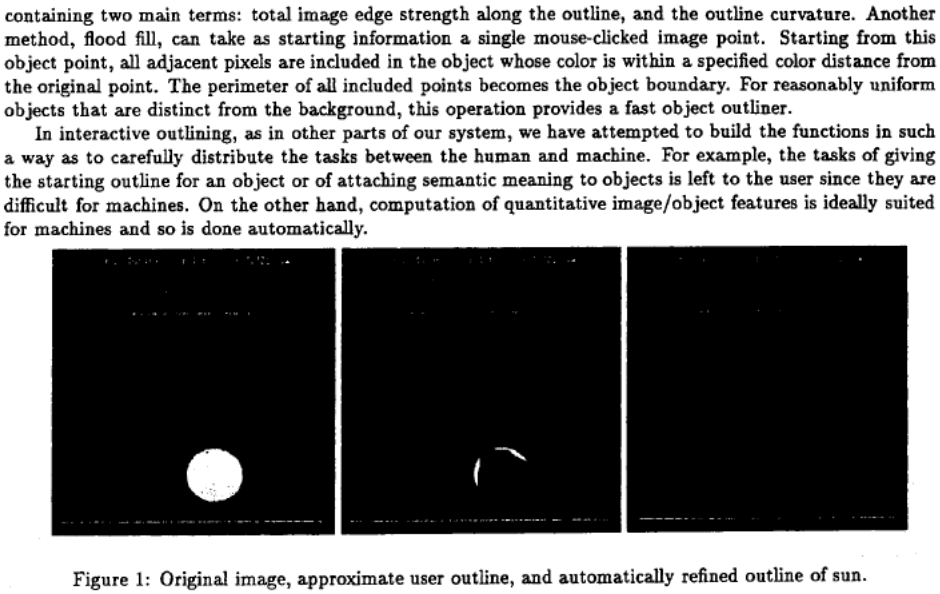 containing two main terms: total image edge strength along the outline, and the outline curvature. Another method, flood fill, can take as starting information a single mouse-clicked image point. Starting from this object point, all adjacent pixels are included in the object whose color is within a specified color distance from the original point. The perimeter of all included points becomes the object boundary. For reasonably uniform objects that are distinct from the background, this operation provides a fast object outliner. In interactive outlining, as in other parts of our system, we have attempted to build the functions in such a way as to carefully distribute the tasks between the human and machine. For example, the tasks of giving the starting outline for an object or of attaching semantic meaning to objects is left to the user since they are difficult for machines. On the other hand, computation of quantitative image/object features is ideally suited for machines and so is done automatically. Figure 1: Original image, approximate user outline, and automatically refined outline of sun. The QBIC Project: Querying Images by Content Using Color, Texture, and Shape (IBM 0002390–404) at Figs. 7–9

Element	'004 Claim Recitation	Exemplary Citations in Reference
		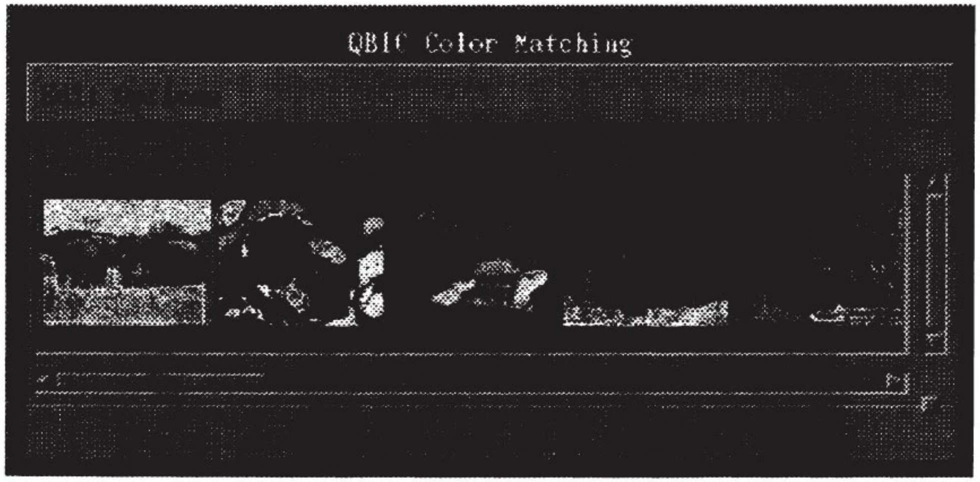 Figure 7: Result of query by example using color only.

Element	'004 Claim Recitation	Exemplary Citations in Reference
		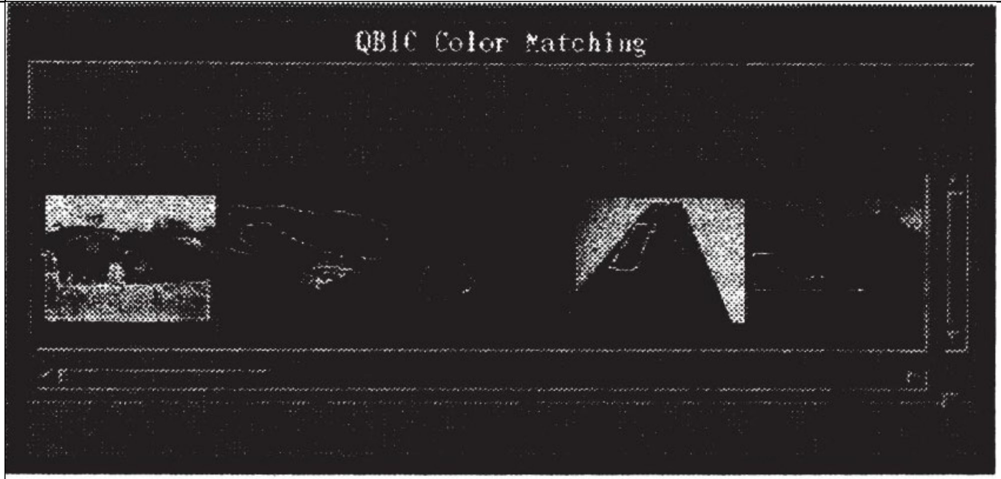 Figure 8: Result of query by example using texture only.

Element	'004 Claim Recitation	Exemplary Citations in Reference
		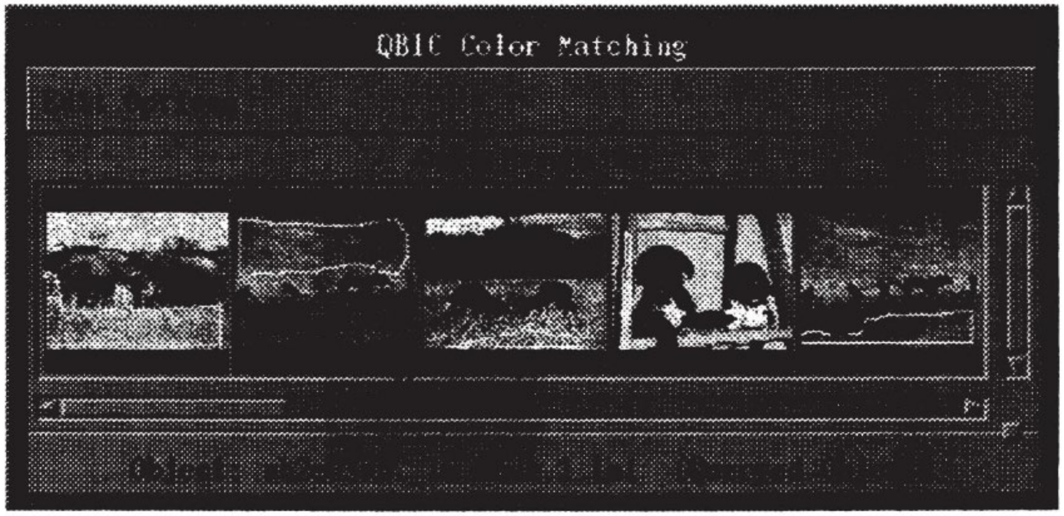 Figure 9: Result of query by example using color and texture. Query by Image and Video Content: The QBIC System (IBM 0000893–902): "To achieve this functionality, QBIC has two main components: database population (the process of creating an image database) and database query. During the population, images and videos are processed to extract features describing their content-colors, textures, shapes, and camera and object motion- and the features are stored in a database. During the query, the user composes a query graphically. Features are generated from the graphical query and then input to a matching engine that finds images or videos from the database with similar features. Figure 2 shows the system architecture." *Id.* at p. 3.

Element	'004 Claim Recitation	Exemplary Citations in Reference
		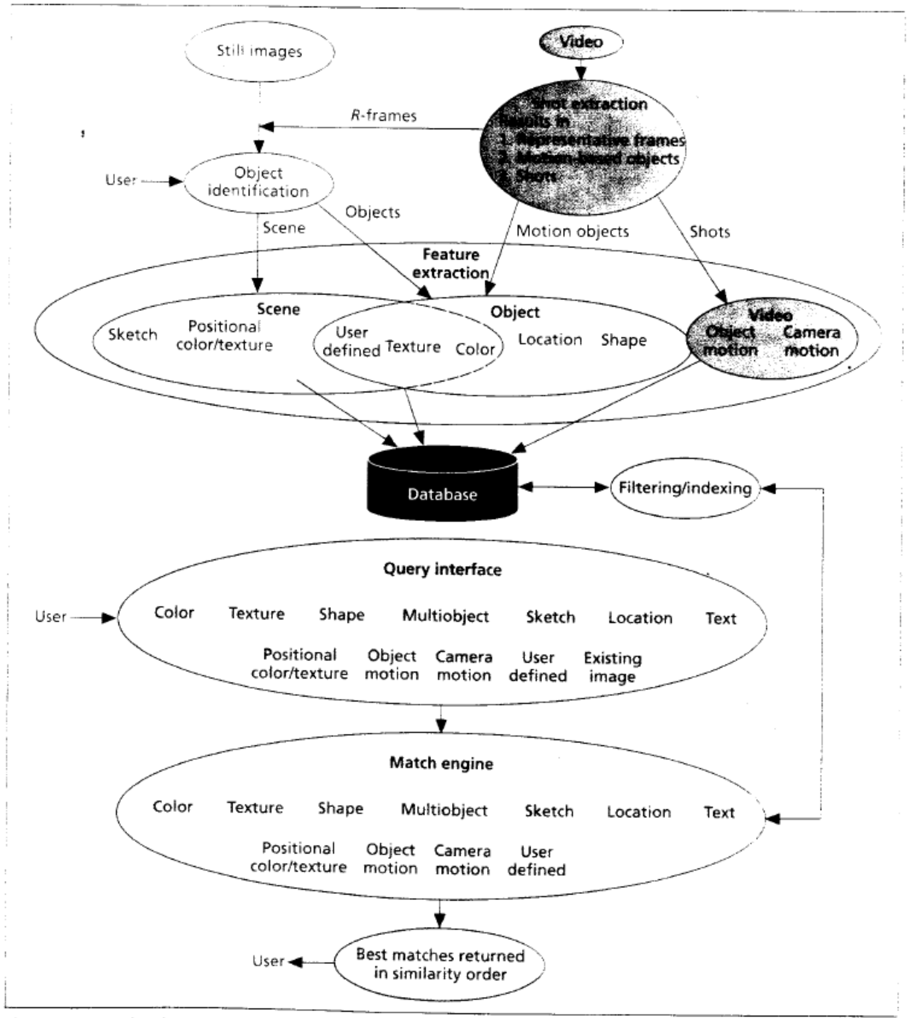 **Figure 2.** QBIC database population (top) and query (bottom) architecture. *Id*. at Fig. 2.

Element	'004 Claim Recitation	Exemplary Citations in Reference
		Image, Audio, and Video Extenders Administration and Programming (IBM 000001–617): "To make an image available for searching by content, you catalog the image. When you catalog an image, the Image Extender analyzes the image, by computing the feature values for the image, and stores the values in a QBIC catalog. When you search for an image by content, your query identifies one or more features for the search (such as average color), a source for each feature (such as an example image), and a target set of cataloged images. The Image Extender computes the feature value of the source and compares it to the cataloged feature values for the target images. It then computes a score that indicates how similar the feature values of the target images are to the source. You can have the Image Extender return the images whose features are most similar to the source. The Image Extender will return the handle of each image and the image score. You can also have the Image Extender return only the score of a single image." *Id.* at pp. 20–21. QBIC Demo (IBM000747) 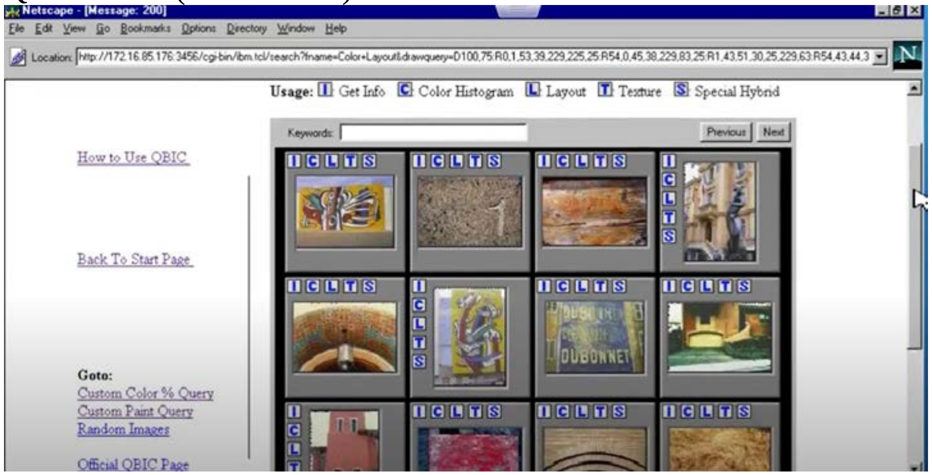 Location: http://172.16.85.176:3456/cgi-bin/ibm.tcl/search?fname=Color+Layout&drawquery=D100,75:R0,1,53,39,229,225,25:R54,0,45,38,229,83,25:R1,43,51,30,25,229,63:R54,43,44,3

Element	'004 Claim Recitation	Exemplary Citations in Reference
		QBIC Demo (IBM000747)  Location: http://172.16.85.176:3456/cgi-bin/ibm.tcl/info?france50%2Ejpg Image france50.jpg Flickner Depo at 131-132 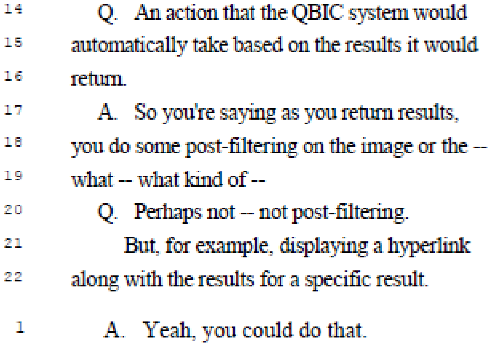 Q. An action that the QBIC system would automatically take based on the results it would return. A. So you're saying as you return results, you do some post-filtering on the image or the -- what -- what kind of -- Q. Perhaps not -- not post-filtering. But, for example, displaying a hyperlink along with the results for a specific result. A. Yeah, you could do that. Flickner 8/29/23 Depo at 53–56

Element	'004 Claim Recitation	Exemplary Citations in Reference
		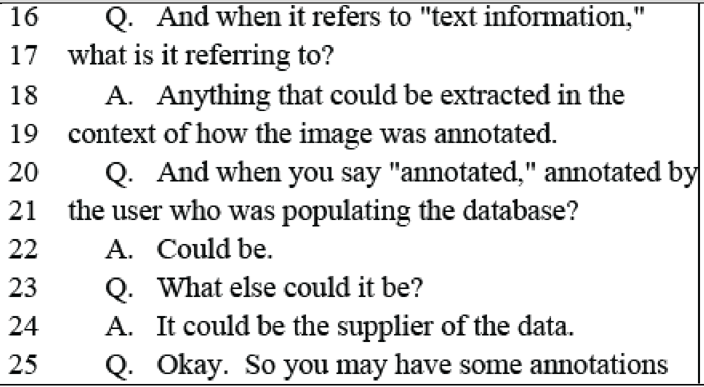 Q. And when it refers to "text information," what is it referring to? A. Anything that could be extracted in the context of how the image was annotated. Q. And when you say "annotated," annotated by the user who was populating the database? A. Could be. Q. What else could it be? A. It could be the supplier of the data. Q. Okay. So you may have some annotations

Element	'004 Claim Recitation	Exemplary Citations in Reference
		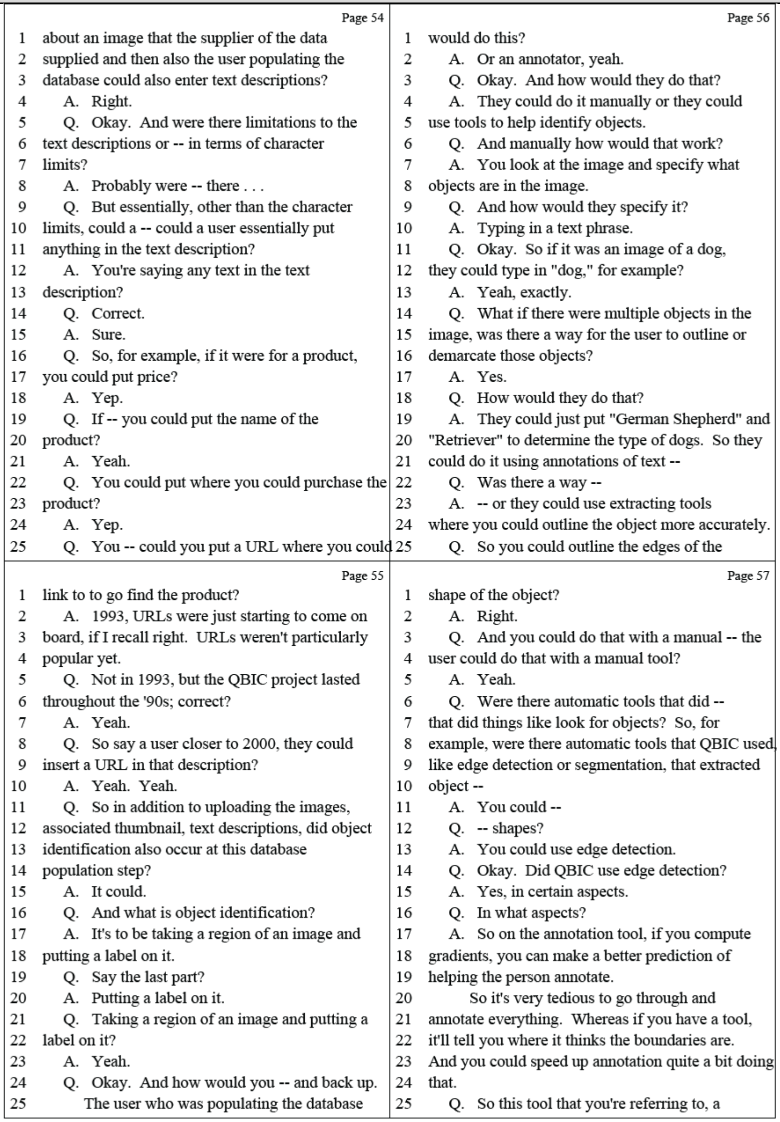

Page 54

about an image that the supplier of the data supplied and then also the user populating the database could also enter text descriptions?
A. Right.
Q. Okay. And were there limitations to the text descriptions or -- in terms of character limits?
A. Probably were -- there . . .
Q. But essentially, other than the character limits, could a -- could a user essentially put anything in the text description?
A. You're saying any text in the text description?
Q. Correct.
A. Sure.
Q. So, for example, if it were for a product, you could put price?
A. Yep.
Q. If -- you could put the name of the product?
A. Yeah.
Q. You could put where you could purchase the product?
A. Yep.
Q. You -- could you put a URL where you could

Page 55

link to to go find the product?
A. 1993, URLs were just starting to come on board, if I recall right. URLs weren't particularly popular yet.
Q. Not in 1993, but the QBIC project lasted throughout the '90s; correct?
A. Yeah.
Q. So say a user closer to 2000, they could insert a URL in that description?
A. Yeah. Yeah.
Q. So in addition to uploading the images, associated thumbnail, text descriptions, did object identification also occur at this database population step?
A. It could.
Q. And what is object identification?
A. It's to be taking a region of an image and putting a label on it.
Q. Say the last part?
A. Putting a label on it.
Q. Taking a region of an image and putting a label on it?
A. Yeah.
Q. Okay. And how would you -- and back up.
The user who was populating the database

Page 56

would do this?
A. Or an annotator, yeah.
Q. Okay. And how would they do that?
A. They could do it manually or they could use tools to help identify objects.
Q. And manually how would that work?
A. You look at the image and specify what objects are in the image.
Q. And how would they specify it?
A. Typing in a text phrase.
Q. Okay. So if it was an image of a dog, they could type in "dog," for example?
A. Yeah, exactly.
Q. What if there were multiple objects in the image, was there a way for the user to outline or demarcate those objects?
A. Yes.
Q. How would they do that?
A. They could just put "German Shepherd" and "Retriever" to determine the type of dogs. So they could do it using annotations of text --
Q. Was there a way --
A. -- or they could use extracting tools where you could outline the object more accurately.
Q. So you could outline the edges of the

Page 57

shape of the object?
A. Right.
Q. And you could do that with a manual -- the user could do that with a manual tool?
A. Yeah.
Q. Were there automatic tools that did -- that did things like look for objects? So, for example, were there automatic tools that QBIC used, like edge detection or segmentation, that extracted object --
A. You could --
Q. -- shapes?
A. You could use edge detection.
Q. Okay. Did QBIC use edge detection?
A. Yes, in certain aspects.
Q. In what aspects?
A. So on the annotation tool, if you compute gradients, you can make a better prediction of helping the person annotate.
So it's very tedious to go through and annotate everything. Whereas if you have a tool, it'll tell you where it thinks the boundaries are. And you could speed up annotation quite a bit doing that.
Q. So this tool that you're referring to, a

Element	'004 Claim Recitation	Exemplary Citations in Reference
		Id. at 68–72 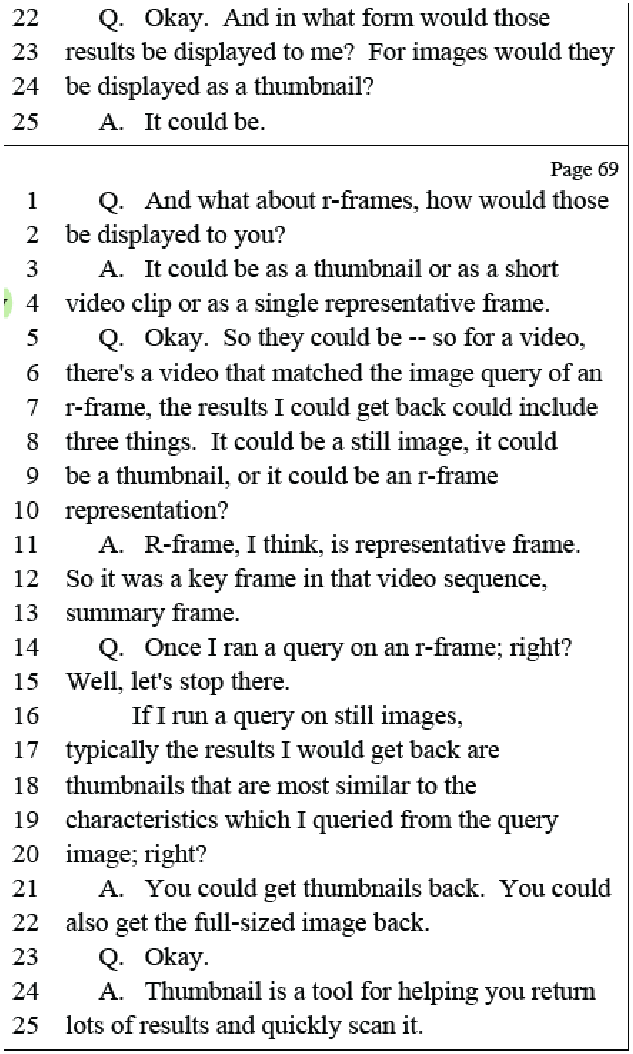 Q. Okay. And in what form would those results be displayed to me? For images would they be displayed as a thumbnail? A. It could be. Page 69 Q. And what about r-frames, how would those be displayed to you? A. It could be as a thumbnail or as a short video clip or as a single representative frame. Q. Okay. So they could be -- so for a video, there's a video that matched the image query of an r-frame, the results I could get back could include three things. It could be a still image, it could be a thumbnail, or it could be an r-frame representation? A. R-frame, I think, is representative frame. So it was a key frame in that video sequence, summary frame. Q. Once I ran a query on an r-frame; right? Well, let's stop there. If I run a query on still images, typically the results I would get back are thumbnails that are most similar to the characteristics which I queried from the query image; right? A. You could get thumbnails back. You could also get the full-sized image back. Q. Okay. A. Thumbnail is a tool for helping you return lots of results and quickly scan it.

Element	'004 Claim Recitation	Exemplary Citations in Reference
		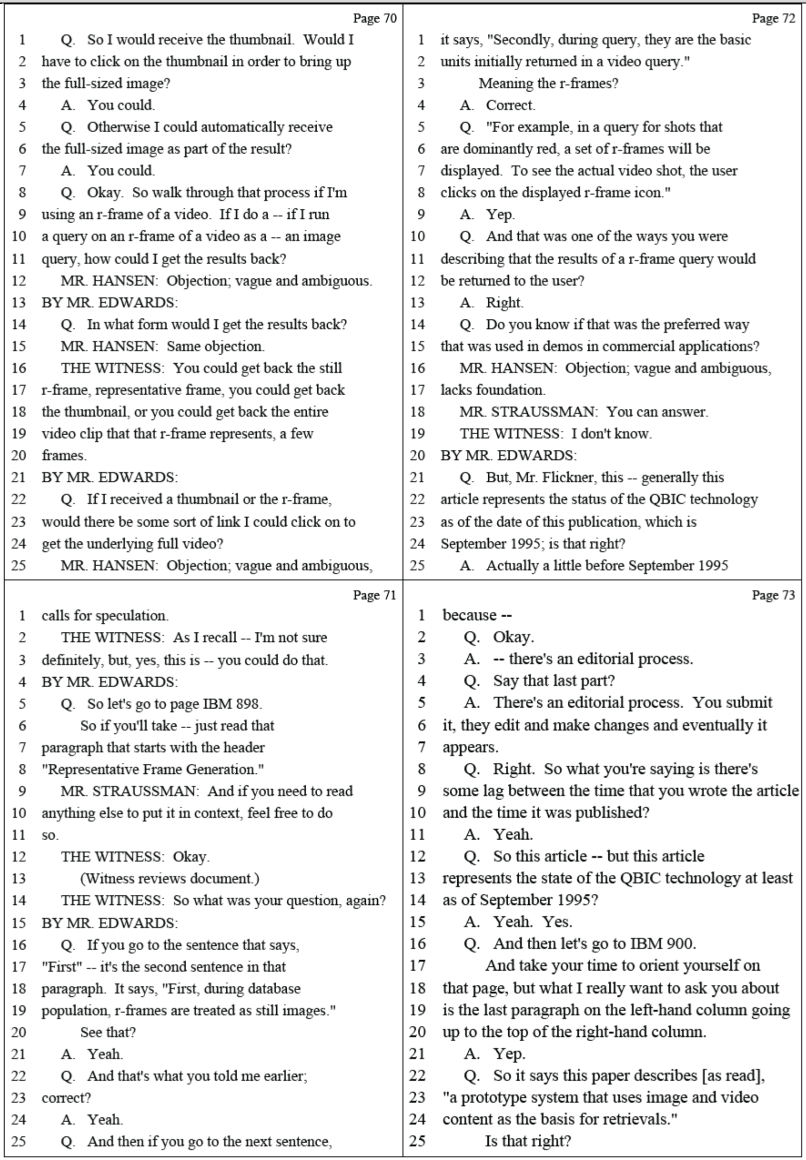

Page 70

Q. So I would receive the thumbnail. Would I have to click on the thumbnail in order to bring up the full-sized image?
A. You could.
Q. Otherwise I could automatically receive the full-sized image as part of the result?
A. You could.
Q. Okay. So walk through that process if I'm using an r-frame of a video. If I do a -- if I run a query on an r-frame of a video as a -- an image query, how could I get the results back?
MR. HANSEN: Objection; vague and ambiguous.
BY MR. EDWARDS:
Q. In what form would I get the results back?
MR. HANSEN: Same objection.
THE WITNESS: You could get back the still r-frame, representative frame, you could get back the thumbnail, or you could get back the entire video clip that that r-frame represents, a few frames.
BY MR. EDWARDS:
Q. If I received a thumbnail or the r-frame, would there be some sort of link I could click on to get the underlying full video?
MR. HANSEN: Objection; vague and ambiguous,

Page 71

calls for speculation.
THE WITNESS: As I recall -- I'm not sure definitely, but, yes, this is -- you could do that.
BY MR. EDWARDS:
Q. So let's go to page IBM 898.
So if you'll take -- just read that paragraph that starts with the header "Representative Frame Generation."
MR. STRAUSSMAN: And if you need to read anything else to put it in context, feel free to do so.
THE WITNESS: Okay.
(Witness reviews document.)
THE WITNESS: So what was your question, again?
BY MR. EDWARDS:
Q. If you go to the sentence that says, "First" -- it's the second sentence in that paragraph. It says, "First, during database population, r-frames are treated as still images."
See that?
A. Yeah.
Q. And that's what you told me earlier; correct?
A. Yeah.
Q. And then if you go to the next sentence,

Page 72

it says, "Secondly, during query, they are the basic units initially returned in a video query."
Meaning the r-frames?
A. Correct.
Q. "For example, in a query for shots that are dominantly red, a set of r-frames will be displayed. To see the actual video shot, the user clicks on the displayed r-frame icon."
A. Yep.
Q. And that was one of the ways you were describing that the results of a r-frame query would be returned to the user?
A. Right.
Q. Do you know if that was the preferred way that was used in demos in commercial applications?
MR. HANSEN: Objection; vague and ambiguous, lacks foundation.
MR. STRAUSSMAN: You can answer.
THE WITNESS: I don't know.
BY MR. EDWARDS:
Q. But, Mr. Flickner, this -- generally this article represents the status of the QBIC technology as of the date of this publication, which is September 1995; is that right?
A. Actually a little before September 1995

Page 73

because --
Q. Okay.
A. -- there's an editorial process.
Q. Say that last part?
A. There's an editorial process. You submit it, they edit and make changes and eventually it appears.
Q. Right. So what you're saying is there's some lag between the time that you wrote the article and the time it was published?
A. Yeah.
Q. So this article -- but this article represents the state of the QBIC technology at least as of September 1995?
A. Yeah. Yes.
Q. And then let's go to IBM 900.
And take your time to orient yourself on that page, but what I really want to ask you about is the last paragraph on the left-hand column going up to the top of the right-hand column.
A. Yep.
Q. So it says this paper describes [as read], "a prototype system that uses image and video content as the basis for retrievals."
Is that right?

Element	'004 Claim Recitation	Exemplary Citations in Reference
		Id. at 118–119 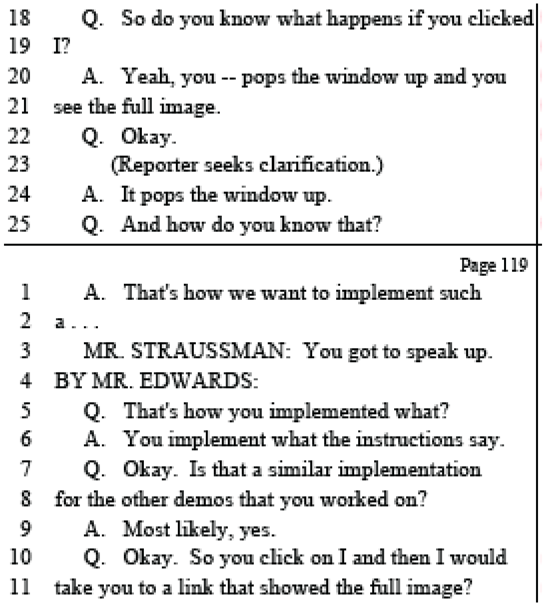 Q. So do you know what happens if you clicked I? A. Yeah, you -- pops the window up and you see the full image. Q. Okay. (Reporter seeks clarification.) A. It pops the window up. Q. And how do you know that? Page 119 A. That's how we want to implement such a . . . MR. STRAUSSMAN: You got to speak up. BY MR. EDWARDS: Q. That's how you implemented what? A. You implement what the instructions say. Q. Okay. Is that a similar implementation for the other demos that you worked on? A. Most likely, yes. Q. Okay. So you click on I and then I would take you to a link that showed the full image? *Id.* at 129 – 131

Element	'004 Claim Recitation	Exemplary Citations in Reference
		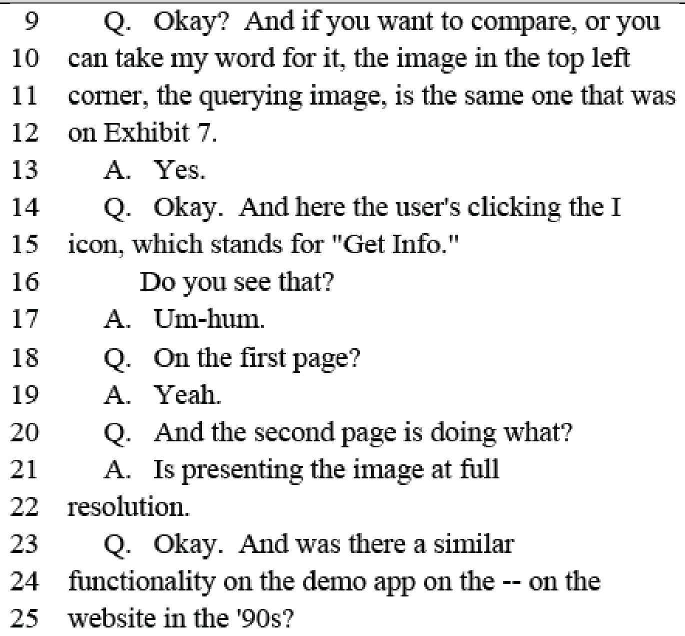 Q. Okay? And if you want to compare, or you can take my word for it, the image in the top left corner, the querying image, is the same one that was on Exhibit 7. A. Yes. Q. Okay. And here the user's clicking the I icon, which stands for "Get Info." Do you see that? A. Um-hum. Q. On the first page? A. Yeah. Q. And the second page is doing what? A. Is presenting the image at full resolution. Q. Okay. And was there a similar functionality on the demo app on the -- on the website in the '90s?

			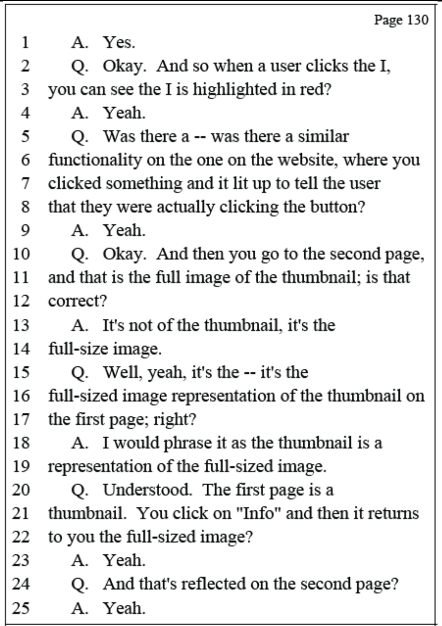Page 130 A. Yes. Q. Okay. And so when a user clicks the I, you can see the I is highlighted in red? A. Yeah. Q. Was there a -- was there a similar functionality on the one on the website, where you clicked something and it lit up to tell the user that they were actually clicking the button? A. Yeah. Q. Okay. And then you go to the second page, and that is the full image of the thumbnail; is that correct? A. It's not of the thumbnail, it's the full-size image. Q. Well, yeah, it's the -- it's the full-sized image representation of the thumbnail on the first page; right? A. I would phrase it as the thumbnail is a representation of the full-sized image. Q. Understood. The first page is a thumbnail. You click on "Info" and then it returns to you the full-sized image? A. Yeah. Q. And that's reflected on the second page? A. Yeah. 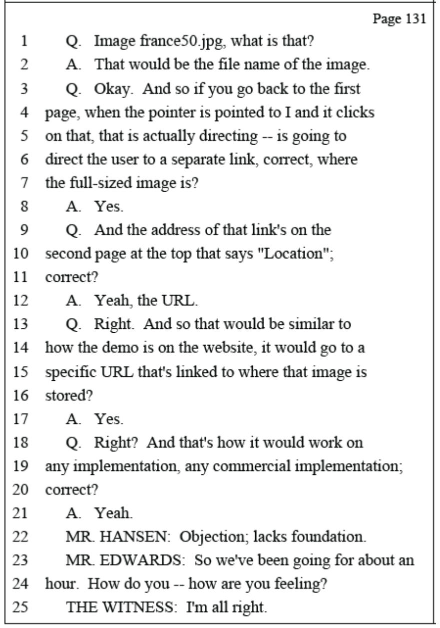Page 131 Q. Image france50.jpg, what is that? A. That would be the file name of the image. Q. Okay. And so if you go back to the first page, when the pointer is pointed to I and it clicks on that, that is actually directing -- is going to direct the user to a separate link, correct, where the full-sized image is? A. Yes. Q. And the address of that link's on the second page at the top that says "Location"; correct? A. Yeah, the URL. Q. Right. And so that would be similar to how the demo is on the website, it would go to a specific URL that's linked to where that image is stored? A. Yes. Q. Right? And that's how it would work on any implementation, any commercial implementation; correct? A. Yeah. MR. HANSEN: Objection; lacks foundation. MR. EDWARDS: So we've been going for about an hour. How do you -- how are you feeling? THE WITNESS: I'm all right.

Element	'004 Claim Recitation	Exemplary Citations in Reference
		Id. at 160–161 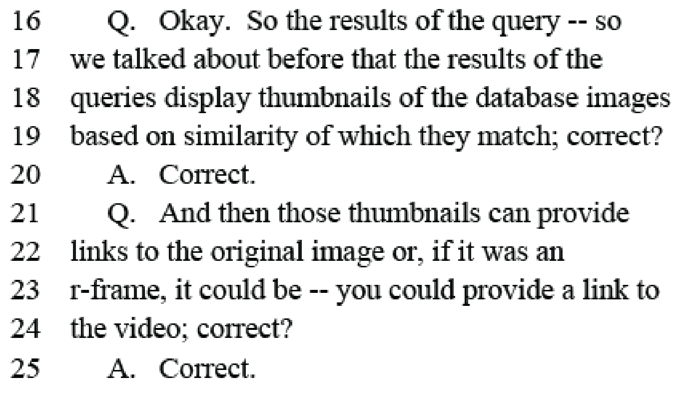 Q. Okay. So the results of the query -- so we talked about before that the results of the queries display thumbnails of the database images based on similarity of which they match; correct? A. Correct. Q. And then those thumbnails can provide links to the original image or, if it was an r-frame, it could be -- you could provide a link to the video; correct? A. Correct. 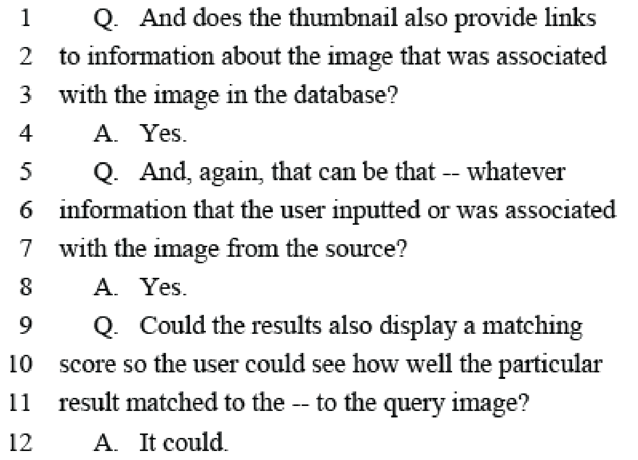 Q. And does the thumbnail also provide links to information about the image that was associated with the image in the database? A. Yes. Q. And, again, that can be that -- whatever information that the user inputted or was associated with the image from the source? A. Yes. Q. Could the results also display a matching score so the user could see how well the particular result matched to the -- to the query image? A. It could. *Id.* at 321

Element	’004 Claim Recitation	Exemplary Citations in Reference
		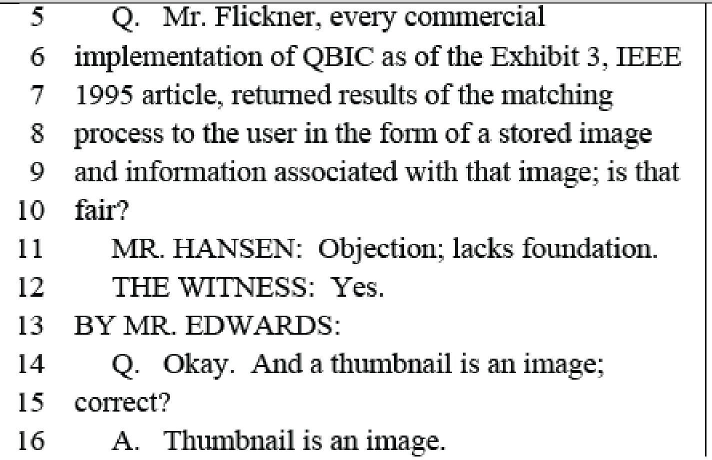 Q. Mr. Flickner, every commercial implementation of QBIC as of the Exhibit 3, IEEE 1995 article, returned results of the matching process to the user in the form of a stored image and information associated with that image; is that fair? MR. HANSEN: Objection; lacks foundation. THE WITNESS: Yes. BY MR. EDWARDS: Q. Okay. And a thumbnail is an image; correct? A. Thumbnail is an image. Flickner Depo at 25-26, 41-42, 50, 162-163, 201-202, 209-210 To the extent it is found that QBIC does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that QBIC does not anticipate this claim, QBIC renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA’s Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine QBIC with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

2. Claim 6

Claim	'004 Claim Recitation	Exemplary Citations in Reference
6	The method of claim 1, wherein the at least one object is identified as the target object based at least in part on at least one user input.	QBIC discloses the limitation. The QBIC Project: Querying Images by Content Using Color, Texture, and Shape (IBM 0002390–404) at 2391-92 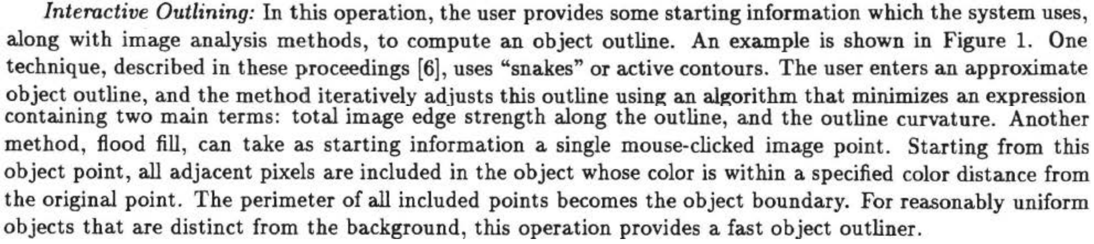 *Interactive Outlining:* In this operation, the user provides some starting information which the system uses, along with image analysis methods, to compute an object outline. An example is shown in Figure 1. One technique, described in these proceedings [6], uses "snakes" or active contours. The user enters an approximate object outline, and the method iteratively adjusts this outline using an algorithm that minimizes an expression containing two main terms: total image edge strength along the outline, and the outline curvature. Another method, flood fill, can take as starting information a single mouse-clicked image point. Starting from this object point, all adjacent pixels are included in the object whose color is within a specified color distance from the original point. The perimeter of all included points becomes the object boundary. For reasonably uniform objects that are distinct from the background, this operation provides a fast object outliner. The QBIC Project: Querying Images by Content Using Color, Texture, and Shape (IBM 0002390–404) at 2396 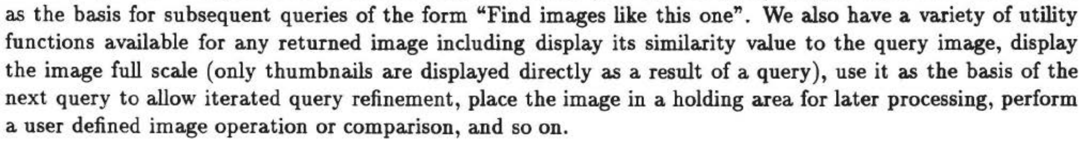 All the above measures can be used individually or in any weighted combination (as selected by the user), allowing, for example, queries based on color only, on color and shape, and so on. When a query is run, results are displayed in order, from best match to *n*th best match (*n* is user-settable). Any returned hit can be used as the basis for subsequent queries of the form "Find images like this one". We also have a variety of utility functions available for any returned image including display its similarity value to the query image, display the image full scale (only thumbnails are displayed directly as a result of a query), use it as the basis of the next query to allow iterated query refinement, place the image in a holding area for later processing, perform a user defined image operation or comparison, and so on. Query by Image and Video Content: The QBIC System (IBM 0000893–902): "To achieve this functionality, QBIC has two main components: database population (the process of creating an image database) and database query. During the population, images and videos are processed to extract features describing their content-colors, textures, shapes, and camera and object motion-and the features are stored in a database. During the query, the user composes a query graphically. Features are generated from the graphical query and then input to a matching engine that finds images or videos from the database with similar features. Figure 2 shows the system architecture." *Id*. at p. 3.

		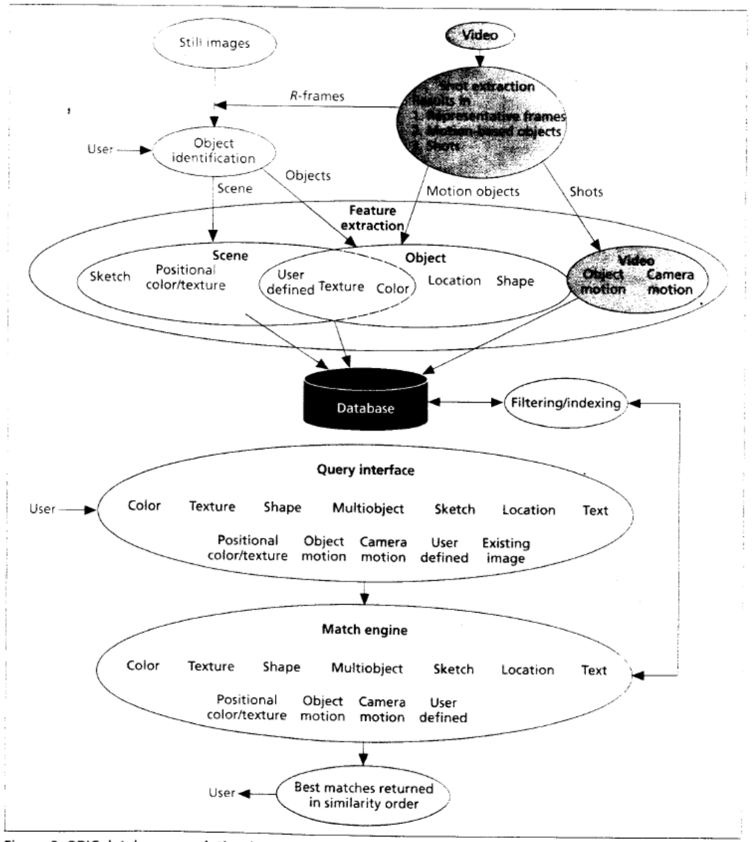 **Figure 2. QBIC database population (top) and query (bottom) architecture.** *Id.* at Fig. 2. "USER INTERFACE. The user interface must be designed to let users easily select content-based properties, allow these properties to be combined with each other and with text or parametric data, and let users reformulate queries and generally navigate the database." *Id.* at p. 9. Image, Audio, and Video Extenders Administration and Programming (IBM 000001–617): "Your applications can even search for images by content. Imagine an application

		that uses visual examples to search for images. With such an application, users could select an example image and have the application find other images that have colors or textures similar to those in the example. With DB2 extenders' Query by Image Content (QBIC) capability, you can create applications that search for images in this visual way." *Id*. at p. 4. "To make an image available for searching by content, you catalog the image. When you catalog an image, the Image Extender analyzes the image, by computing the feature values for the image, and stores the values in a QBIC catalog. When you search for an image by content, your query identifies one or more features for the search (such as average color), a source for each feature (such as an example image), and a target set of cataloged images. The Image Extender computes the feature value of the source and compares it to the cataloged feature values for the target images. It then computes a score that indicates how similar the feature values of the target images are to the source. You can have the Image Extender return the images whose features are most similar to the source. The Image Extender will return the handle of each image and the image score. You can also have the Image Extender return only the score of a single image." *Id*. at pp. 20–21.

		QBIC Demo (IBM000747) 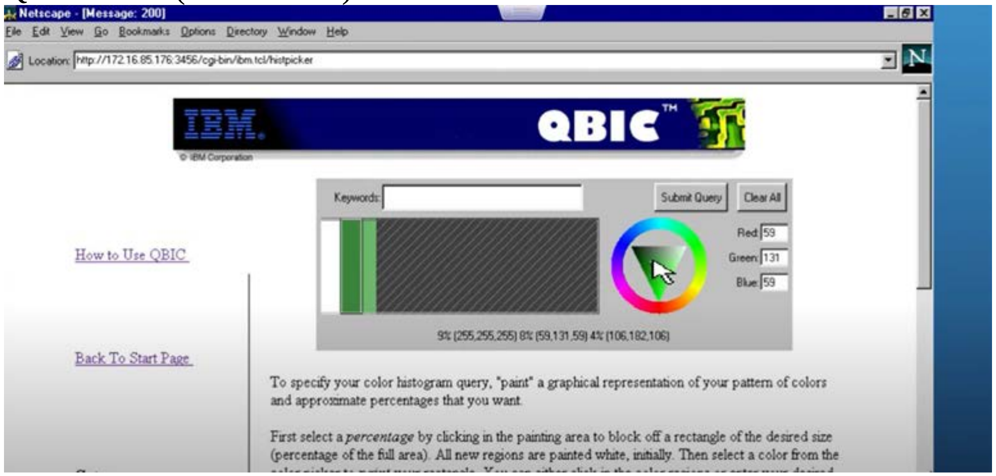 Location: http://172.16.85.176:3456/cgi-bin/ibm.tcl/histpicker Keywords: Red: 59 Green: 131 Blue: 59 9% (255,255,255) 8% (59,131,59) 4% (106,182,106) To specify your color histogram query, "paint" a graphical representation of your pattern of colors and approximate percentages that you want. First select a *percentage* by clicking in the painting area to block off a rectangle of the desired size (percentage of the full area). All new regions are painted white, initially. Then select a color from the QBIC Demo (IBM000747) 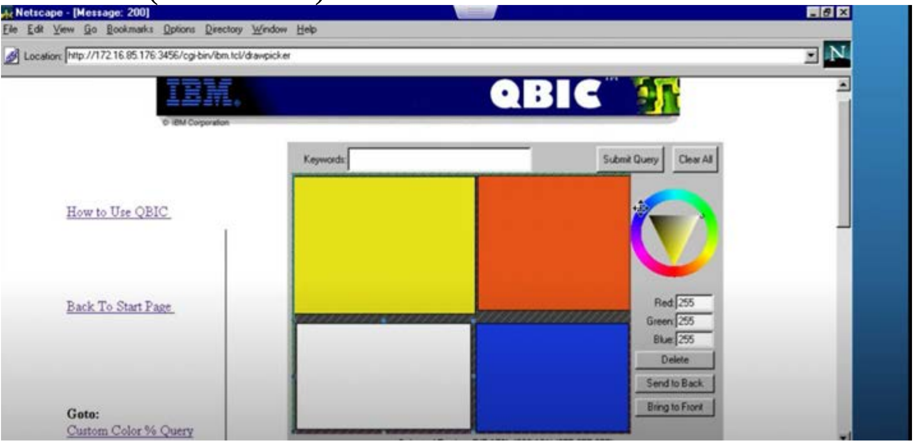 Location: http://172.16.85.176:3456/cgi-bin/ibm.tcl/drawpicker Keywords: Red: 255 Green: 255 Blue: 255

		QBIC Demo (IBM000747) 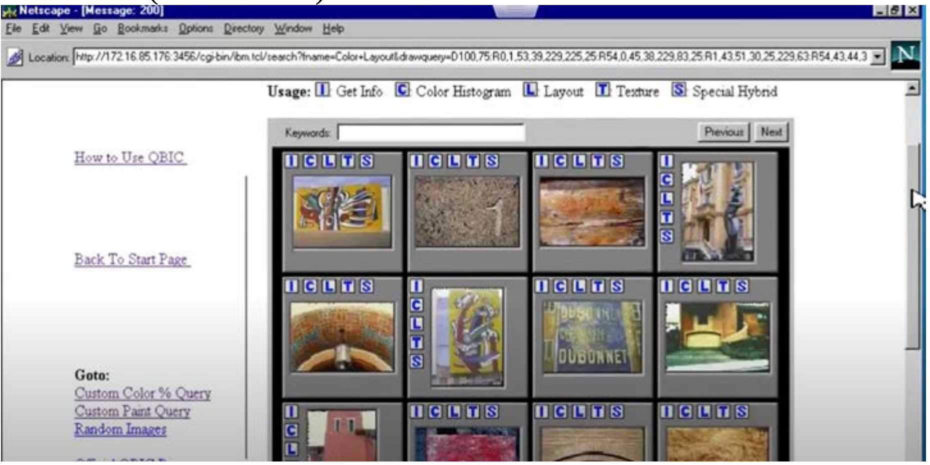 Location: http://172.16.85.176:3456/cgi-bin/ibm.tcl/search?fname=Color+Layout&drawquery=D100,75:R0,1,53,39,229,225,25:R54,0,45,38,229,83,25:R1,43,51,30,25,229,63:R54,43,44,3 Flickner Depo at 49-50 Q. Okay. And what about the L button, it says "Layout," right? A. Yeah. Q. What does that mean? A. You did a crude sketch. Q. Sorry. What was that? A. You would do a crude sketch. Q. A crude sketch. What is a -- what do you mean by a crude sketch? A. You would -- it would pop -- pop up a drawing box, and you could draw lines and -- and input different colors, if I recall correctly.

		Flickner Depo at 138 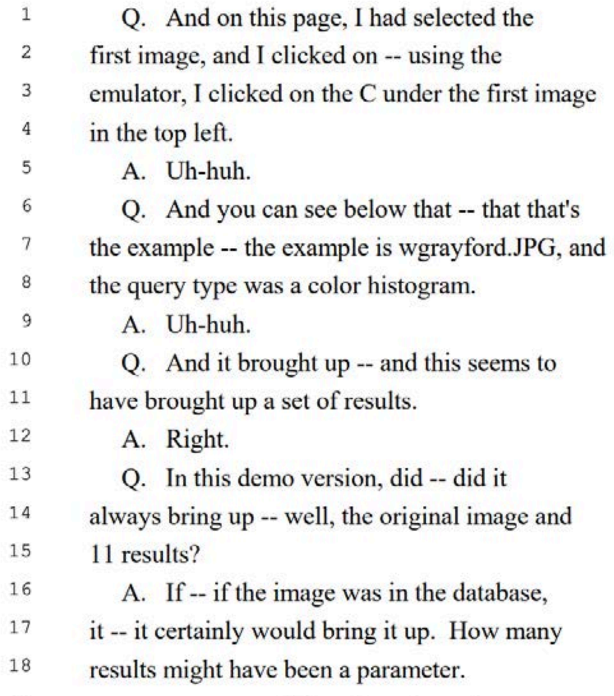 Q. And on this page, I had selected the first image, and I clicked on -- using the emulator, I clicked on the C under the first image in the top left. A. Uh-huh. Q. And you can see below that -- that that's the example -- the example is wgrayford.JPG, and the query type was a color histogram. A. Uh-huh. Q. And it brought up -- and this seems to have brought up a set of results. A. Right. Q. In this demo version, did -- did it always bring up -- well, the original image and 11 results? A. If -- if the image was in the database, it -- it certainly would bring it up. How many results might have been a parameter. Flickner 8/29/23 Depo at 100 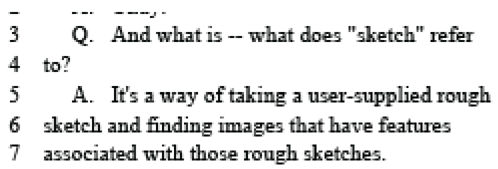 Q. And what is -- what does "sketch" refer to? A. It's a way of taking a user-supplied rough sketch and finding images that have features associated with those rough sketches. *Id.* at 119–124

			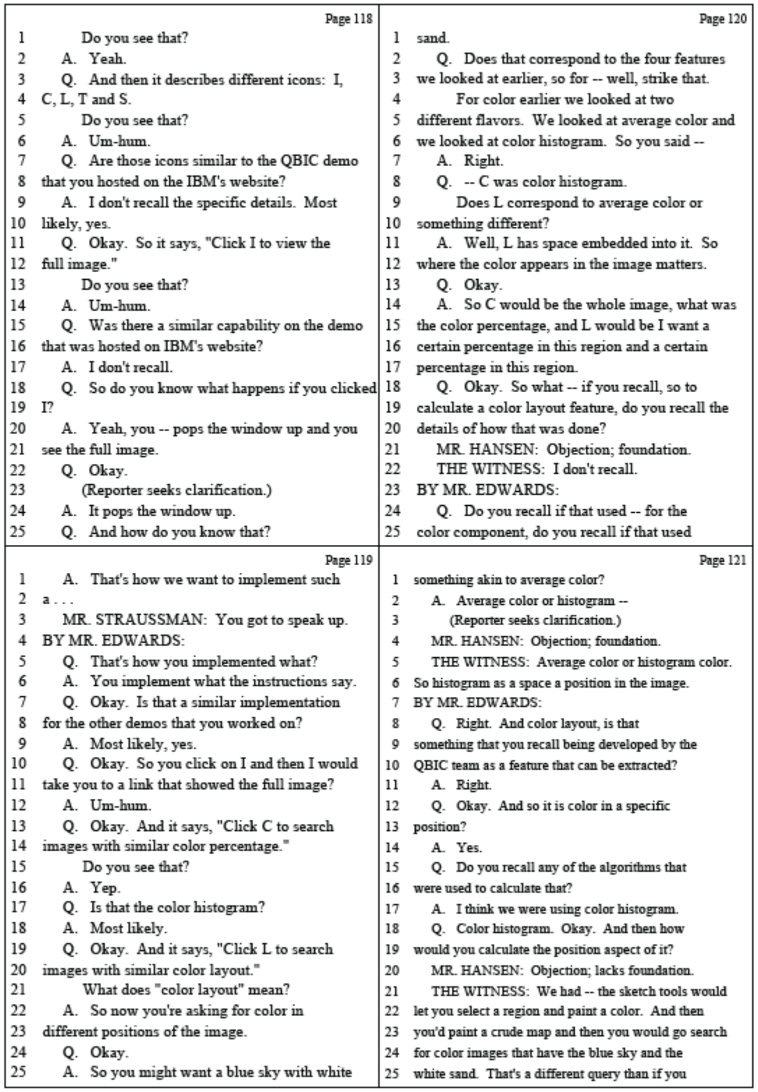

Page 118

Do you see that?
A. Yeah.
Q. And then it describes different icons: I, C, L, T and S.
Do you see that?
A. Um-hum.
Q. Are those icons similar to the QBIC demo that you hosted on the IBM's website?
A. I don't recall the specific details. Most likely, yes.
Q. Okay. So it says, "Click I to view the full image."
Do you see that?
A. Um-hum.
Q. Was there a similar capability on the demo that was hosted on IBM's website?
A. I don't recall.
Q. So do you know what happens if you clicked I?
A. Yeah, you -- pops the window up and you see the full image.
Q. Okay.
(Reporter seeks clarification.)
A. It pops the window up.
Q. And how do you know that?

Page 119

A. That's how we want to implement such a . . .
MR. STRAUSSMAN: You got to speak up.
BY MR. EDWARDS:
Q. That's how you implemented what?
A. You implement what the instructions say.
Q. Okay. Is that a similar implementation for the other demos that you worked on?
A. Most likely, yes.
Q. Okay. So you click on I and then I would take you to a link that showed the full image?
A. Um-hum.
Q. Okay. And it says, "Click C to search images with similar color percentage."
Do you see that?
A. Yep.
Q. Is that the color histogram?
A. Most likely.
Q. Okay. And it says, "Click L to search images with similar color layout."
What does "color layout" mean?
A. So now you're asking for color in different positions of the image.
Q. Okay.
A. So you might want a blue sky with white

Page 120

sand.
Q. Does that correspond to the four features we looked at earlier, so for -- well, strike that.
For color earlier we looked at two different flavors. We looked at average color and we looked at color histogram. So you said --
A. Right.
Q. -- C was color histogram.
Does L correspond to average color or something different?
A. Well, L has space embedded into it. So where the color appears in the image matters.
Q. Okay.
A. So C would be the whole image, what was the color percentage, and L would be I want a certain percentage in this region and a certain percentage in this region.
Q. Okay. So what -- if you recall, so to calculate a color layout feature, do you recall the details of how that was done?
MR. HANSEN: Objection; foundation.
THE WITNESS: I don't recall.
BY MR. EDWARDS:
Q. Do you recall if that used -- for the color component, do you recall if that used

Page 121

something akin to average color?
A. Average color or histogram --
(Reporter seeks clarification.)
MR. HANSEN: Objection; foundation.
THE WITNESS: Average color or histogram color. So histogram as a space a position in the image.
BY MR. EDWARDS:
Q. Right. And color layout, is that something that you recall being developed by the QBIC team as a feature that can be extracted?
A. Right.
Q. Okay. And so it is color in a specific position?
A. Yes.
Q. Do you recall any of the algorithms that were used to calculate that?
A. I think we were using color histogram.
Q. Color histogram. Okay. And then how would you calculate the position aspect of it?
MR. HANSEN: Objection; lacks foundation.
THE WITNESS: We had -- the sketch tools would let you select a region and paint a color. And then you'd paint a crude map and then you would go search for color images that have the blue sky and the white sand. That's a different query than if you

		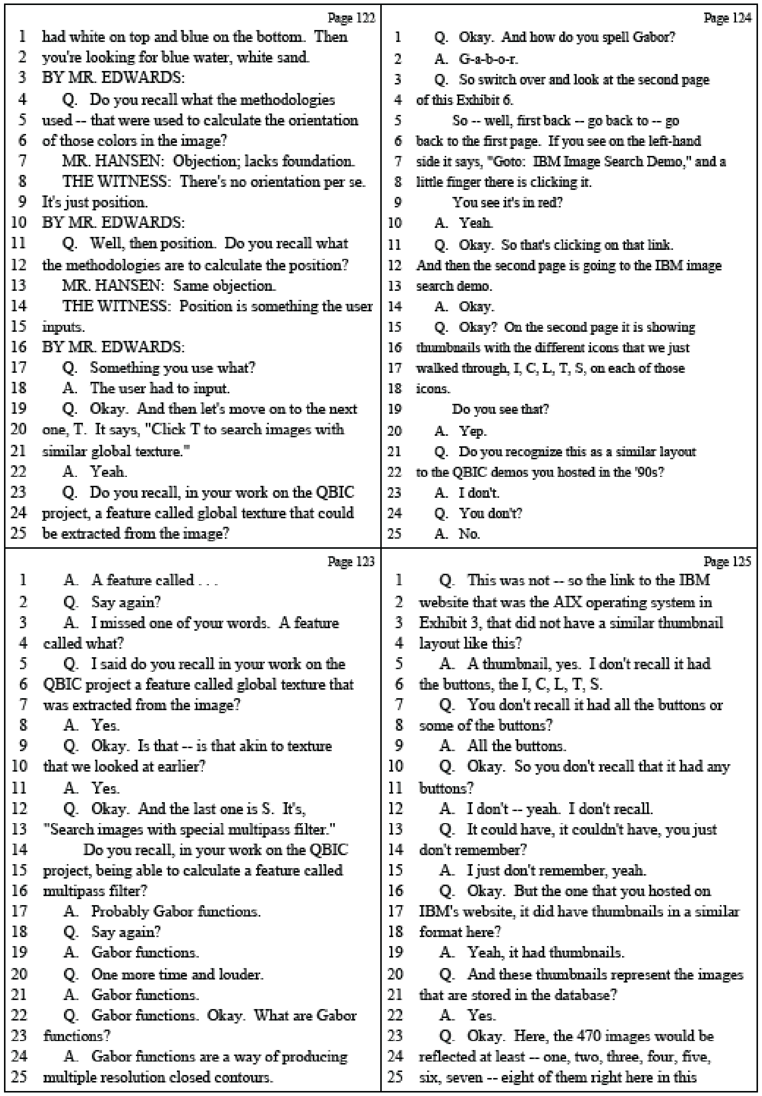 Page 122 had white on top and blue on the bottom. Then you're looking for blue water, white sand. BY MR. EDWARDS: Q. Do you recall what the methodologies used -- that were used to calculate the orientation of those colors in the image? MR. HANSEN: Objection; lacks foundation. THE WITNESS: There's no orientation per se. It's just position. BY MR. EDWARDS: Q. Well, then position. Do you recall what the methodologies are to calculate the position? MR. HANSEN: Same objection. THE WITNESS: Position is something the user inputs. BY MR. EDWARDS: Q. Something you use what? A. The user had to input. Q. Okay. And then let's move on to the next one, T. It says, "Click T to search images with similar global texture." A. Yeah. Q. Do you recall, in your work on the QBIC project, a feature called global texture that could be extracted from the image? Page 123 A. A feature called . . . Q. Say again? A. I missed one of your words. A feature called what? Q. I said do you recall in your work on the QBIC project a feature called global texture that was extracted from the image? A. Yes. Q. Okay. Is that -- is that akin to texture that we looked at earlier? A. Yes. Q. Okay. And the last one is S. It's, "Search images with special multipass filter." Do you recall, in your work on the QBIC project, being able to calculate a feature called multipass filter? A. Probably Gabor functions. Q. Say again? A. Gabor functions. Q. One more time and louder. A. Gabor functions. Q. Gabor functions. Okay. What are Gabor functions? A. Gabor functions are a way of producing multiple resolution closed contours. Page 124 Q. Okay. And how do you spell Gabor? A. G-a-b-o-r. Q. So switch over and look at the second page of this Exhibit 6. So -- well, first back -- go back to -- go back to the first page. If you see on the left-hand side it says, "Goto: IBM Image Search Demo," and a little finger there is clicking it. You see it's in red? A. Yeah. Q. Okay. So that's clicking on that link. And then the second page is going to the IBM image search demo. A. Okay. Q. Okay? On the second page it is showing thumbnails with the different icons that we just walked through, I, C, L, T, S, on each of those icons. Do you see that? A. Yep. Q. Do you recognize this as a similar layout to the QBIC demos you hosted in the '90s? A. I don't. Q. You don't? A. No. Page 125 Q. This was not -- so the link to the IBM website that was the AIX operating system in Exhibit 3, that did not have a similar thumbnail layout like this? A. A thumbnail, yes. I don't recall it had the buttons, the I, C, L, T, S. Q. You don't recall it had all the buttons or some of the buttons? A. All the buttons. Q. Okay. So you don't recall that it had any buttons? A. I don't -- yeah. I don't recall. Q. It could have, it couldn't have, you just don't remember? A. I just don't remember, yeah. Q. Okay. But the one that you hosted on IBM's website, it did have thumbnails in a similar format here? A. Yeah, it had thumbnails. Q. And these thumbnails represent the images that are stored in the database? A. Yes. Q. Okay. Here, the 470 images would be reflected at least -- one, two, three, four, five, six, seven -- eight of them right here in this *Id.* at 159–160

		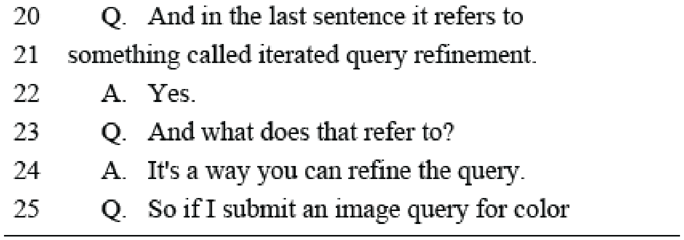Q. And in the last sentence it refers to something called iterated query refinement. A. Yes. Q. And what does that refer to? A. It's a way you can refine the query. Q. So if I submit an image query for color 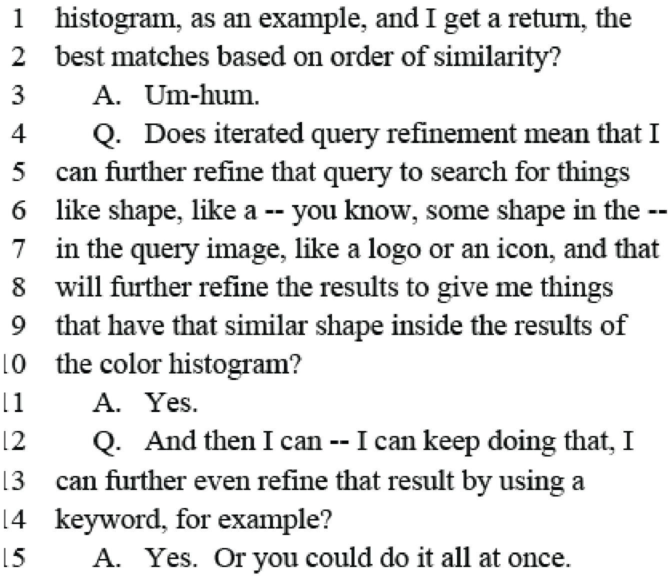histogram, as an example, and I get a return, the best matches based on order of similarity? A. Um-hum. Q. Does iterated query refinement mean that I can further refine that query to search for things like shape, like a -- you know, some shape in the -- in the query image, like a logo or an icon, and that will further refine the results to give me things that have that similar shape inside the results of the color histogram? A. Yes. Q. And then I can -- I can keep doing that, I can further even refine that result by using a keyword, for example? A. Yes. Or you could do it all at once. Flickner 8/29/23 Depo at Ex. 6 at 1

		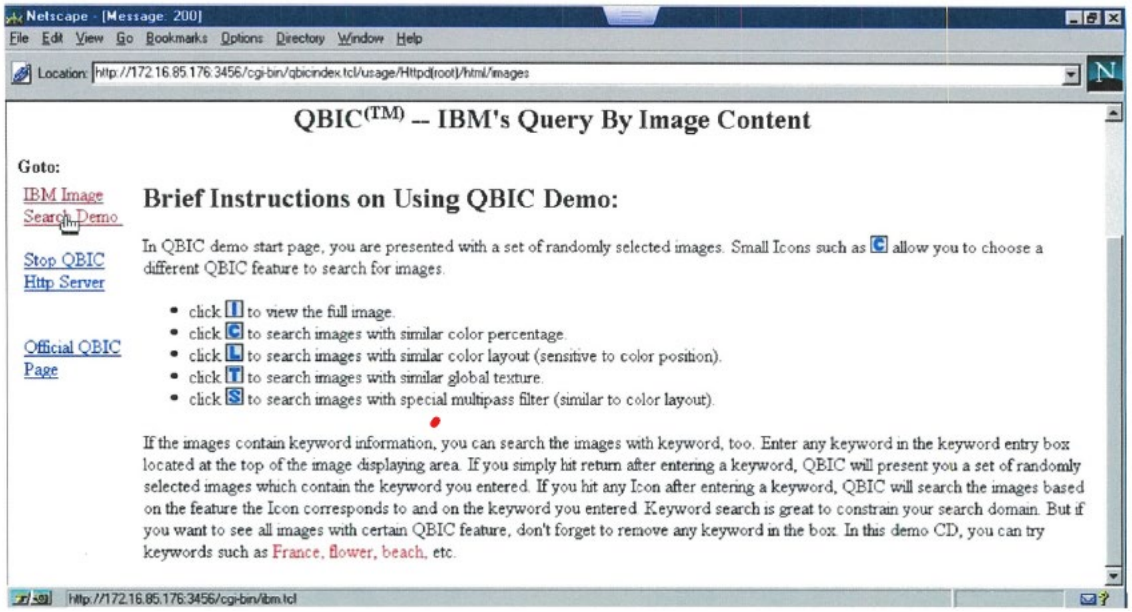 QBIC$^{(TM)}$ -- IBM's Query By Image Content **Goto:** IBM Image Search Demo Stop QBIC Http Server Official QBIC Page **Brief Instructions on Using QBIC Demo:** In QBIC demo start page, you are presented with a set of randomly selected images. Small Icons such as C allow you to choose a different QBIC feature to search for images. • click I to view the full image. • click C to search images with similar color percentage. • click L to search images with similar color layout (sensitive to color position). • click T to search images with similar global texture. • click S to search images with special multipass filter (similar to color layout). If the images contain keyword information, you can search the images with keyword, too. Enter any keyword in the keyword entry box located at the top of the image displaying area. If you simply hit return after entering a keyword, QBIC will present you a set of randomly selected images which contain the keyword you entered. If you hit any Icon after entering a keyword, QBIC will search the images based on the feature the Icon corresponds to and on the keyword you entered. Keyword search is great to constrain your search domain. But if you want to see all images with certain QBIC feature, don't forget to remove any keyword in the box. In this demo CD, you can try keywords such as France, flower, beach, etc. *Id.* at 2.

<table>
<tr>
<td></td>
<td></td>
<td>
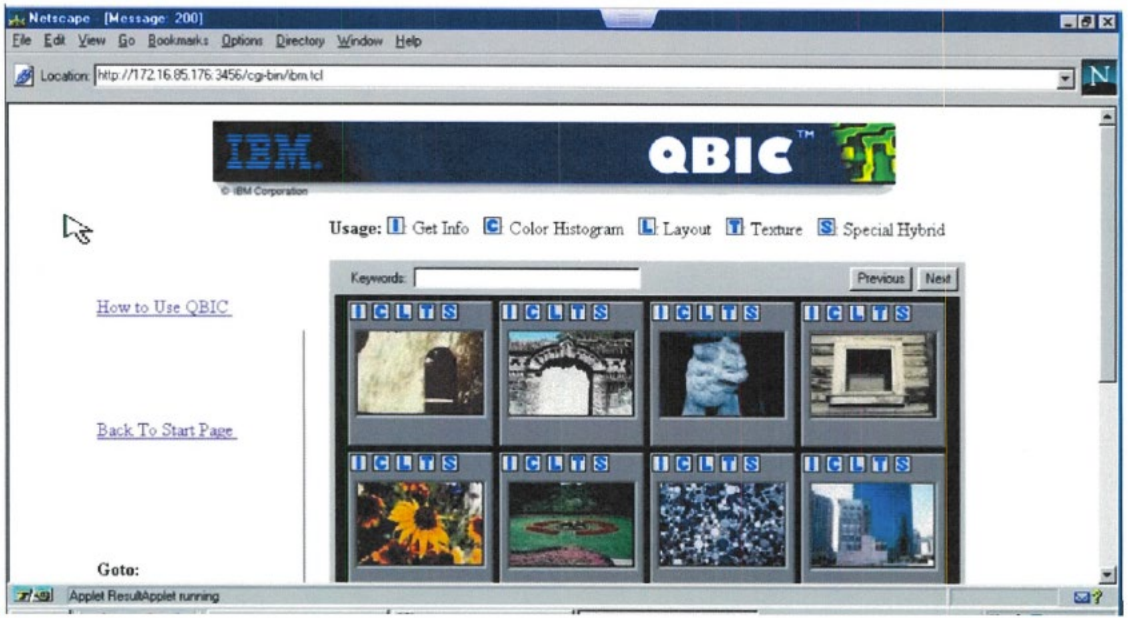

Flickner 8/29/23 Depo Ex. 7 at 1
</td>
</tr>
</table>

		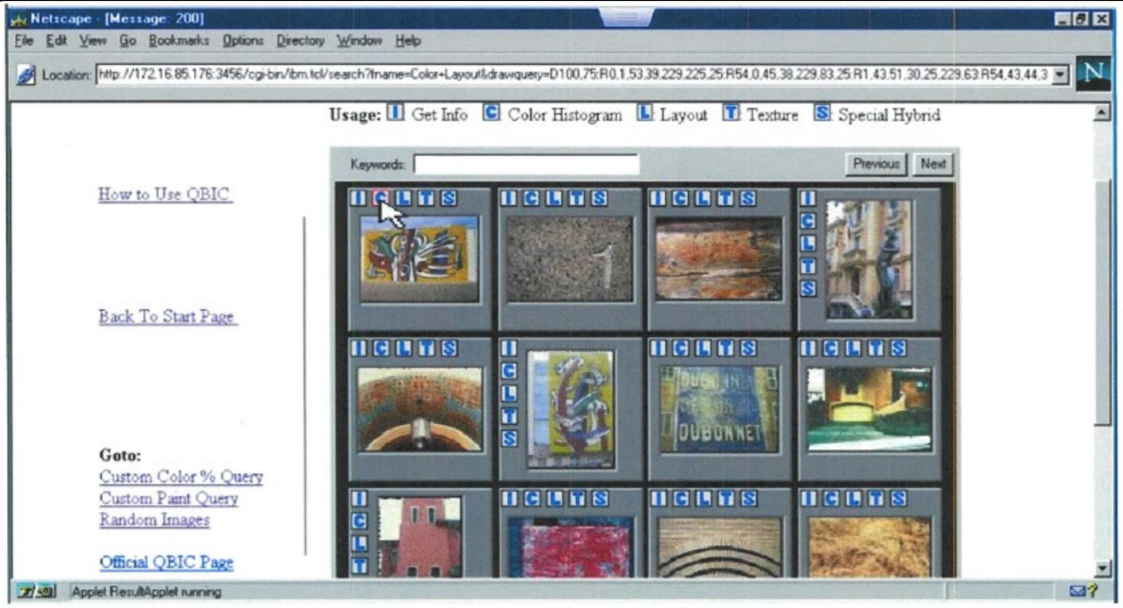 Location: http://172.16.85.176:3456/cgi-bin/ibm.tcl/search?fname=Color+Layout&drawquery=D100,75:R0,1,53,39,229,225,25:R54,0,45,38,229,83,25:R1,43,51,30,25,229,63:R54,43,44,3 *Id.* at 2

		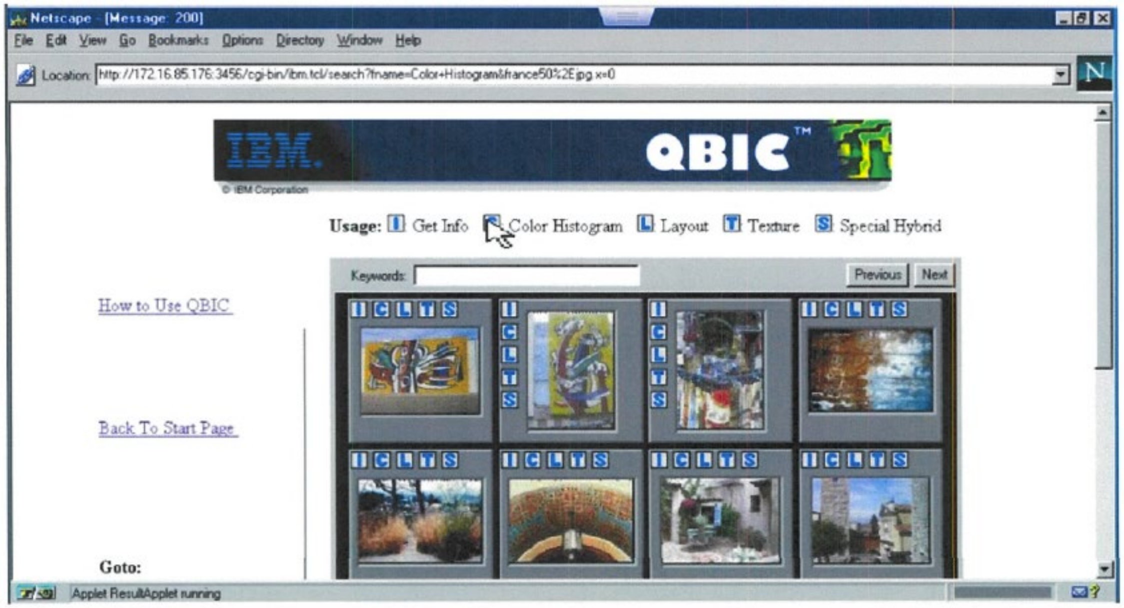 Location: http://172.16.85.176:3456/cgi-bin/ibm.tcl/search?fname=Color+Histogram&france50%2Ejpg.x=0 Flickner 8/29/23 Depo at Ex. 8

		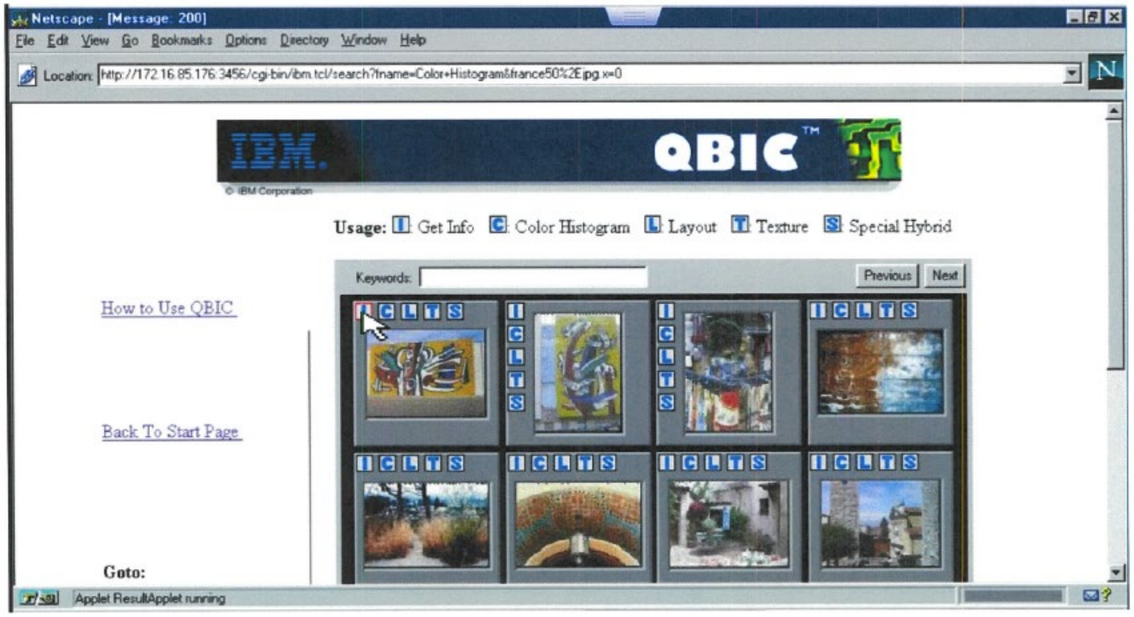 http://172.16.85.176:3456/cgi-bin/ibm.tcl/search?fname=Color+Histogram&france50%2Ejpg.x=0 Applet ResultApplet running *Id.* at 2

		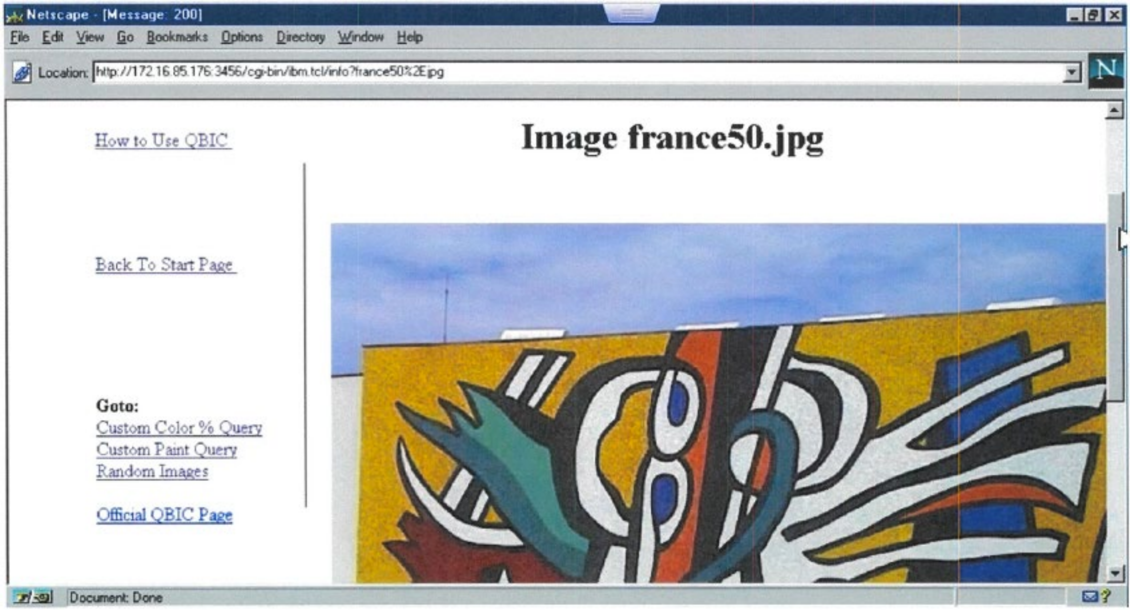 Netscape - [Message: 200] File Edit View Go Bookmarks Options Directory Window Help Location: http://172.16.85.176:3456/cgi-bin/ibm.tcl/info?france50%2Ejpg How to Use QBIC Back To Start Page Goto: Custom Color % Query Custom Paint Query Random Images Official QBIC Page Image france50.jpg Document: Done *Id.* at 3

		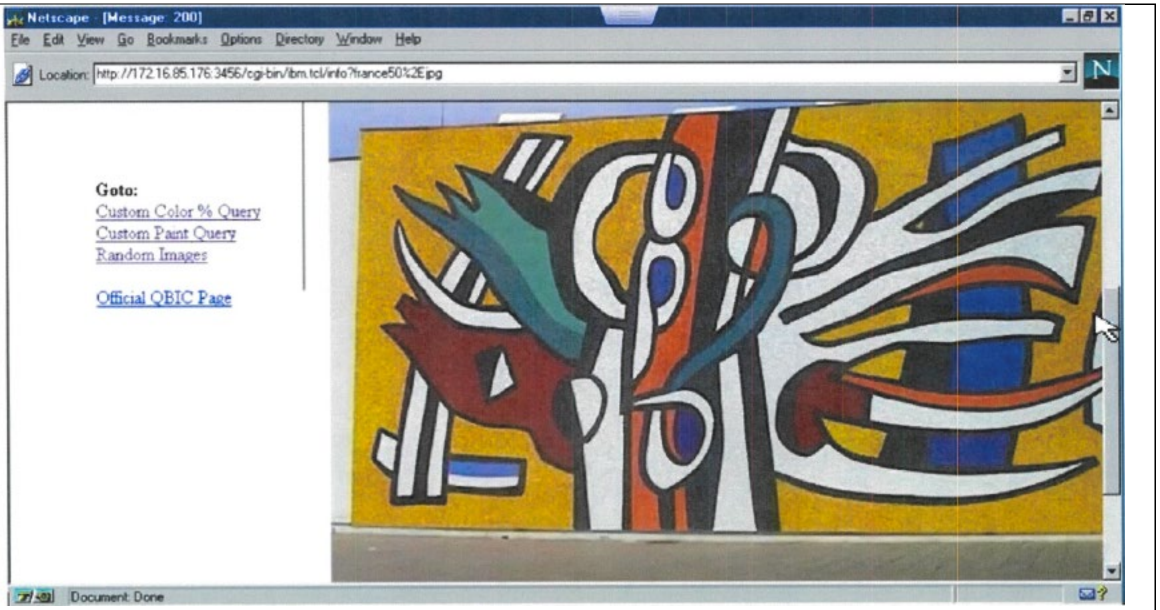 Location: http://172.16.85.176:3456/cgi-bin/ibm.tcl/info?france50%2Ejpg *See also* Flickner Depo at 62-63, 131-132 To the extent it is found that QBIC does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that QBIC does not anticipate this claim, QBIC renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine QBIC with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

3. Claim 18

Claim	'004 Claim Recitation	Exemplary Citations in Reference
18	The method of claim 1, further comprising communicating at least one of banking account information, user data and terminal specific data to a banking company.	QBIC discloses this limitation. Query by Image and Video Content: The QBIC System (IBM 0000893–902): "INTEGRATION WITH TEXT AND PARAMETRIC ANNOTATION. Query by visual content complements and extends existing query methods. Systems must be able to integrate queries combining date, subject matter, price, and availability with content properties such as color, texture, and shape." *Id.* at p. 8. IBM0002390 at 2391-92 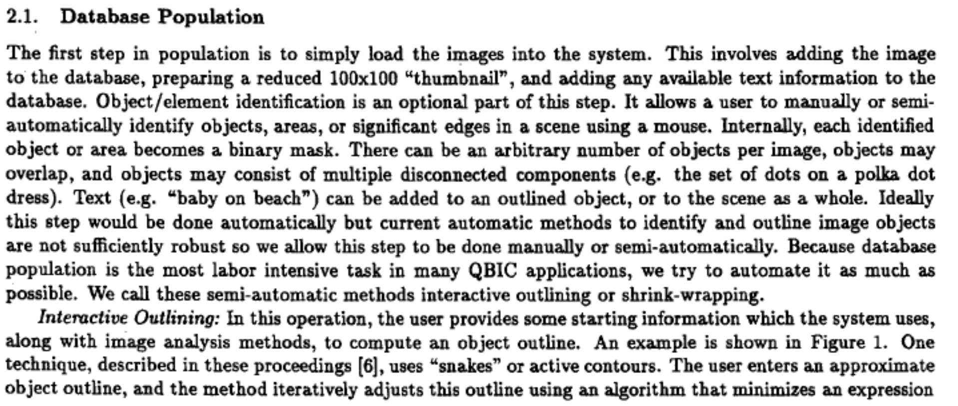 **2.1. Database Population** The first step in population is to simply load the images into the system. This involves adding the image to the database, preparing a reduced 100x100 "thumbnail", and adding any available text information to the database. Object/element identification is an optional part of this step. It allows a user to manually or semi-automatically identify objects, areas, or significant edges in a scene using a mouse. Internally, each identified object or area becomes a binary mask. There can be an arbitrary number of objects per image, objects may overlap, and objects may consist of multiple disconnected components (e.g. the set of dots on a polka dot dress). Text (e.g. "baby on beach") can be added to an outlined object, or to the scene as a whole. Ideally this step would be done automatically but current automatic methods to identify and outline image objects are not sufficiently robust so we allow this step to be done manually or semi-automatically. Because database population is the most labor intensive task in many QBIC applications, we try to automate it as much as possible. We call these semi-automatic methods interactive outlining or shrink-wrapping. *Interactive Outlining:* In this operation, the user provides some starting information which the system uses, along with image analysis methods, to compute an object outline. An example is shown in Figure 1. One technique, described in these proceedings [6], uses "snakes" or active contours. The user enters an approximate object outline, and the method iteratively adjusts this outline using an algorithm that minimizes an expression

		containing two main terms: total image edge strength along the outline, and the outline curvature. Another method, flood fill, can take as starting information a single mouse-clicked image point. Starting from this object point, all adjacent pixels are included in the object whose color is within a specified color distance from the original point. The perimeter of all included points becomes the object boundary. For reasonably uniform objects that are distinct from the background, this operation provides a fast object outliner. In interactive outlining, as in other parts of our system, we have attempted to build the functions in such a way as to carefully distribute the tasks between the human and machine. For example, the tasks of giving the starting outline for an object or of attaching semantic meaning to objects is left to the user since they are difficult for machines. On the other hand, computation of quantitative image/object features is ideally suited for machines and so is done automatically.  Figure 1: Original image, approximate user outline, and automatically refined outline of sun. Image, Audio, and Video Extenders Administration and Programming (IBM 000001–617): "DB2 (DB2) Universal Database (UDB) is a powerful, object-relational database manager. It stores and protects traditional numeric and character data, as well as large, complex objects (LOBs). DB2 extenders help you exploit DB2 's object-relational features. The extenders define distinct data types and special functions for image, audio, video, and text objects. … Because of that, the extenders give your applications a single point of access to any or all of these types of data, along with traditional numeric and character data." *Id.* at p. 3. "A trigger defines a set of actions that are activated by a change to a table. Triggers can be used to perform actions such as validating input data, automatically generating a value for a newly inserted row, reading from other

> tables for cross-referencing purposes, or writing to other tables for auditing purposes. Triggers are often used for integrity checking or to enforce business rules." *Id*. at 16–17.

Managing Enterprise Information Portal (IBM 000777–880):

> "Many paper-intensive enterprises, such as insurance companies and financial institutions, administer large volumes of business-related content. The need for an enterprise solution for managing and accessing business information spans many industries. A content server is a software system that stores multimedia, business forms, documents, and related data, along with metadata that allows employees to process and work with the content." *Id*. at p. 3.
>
> "The image search server uses IBM's QBIC (query by image content) technology to help you search for objects by certain visual properties, such as color and texture. The image search server analyzes images and stores the image information in a database. Then users can run image queries, which use the visual properties of images, to match colors, textures, and their positions without describing them in words. You can combine content-based queries with text and keyword searches for more powerful retrieval of image and multimedia data." *Id*. at p. 63.

Flickner Depo at 131-132.

Flickner 8/29/23 Depo at 197–198

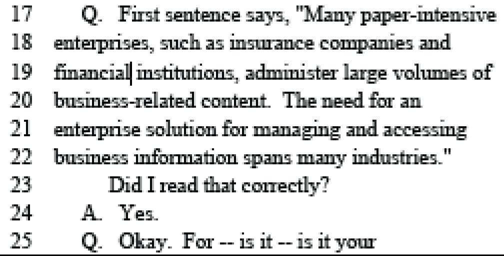

Q. First sentence says, "Many paper-intensive enterprises, such as insurance companies and financial institutions, administer large volumes of business-related content. The need for an enterprise solution for managing and accessing business information spans many industries."
Did I read that correctly?
A. Yes.
Q. Okay. For -- is it -- is it your

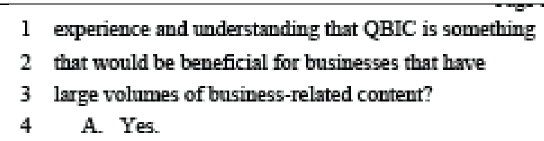

experience and understanding that QBIC is something
that would be beneficial for businesses that have
large volumes of business-related content?
A. Yes.

Id. at 307

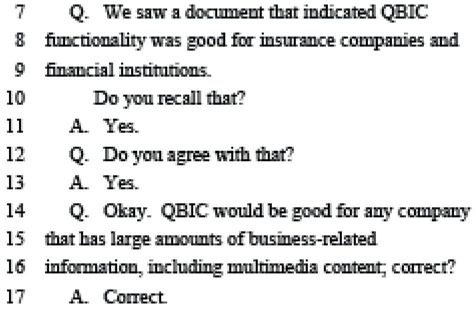

Q. We saw a document that indicated QBIC
functionality was good for insurance companies and
financial institutions.
Do you recall that?
A. Yes.
Q. Do you agree with that?
A. Yes.
Q. Okay. QBIC would be good for any company
that has large amounts of business-related
information, including multimedia content; correct?
A. Correct.

To the extent it is found that QBIC does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that QBIC does not anticipate this claim, QBIC renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine QBIC with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results. For example, a person of ordinary skill in the art would have been motivated to improve, modify, or combine QBIC with prior art reference teachings to arrive at this limitation at least because QBIC was designed for applications that included large amounts of business-related information, including multi-media content, including for financial

		institutions. *See, e.g.*, IBM 000789; Flickner Depo Tr. at 191:16-192:8, 293:14-294:8.

Nantworks, LLC v. Bank of America Corporation
Central District of California, No. 2:20-cv-7872
Invalidity Chart for U.S. Patent 8,478,036 based on
U.S. Patent No. 6,947,571 to Rhoads ("Rhoads")

Exhibit C-3

Bank of America Corporation ("BofA" or "Defendant") provides this chart subject to all reservations, objections, statements, and disclaimers set forth herein and in BofA's Final Invalidity Contentions Cover Pleading, as well as any amendment, supplement, or modification thereof, which are incorporated herein by reference in their entirety.

On information and belief, the application that matured into Rhoads was filed on May 15, 2000, claiming the priority date of May 19, 1999. Rhoads published at least as of September 20, 2005, and is available as prior art at least under 35 U.S.C. § 102(e).

As shown in the chart below, Rhoads anticipates asserted claim 1 of the '036 Patent. To the extent Rhoads is found not to expressly disclose certain limitations in the asserted claims, such limitations are inherent. To the extent Rhoads is found not to anticipate any asserted claims or claim elements of the '036 Patent, the reference nevertheless renders those claims or claim elements obvious under 35 U.S.C. § 103, either alone or in combination with other art identified herein or in the cover pleading of BofA's Final Invalidity Contentions, for the reasons identified in the cover pleading. These Final Invalidity Contentions are not an admission by BofA that the accused products, including any current or past versions of the accused products, are covered by, or infringe the asserted claims, but are based instead on the recognition that if the claims are interpreted to be broad enough to encompass the accused products, the claims must also be construed to have that same scope when considering whether they are invalid.

The following chart is based on the Court's claim construction rulings (Dkt. Nos. 145, 153, 236), and to the extent not contradictory, Plaintiffs' apparent interpretations of the claims in Plaintiff's "Preliminary" Final Infringement Contentions. BofA notes that in many instances, Plaintiff's "Preliminary" Final Infringement Contentions ignore and/or contradict the Court's claim construction rulings. BofA does not accept the assumptions concerning the scope and meaning implicit in Plaintiff's Preliminary Final Infringement Contentions, to the extent those assumptions are discernible, and reserves the right to challenge Plaintiff's proposed (or implied) constructions. Furthermore, Plaintiffs' "Preliminary" Final Contentions fail to comply with the level of specificity required by SPR 2.1.2 and 2.1.3, and include improper reservation of rights to supplement with "additional ways in which BoA Accused Products infringe." Any alleged deficient discovery is of Plaintiffs own making. Most if not all of the deficiencies in Plaintiffs' Preliminary Final Contentions have been documented by Defendants through multiple letters since at least May 18, 2021 and up to Bank of America's most recent letter of August 19, 2023. Plaintiffs' shifting infringement contentions, deficiencies in its infringement contentions, and changing claim scope have prejudiced BofA from being able to reasonably prepare its defenses, including its invalidity contentions. BofA also reserves the right to revise and supplement these charts if and when Plaintiff is permitted to provide revised or supplemental Infringement Contentions. Citations given in the chart below are merely representative of the respective elements to provide notice to Plaintiffs' of BofA's invalidity theories and are not meant to be exhaustive.

1. Claim 1

Element	'036 Claim Recitation	Exemplary Citations in Reference
[1Pre]	A content provisioning system comprising:	Rhoads discloses a content provisioning system. Rhoads at Abstract • "A cell phone is equipped with a 2D optical sensor, enabling a variety of applications. For example, such a phone may also be provided with a digital watermark decoder, permitting decoding of steganographic data on imaged objects. Movement of a phone may be inferred by sensing movement of an imaged pattern across the optical sensor's field of view, allowing use of the phone as a gestural input device through which a user can signal instructions to a computer-based process. A variety of other arrangements by which electronic devices can interact with the physical world are also detailed, e.g., involving sensing and responding to digital watermarks, bar codes, RFIDs, etc." Rhoads at 5:54-56. • "Basically, the technology detailed in this disclosure may be regarded as enhanced systems by which users can interact with computer-based devices." Rhoads at 9:35–41 • "The Steganographic decoding process may entail three Steps. In the first, the object is located. In the Second, the objects orientation is discerned. In the third, the Bedoop data is extracted from the image data corresponding to the Bedoop object." Rhoads at Fig. 2 (and accompanying description)

Element	'036 Claim Recitation	Exemplary Citations in Reference
		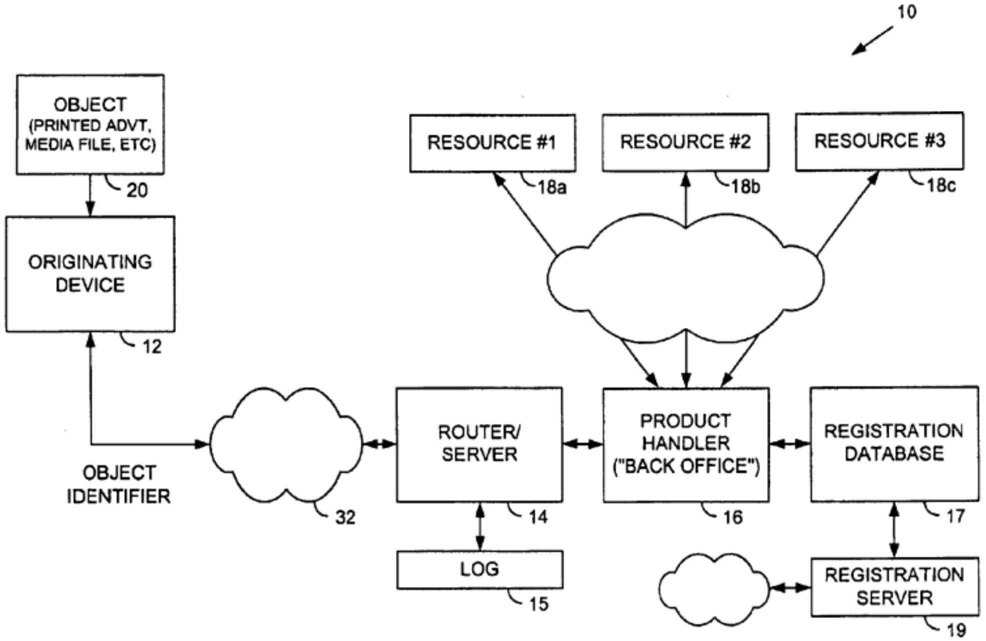 **FIG. 2** Rhoads at 49:15-18 - "Referring to FIG. 2, a system 10 according to the exemplary embodiment includes an originating device 12, a router/server 14, a product handler 16, a registration database 17, and one or more remote resources 18." Rhoads at 50:21-26 - "As more particularly detailed below, the handler 16 provides a response in accordance with the particular watermark payload. The response may be provided directly by the product handler to the device 12, or the handler may respond by communicating with a remote resource 18 (which may be, e.g., a data repository or service provider)."

Element	'036 Claim Recitation	Exemplary Citations in Reference
		To the extent it is found that Rhoads does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Rhoads does not anticipate this claim, Rhoads renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Rhoads with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1A]	a target database storing known targets of different types and recognition parameters associated with the known targets;	Rhoads discloses a target database storing known targets of different types and recognition parameters associated with the known targets. Rhoads at 9:35–10:21 • "The Steganographic decoding process may entail three Steps. In the first, the object is located. In the Second, the objects orientation is discerned. In the third, the Bedoop data is extracted from the image data corresponding to the Bedoop object. The first step, object location, can be assisted by various clues. One is the placement of the object; typically the center of the image field will be a point on the object. The surrounding data can then be analyzed to try and discern the object's boundaries. Another location technique is slight movement. Although the user will typically try to hold the object still, there will usually be some jitter of the Bedoop object within the image frame (e.g., a few pixels back and forth). Background visual clutter, in contrast, will typically be stationary. Such movement may thus be sensed and used to identify the Bedoop object from

Element	'036 Claim Recitation	Exemplary Citations in Reference
		within the image data. Still another object-location clue is object shape. Many Bedoop objects are rectangular in shape (or trapezoidal as viewed by the camera). Straight edge boundaries can thus be used to define an area of likely Bedoop data. Color is a further object identification clue that may be useful in some contexts. Yet another object location clue is spatial frequency. In imaging systems with well defined focal zones, undesired visual clutter may be at focal distances that results in blurring. The Bedoop object, in contrast, will be in focus and may be characterized by fine detail. Analyzing the image data for the high frequencies associated with fine detail can be used to distinguish the intended object from others. Characteristic markings on the object (as discussed below in connection with determining object orientation) can also be sensed and used in locating the object. Once the Bedoop object has been located within the image data, masking can be applied (if desired) to eliminate image data not corresponding to the intended object. The next step in the decoding process, determining orientation of the Bedoop data, can likewise be discerned by reference to visual clues. For example, some objects include subliminal graticule data, or other calibration data, steganographically encoded with the Bedoop data to aid in determining orientation. Others can employ overt markings, either placed for that sole purpose (e.g. reference lines or fiducials), or serving another purpose as well (e.g. lines of text), to discern orientation. Edge-detection algorithms can also be employed to deduce the orientation of the object by reference to its

Element	'036 Claim Recitation	Exemplary Citations in Reference
		edges." Rhoads at 16:66 – 17:10 • "Inside the building the candidate may encounter a security guard. Seeing an unfamiliar person, the guard may visually compare the photo on the badge with the candidate's face. Additionally, the guard can present the badge to a portable Bedoop device, or to one of many Bedoop systems scattered through the building (e g., at every telephone). The Bedoop system extracts the Bedoop data from the card (i.e., from the DMV photograph), interrogates the DMV's DNS server with this Bedoop data, and receives in reply the name of the person depicted in the photograph. (If the Bedoop system is a telephone, the name may be displayed on a small LCD display commonly provided on telephones.)" Rhoads at 22:9–27 • "To aid in identification, the Bedoop objects (e.g., badges) can be colored a distinctive color, permitting the system to more easily identify candidate objects from other items within the optical capture devices' field of view. Or the object can be provided with a retro-reflective coating, and the elevator can be equipped with one or more illumination sources of known spectral or temporal quality (e.g., constant infra red, or constant illumination with a single- or multi-line spectrum, or a pulsed light source of known periodicity; LEDs or semiconductor lasers, each with an associated diffuser, can be used for each the foregoing and can be paired with the image capture devices). Other such tell-tale clues can likewise be used to aid in object location. In all such cases, the optical capture device can sense the tell-tale clue(s) using a wide field of view sensor. The device can then be physically or electronically steered, and/or zoomed, to acquire a higher resolution image of the digitally-encoded object suitable for decoding."

Element	'036 Claim Recitation	Exemplary Citations in Reference
		Rhoads at 23:33–54 • "When a recipient of a business card holds it in front of a Bedoop sensor, the operating system on the local system launches a local Bedoop application. That local Bedoop application, in turn, establishes an external internet connection to a remote business card server. The address of that server may already be known to the local Bedoop application (e.g., having been stored from previous use), or the local Bedoop system can traverse the above-described public network of DNS servers to reach the business card server. A database on the business card name server maintains a large collection of business card data, one database record per UID. When that server receives Bedoop data from a local Bedoop system, it parses out the UID and accesses the corresponding database record. This record typically includes more information than is commonly printed on conventional business cards. Sample fields from the record may include, for example, name, title, office phone, office fax, home phone, home fax, cellular phone, email address, company name, corporate web page address, personal web page address, secretary's name, spouse's name, and birthday. This record is transmitted back to the originating Bedoop system." Rhoads at 48:32–53 • "As noted, the following disclosure focuses on one particular application—a system for linking print media to electronic content It should be reiterated, however, that the technology is not so limited, and may more generally be viewed as a system for linking any object (physical or electronic) to a corresponding networked or local resource. In accordance with an exemplary application of the below-detailed technology, digital watermarking is employed to convey a plural bit identifier within print media, such as magazine advertisements or articles, direct mail coupons or catalogs, bank- or credit-cards, and business cards.

Element	'036 Claim Recitation	Exemplary Citations in Reference
		This identifier is read by software in a user's computing device and forwarded to a remote database. The remote database identifies a URL corresponding to the identifier, and provides the URL back to the user's computer—permitting a browser on the user's computer to display the URL-identified web page. That web page can provide additional information or services—more timely and/or more extensive than that provided by the print material. By such arrangement, more efficient internet navigation and access is provided to consumers, and more effective means for linking readers to e-commerce points of sale is provided to advertisers." Rhoads at 50:27-60 • "In the former case, the handler 16 may identify a URL corresponding to the watermark (using the database 17), and return the URL to the application 28 c. Application 28 c can then pass the URL to a web browser 28 b in the device 12, and initiate a link to the internet site identified by the URL. Or the handler may have some locally stored data (e.g., audio or video, or software updates) and send it to the device 12 in response to the watermark. In the latter case, the handler 16 does not respond directly to the device 12. Instead, the handler responds by communicating with a remote resource 18. The communication can be as simple as logging receipt of the watermark message in a remote repository. Or it can be to authenticate device 12 (or a user thereof) to a remote resource in anticipation of a further transaction (e.g., the communication can form part of an on-line licensing or digital rights management transaction). Or the communication can request the remote resource to provide data or a service back to device 12 or to another destination (e.g., to initiate an FTP file transfer, or to request that a song selection identified by the watermark be downloaded to a user's personal music library, or to update software installed on device 12).

Element	'036 Claim Recitation	Exemplary Citations in Reference
		In still other cases, hybrids of the two foregoing cases can be employed, e.g., handler 16 can send some data back to device 12, while also communicating with a remote resource 18. In some cases, the response returned to the device 12 by handler 6 (or a remote resource 18) can serve to trigger some further action by the device 12. For example, the response returned to device 12 can include a WindowsMedia audio file, together with a request that the device 12 launch the WindowsMedia player installed on the device. (The launching of a browser pointed to a URL is another example of such triggering.)" Rhoads at 51:18-37 • "When shown a watermarked image, the application 28 c analyzes the image and extracts the embedded watermark payload (more particularly detailed below) from the image. The application sends some or all of this information in a message format to the router 14. The router 14 decodes the received message, looking for vendor and product information. Based on this information, it passes the message to a corresponding product handler 16. The product handler receives the message and attempts to match the detected watermark serial number to a registered watermark serial number earlier stored in the database 17. If a match is found, the product handler performs the desired action. As noted, typical actions include returning a URL for web redirection, serving up an HTML page for initial user navigation, initiating software downloads, etc. If a match is not found, the product handler returns an error code and message to the application 28 c. If a match is found, but the corresponding action is unavailable, incomplete, inactive or invalid, the product handler returns an error code and message to the calling application."

Element	'036 Claim Recitation	Exemplary Citations in Reference
		Rhoads at 52:38-51 • "In order for the system to identify the response (e.g., a URL) that corresponds to an object identifier (e.g., a watermark), this data must first be associated within the database 17 in association with the watermark to which it corresponds. The watermark registration process captures some basic identification information used later to validate the incoming message, and identifies the associated information/action In the illustrated example the identification information includes: Customer Account, Object and associated attributes (name, description, expiration, etc.), Action, and Registered Serial Number (for registration updates) The Customer Account identifies the watermark registrant. In most cases, this is also the party to be billed for services. For validation and security reasons, the Customer Account is required to be a known, existing account. Account information, including the account's password, is maintained by an Account Management system. The Object and associated attributes identifies the object to be watermarked. The object attributes typically include the name and description of the object and a list of accounts authorized to access the object's registration. These authorized "supporting" accounts are typically the ad agencies, pre-press houses, etc. involved in the watermark embedding process in the print advertising example contemplated herein. The Action defines the response the customer desires when the watermark is detected. It varies byproduct, but in the illustrative embodiment involves the return of some additional information regarding the watermarked object. In the illustrative system, the action is return of a URL or HTML to be used to display a web page associated with the watermarked object. For other products, the desired response may be display of the object's owner & rights

Element	**'036 Claim Recitation**	**Exemplary Citations in Reference**
		information, software/data downloads, delivery of streamed audio or video, presentation of an advertisement, initiation of object-based actions, etc. The Registered Serial Number forms the last component of the registration. It is this assigned vendor and product-unique identifier that allows the system to acquire the specific information/action for the object in question." To the extent it is found that Rhoads does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Rhoads does not anticipate this claim, Rhoads renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Rhoads with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1B]	an identification platform coupled with the target database, and that	Rhoads discloses an identification platform coupled with the target database. Rhoads at Fig. 2 (and accompanying description)

Element	'036 Claim Recitation	Exemplary Citations in Reference
		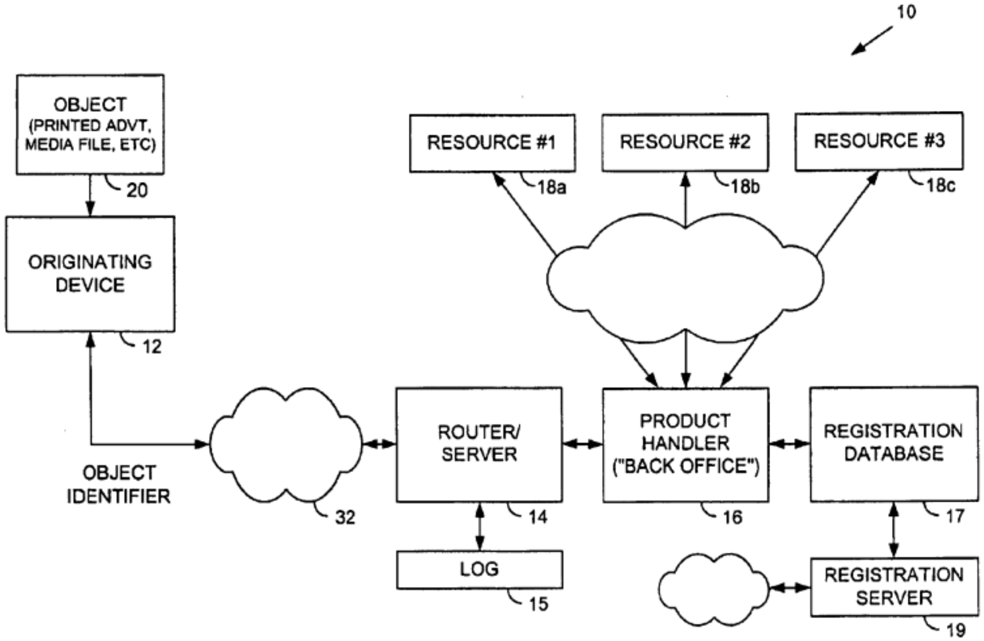 **FIG. 2** Rhoads at 50:21-48 • "As more particularly detailed below, the handler 16 provides a response in accordance with the particular watermark payload. The response may be provided directly by the product handler to the device 12, or the handler may respond by communicating with a remote resource 18 (which may be, e.g., a data repository or service provider). In the former case, the handler 16 may identify a URL corresponding to the watermark (using the database 17), and return the URL to the application 28 c. Application 28 c can then pass the URL to a web browser 28 b in the device 12, and initiate a link to the internet site identified by the URL. Or the handler may have some locally stored data (e.g., audio or video, or software updates) and send it to the device 12 in response to the watermark.

Element	'036 Claim Recitation	Exemplary Citations in Reference
		In the latter case, the handler 16 does not respond directly to the device 12. Instead, the handler responds by communicating with a remote resource 18. The communication can be as simple as logging receipt of the watermark message in a remote repository. Or it can be to authenticate device 12 (or a user thereof) to a remote resource in anticipation of a further transaction (e.g., the communication can form part of an on-line licensing or digital rights management transaction). Or the communication can request the remote resource to provide data or a service back to device 12 or to another destination (e.g., to initiate an FTP file transfer, or to request that a song selection identified by the watermark be downloaded to a user's personal music library, or to update software installed on device 12). In still other cases, hybrids of the two foregoing cases can be employed, e.g., handler 16 can send some data back to device 12, while also communicating with a remote resource 18. In some cases, the response returned to the device 12 by handler 6 (or a remote resource 18) can serve to trigger some further action by the device 12. For example, the response returned to device 12 can include a WindowsMedia audio file, together with a request that the device 12 launch the WindowsMedia player installed on the device. (The launching of a browser pointed to a URL is another example of such triggering.)" To the extent it is found that Rhoads does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Rhoads does not anticipate this claim, Rhoads renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Rhoads with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar

Element	’036 Claim Recitation	Exemplary Citations in Reference
		known problems, according to known methods which would yield predictable results.
[1B.i]	communicates with a mobile device capable of acquiring a digital representation of a scene containing at least a portion of a target;	Rhoads discloses communicates with a mobile device capable of acquiring a digital representation of a scene containing at least a portion of a target. Rhoads at Fig. 2 (and accompanying description) 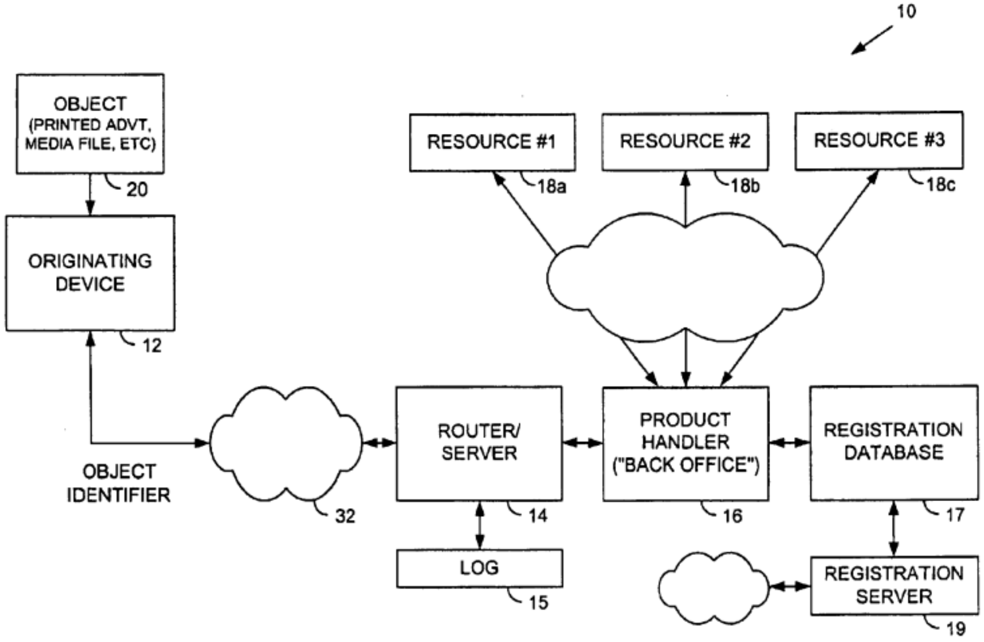 FIG. 2 Rhoads at 3:41–51 • “According to another aspect, the invention includes a promotional method comprising: (a) presenting an object within the field of view of an optical sensor device, the object being selected from the list consisting of a retail product, packaging for a retail product, or printed advertising; (b) acquiring

Element	'036 Claim Recitation	Exemplary Citations in Reference
		optical data corresponding to the object; (c) decoding plural-bit digital data from the optical data; (d) submitting at least some of said decoded data to a remote computer; and (e) determining at the remote computer whether a prize should be awarded in response to submission of said decoded data." Rhoads at 7:24-42 • "Device 12 interacts with an object 20. The object can be electronic or not. Electronic objects can include computer files, representations of audio, video, or still imagery (e.g., files or in streaming form), etc. Non-electronic objects can include physical objects such as newspapers, magazine pages, posters, product packaging, event tickets, credit cards, paper currency, etc. Non-electronic objects can also include sounds produced by loudspeakers. When used with nonelectronic objects, device 12 (FIG. 2) typically includes some form of sensor or transducer 22 to produce electronic signals or data corresponding to the object. Examples include CCD- or CMOS-based optical sensors (either as part of still- or video cameras, flatbed scanners, mice, or otherwise), microphones, barcode scanners, RF ID sensors, mag stripe readers, etc. In such cases, the sensor 22 may be coupled to associated interface electronics 24, which in turn may be coupled to device driver software 26, which in turn may be coupled to one or more application programs 28." Rhoads at Fig. 11 (and accompanying description)

Element	'036 Claim Recitation	Exemplary Citations in Reference
		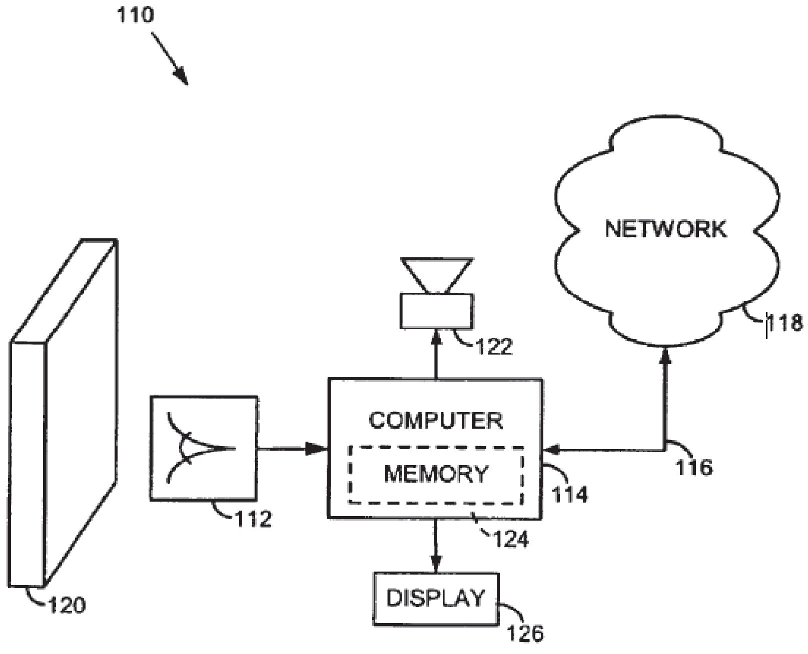 FIG. 11 Rhoads at 7:13-29 • "Referring to FIG. 11, a basic embodiment 110 of the present technology includes an optical sensor 112, a computer 114, and a network connection 116 to the internet 118. The illustrated optical sensor 112 is a digital camera having a resolution of 320 by 200 pixels (color or black and white) that stares out, grabbing frames of image data five times per second and storing same in one or more frame buffers. These frames of image data are analyzed

Element	'036 Claim Recitation	Exemplary Citations in Reference
		by a computer 114 for the presence of Bedoop data. (Essentially, Bedoop data is any form of plural-bit data encoding recognized by the system 110—data which, in many embodiments, initiates some action.) Once detected, the system responds in accordance with the detected Bedoop data (e.g., by initiating some local action, or by communication with a remote computer, such as over the internet, via an online service such as AOL, or using point-to-point dial-up communications, as with a bulletin board system)." Rhoads 8:66 – 9:12 • "In still other embodiments, the camera (or other sensor) can be equipped with one or more auxiliary fixed-focus lenses that can be selectively used, depending on the particular application. Some such embodiments have a first fixed focused lens that always overlies the sensor, with which one or more auxiliary lenses can be optically cascaded (e.g., by hinge or slide arrangements). Such arrangements are desirable, e.g., when a camera is not a dedicated Bedoop sensor but also performs other imaging tasks. When the camera is to be used for Bedoop, the auxiliary lens is positioned (e.g., flipped into) place, changing the focal length of the first lens (which may by unsuitably long for Bedoop purposes, such as infinity) to an appropriate Bedoop imaging range (such as one foot)." Rhoads at Fig. 3 (and accompanying description)

Element	'036 Claim Recitation	Exemplary Citations in Reference
		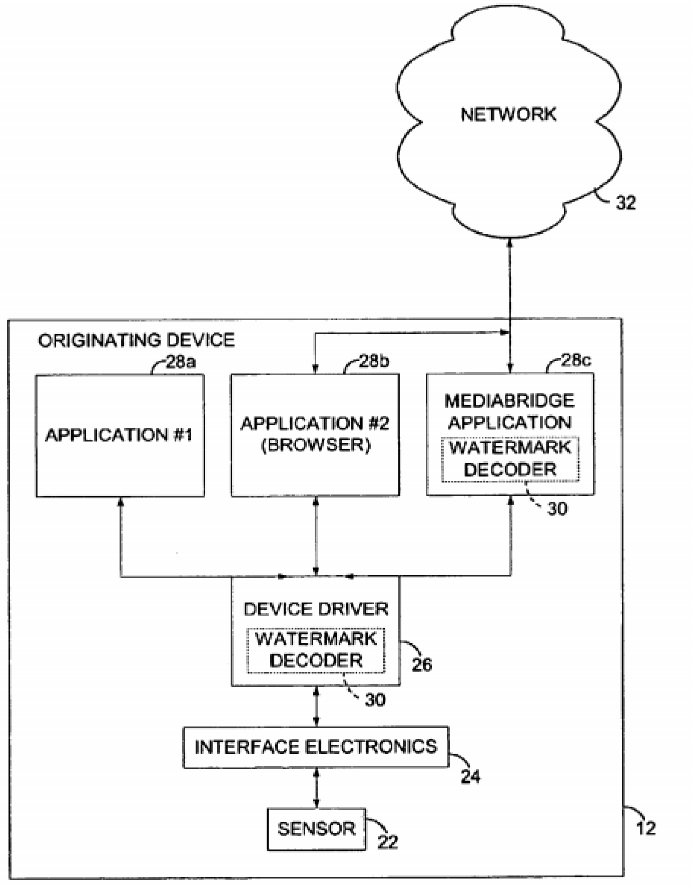 **FIG. 3** Rhoads at 17:46–54 • "The building can be provided with an image sensor (such as a video camera or the like), an associated Bedoop. detection system, and the proximity

Element	'036 Claim Recitation	Exemplary Citations in Reference
		detector. When a user wearing the badge approaches, the proximity detector signals the camera to capture image data. The Bedoop detection system identifies the badge photograph (e.g., by clues as are described in the prior applications, or without such aids), captures optical data, and decodes same to extract the steganographically-embedded data hidden therein." Rhoads at 46:47-53 • "The provision of image Sensors in cell phones enables other functionality. One is the capture of still or video imagery. Such image data can be compressed (typically by lossy processes such as MPEG, JPEG, or the like, implemented with dedicated hardware CODECs) and transmitted with the audio data. The Screens on Such phones can likewise be used for display of incoming image or Video data." Rhoads at 49:24-27 • "Device 12 interacts with an object 20. The object can be electronic or not. Electronic objects can include computer files, representations of audio, video, or still imagery (e.g., files or in streaming form), etc." To the extent it is found that Rhoads does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Rhoads does not anticipate this claim, Rhoads renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Rhoads with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar

Element	'036 Claim Recitation	Exemplary Citations in Reference
		known problems, according to known methods which would yield predictable results.
[1B.ii]	receives the digital representation from the mobile device; and	Rhoads discloses receives the digital representation from the mobile device. Rhoads at 5:21-23 - "As more particularly detailed below, the handler 16 provides a response in accordance with the particular watermark payload.". Rhoads at Fig. 2 (and accompanying description) 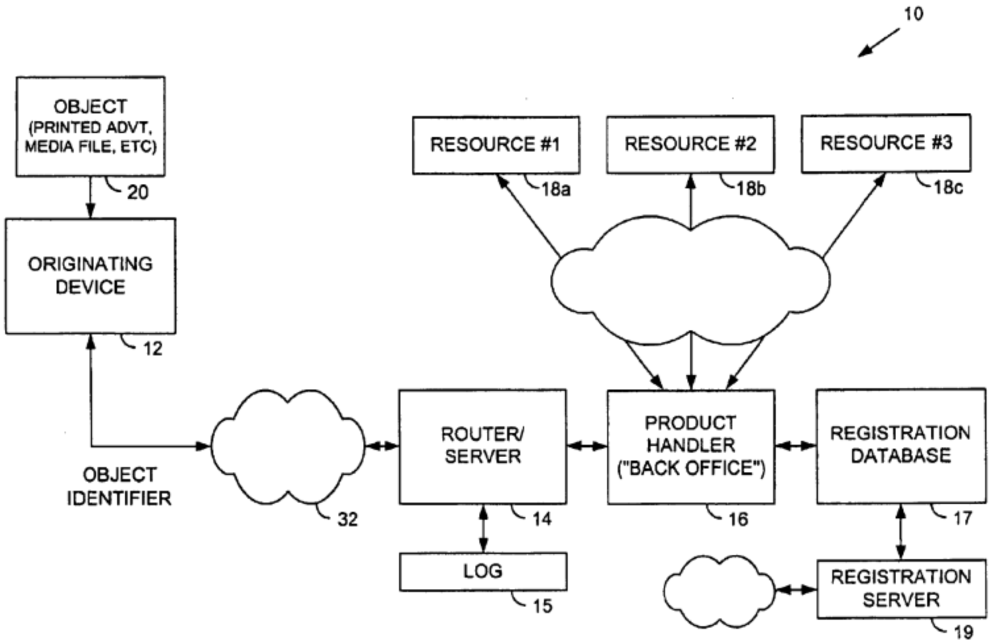 FIG. 2 Rhoads at 12:23 – 33

Element	'036 Claim Recitation	Exemplary Citations in Reference
		• "Sometimes the computer system may encounter a Bedoop object for which it does not have a registered application program. In such case, a default Bedoop application can be invoked. This default application can, e.g., establish an internet link to a remote server computer (or a network of such computers), and can transmit the Bedoop data (or a part of the Bedoop data) to that remote computer. The remote server can undertake the response itself, it can instruct the originating computer how to respond appropriately, or it can undertake some combination of these two responses. (Such arrangements are further considered below.)" Rhoads at 16:66 – 17:10 • "Inside the building the candidate may encounter a security guard. Seeing an unfamiliar person, the guard may visually compare the photo on the badge with the candidate's face. Additionally, the guard can present the badge to a portable Bedoop device, or to one of many Bedoop systems scattered through the building (e g., at every telephone). The Bedoop system extracts the Bedoop data from the card (i.e., from the DMV photograph), interrogates the DMV's DNS server with this Bedoop data, and receives in reply the name of the person depicted in the photograph. (If the Bedoop system is a telephone, the name may be displayed on a small LCD display commonly provided on telephones.)" To the extent it is found that Rhoads does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Rhoads does not anticipate this claim, Rhoads renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Rhoads with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because

Element	'036 Claim Recitation	Exemplary Citations in Reference
		each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results. Rhoads at 49:19-47 • "Device 12 interacts with an object 20. The object can be electronic or not. Electronic objects can include computer files, representations of audio, video, or still imagery (e.g., files or in streaming form), etc. Non-electronic objects can include physical objects such as newspapers, magazine pages, posters, product packaging, event tickets, credit cards, paper currency, etc. Non-electronic objects can also include sounds produced by loudspeakers. When used with nonelectronic objects, device 12 (FIG. 2) typically includes some form of sensor or transducer 22 to produce electronic signals or data corresponding to the object. Examples include CCD- or CMOS-based optical sensors (either as part of still- or video cameras, flatbed scanners, mice, or otherwise), microphones, barcode scanners, RF ID sensors, mag stripe readers, etc. In such cases, the sensor 22 may be coupled to associated interface electronics 24, which in turn may be coupled to device driver software 26, which in turn may be coupled to one or more application programs 28. Device driver software 26 serves as a software interface, communicating at a relatively high level with the application programs 28 (e.g., through API instructions whose content and format are standardized to facilitate application programming), and at a relatively low level with the interface electronics 24." Rhoads at 49:48-65 • "The detailed embodiment contemplates that the object 20 is a magazine advertisement encoded with a steganographic watermark conveying a plural bit object identifier. The watermark is hidden in the advertisement's image

Element	'036 Claim Recitation	Exemplary Citations in Reference
		in a manner indiscernable to human observers, but detectable by computer analysis. That analysis is performed by a watermark detector 30. Watermark detector 30 can be implemented in various different locations in the system of FIG. 1. Typically, the detector is implemented in the originating device 12, e.g., in the driver software 26, or in application software 28 c that serves to link to external resources based on detected watermarks. But it may be implemented elsewhere, e.g., in hardware in the interface electronics 24, in an operating system associated with the device, or outside device 12 altogether. Some systems may have plural watermark detectors, implemented at different locations throughout the system." Rhoads at 50:21–51:3 • "As more particularly detailed below, the handler 16 provides a response in accordance with the particular watermark payload. The response may be provided directly by the product handler to the device 12, or the handler may respond by communicating with a remote resource 18 (which may be, e.g., a data repository or service provider). In the former case, the handler 16 may identify a URL corresponding to the watermark (using the database 17), and return the URL to the application 28 c. Application 28 c can then pass the URL to a web browser 28 b in the device 12, and initiate a link to the internet site identified by the URL. Or the handler may have some locally stored data (e.g., audio or video, or software updates) and send it to the device 12 in response to the watermark. In the latter case, the handler 16 does not respond directly to the device 12. Instead, the handler responds by communicating with a remote resource 18. The communication can be as simple as logging receipt of the watermark message in a remote repository. Or it can be to authenticate device 12 (or a user thereof) to a remote resource in anticipation of a further transaction (e.g., the communication can form part of an on-line licensing or digital

Element	'036 Claim Recitation	Exemplary Citations in Reference
		rights management transaction). Or the communication can request the remote resource to provide data or a service back to device 12 or to another destination (e.g., to initiate an FTP file transfer, or to request that a song selection identified by the watermark be downloaded to a user's personal music library, or to update software installed on device 12). In still other cases, hybrids of the two foregoing cases can be employed, e.g., handler 16 can send some data back to device 12, while also communicating with a remote resource 18. In some cases, the response returned to the device 12 by handler 6 (or a remote resource 18) can serve to trigger some further action by the device 12. For example, the response returned to device 12 can include a WindowsMedia audio file, together with a request that the device 12 launch the WindowsMedia player installed on the device. (The launching of a browser pointed to a URL is another example of such triggering.) The illustrated product handler 16 comprises essentially the same hardware elements as the router 14, e.g., CPU, memory, etc. Although FIG. 2 shows just one product handler, several product handlers can be included in the system—either co-located of geographically distributed. Different handlers can be dedicated to different functions (e.g., serving URLs, serving music, etc.) or to different watermark sources (e.g., one responds to watermarks found in audio, another responds to watermarks found in print advertising, etc.)." Rhoads at Fig. 12 (and accompanying description)

Element	'036 Claim Recitation	Exemplary Citations in Reference
		FIG. 12 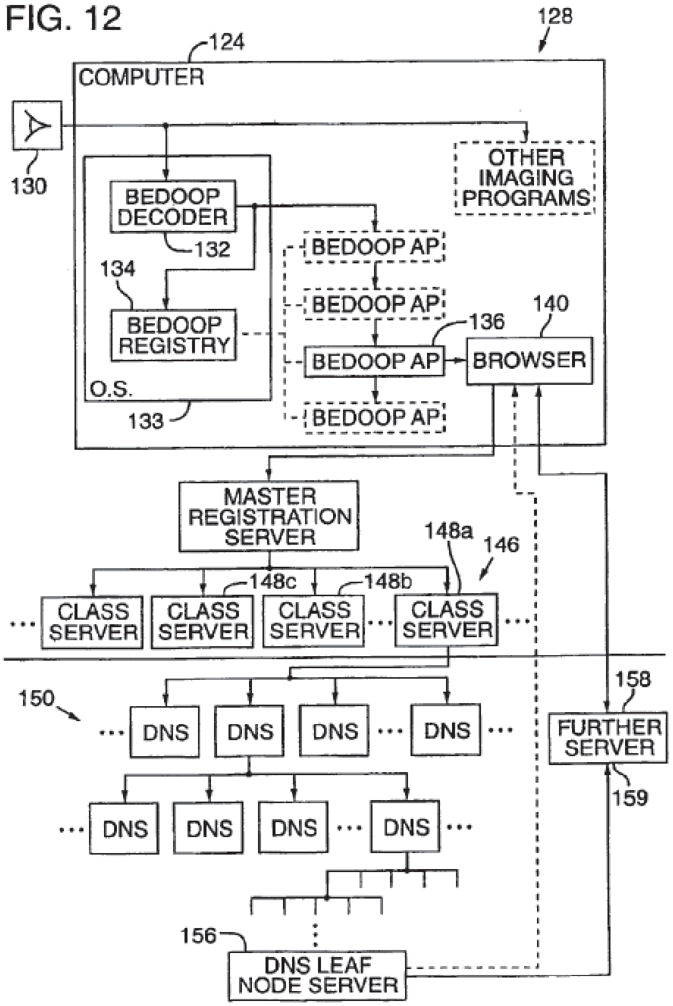 To the extent it is found that Rhoads does not expressly disclose this limitation, the

Element	'036 Claim Recitation	Exemplary Citations in Reference
		limitation is inherent. And to the extent it is found that Rhoads does not anticipate this claim, Rhoads renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Rhoads with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1B.iii]	recognizes the target as a known target from the target database based on comparing parameters derived from the digital representation to recognition parameters associated with the known targets; and	Rhoads discloses recognizes the target as a known target from the target database based on comparing parameters derived from the digital representation to recognition parameters associated with the known targets. Rhoads at 9:35–10:21 • "The Steganographic decoding process may entail three Steps. In the first, the object is located. In the Second, the objects orientation is discerned. In the third, the Bedoop data is extracted from the image data corresponding to the Bedoop object. The first step, object location, can be assisted by various clues. One is the placement of the object; typically the center of the image field will be a point on the object. The surrounding data can then be analyzed to try and discern the object's boundaries. Another location technique is slight movement. Although the user will typically try to hold the object still, there will usually be some jitter of the Bedoop object within the image frame (e.g., a few pixels back and forth). Background visual clutter, in contrast, will typically be stationary. Such

Element	'036 Claim Recitation	Exemplary Citations in Reference
		movement may thus be sensed and used to identify the Bedoop object from within the image data. Still another object-location clue is object shape. Many Bedoop objects are rectangular in shape (or trapezoidal as viewed by the camera). Straight edge boundaries can thus be used to define an area of likely Bedoop data. Color is a further object identification clue that may be useful in some contexts. Yet another object location clue is spatial frequency. In imaging systems with well defined focal zones, undesired visual clutter may be at focal distances that results in blurring. The Bedoop object, in contrast, will be in focus and may be characterized by fine detail. Analyzing the image data for the high frequencies associated with fine detail can be used to distinguish the intended object from others. Characteristic markings on the object (as discussed below in connection with determining object orientation) can also be sensed and used in locating the object. Once the Bedoop object has been located within the image data, masking can be applied (if desired) to eliminate image data not corresponding to the intended object. The next step in the decoding process, determining orientation of the Bedoop data, can likewise be discerned by reference to visual clues. For example, some objects include subliminal graticule data, or other calibration data, steganographically encoded with the Bedoop data to aid in determining orientation. Others can employ overt markings, either placed for that sole purpose (e.g. reference lines or fiducials), or serving another purpose as well (e.g. lines of text), to discern orientation. Edge-detection algorithms can also

Element	'036 Claim Recitation	Exemplary Citations in Reference
		be employed to deduce the orientation of the object by reference to its edges." Rhoads at 16:66 – 17:10 - "Inside the building the candidate may encounter a security guard. Seeing an unfamiliar person, the guard may visually compare the photo on the badge with the candidate's face. Additionally, the guard can present the badge to a portable Bedoop device, or to one of many Bedoop systems scattered through the building (e g., at every telephone). The Bedoop system extracts the Bedoop data from the card (i.e., from the DMV photograph), interrogates the DMV's DNS server with this Bedoop data, and receives in reply the name of the person depicted in the photograph. (If the Bedoop system is a telephone, the name may be displayed on a small LCD display commonly provided on telephones.)" Rhoads at 22:9–27 - "To aid in identification, the Bedoop objects (e.g., badges) can be colored a distinctive color, permitting the system to more easily identify candidate objects from other items within the optical capture devices' field of view. Or the object can be provided with a retro-reflective coating, and the elevator can be equipped with one or more illumination sources of known spectral or temporal quality (e.g., constant infra red, or constant illumination with a single- or multi-line spectrum, or a pulsed light source of known periodicity; LEDs or semiconductor lasers, each with an associated diffuser, can be used for each the foregoing and can be paired with the image capture devices). Other such tell-tale clues can likewise be used to aid in object location. In all such cases, the optical capture device can sense the tell-tale clue(s) using a wide field of view sensor. The device can then be physically or electronically steered, and/or zoomed, to acquire a higher resolution image of the digitally-encoded object suitable for decoding."

Element	'036 Claim Recitation	Exemplary Citations in Reference
		Rhoads at 23:33–54 • "When a recipient of a business card holds it in front of a Bedoop sensor, the operating system on the local system launches a local Bedoop application. That local Bedoop application, in turn, establishes an external internet connection to a remote business card server. The address of that server may already be known to the local Bedoop application (e.g., having been stored from previous use), or the local Bedoop system can traverse the above-described public network of DNS servers to reach the business card server. A database on the business card name server maintains a large collection of business card data, one database record per UID. When that server receives Bedoop data from a local Bedoop system, it parses out the UID and accesses the corresponding database record. This record typically includes more information than is commonly printed on conventional business cards. Sample fields from the record may include, for example, name, title, office phone, office fax, home phone, home fax, cellular phone, email address, company name, corporate web page address, personal web page address, secretary's name, spouse's name, and birthday. This record is transmitted back to the originating Bedoop system." Rhoads at 48:32–53 • "As noted, the following disclosure focuses on one particular application—a system for linking print media to electronic content It should be reiterated, however, that the technology is not so limited, and may more generally be viewed as a system for linking any object (physical or electronic) to a corresponding networked or local resource. In accordance with an exemplary application of the below-detailed technology, digital watermarking is employed to convey a plural bit identifier within print media, such as magazine advertisements or articles,

Element	'036 Claim Recitation	Exemplary Citations in Reference
		direct mail coupons or catalogs, bank- or credit-cards, and business cards. This identifier is read by software in a user's computing device and forwarded to a remote database. The remote database identifies a URL corresponding to the identifier, and provides the URL back to the user's computer—permitting a browser on the user's computer to display the URL-identified web page. That web page can provide additional information or services—more timely and/or more extensive than that provided by the print material. By such arrangement, more efficient internet navigation and access is provided to consumers, and more effective means for linking readers to e-commerce points of sale is provided to advertisers." Rhoads at 51:18-37 • "When shown a watermarked image, the application 28 c analyzes the image and extracts the embedded watermark payload (more particularly detailed below) from the image. The application sends some or all of this information in a message format to the router 14. The router 14 decodes the received message, looking for vendor and product information. Based on this information, it passes the message to a corresponding product handler 16. The product handler receives the message and attempts to match the detected watermark serial number to a registered watermark serial number earlier stored in the database 17. If a match is found, the product handler performs the desired action. As noted, typical actions include returning a URL for web redirection, serving up an HTML page for initial user navigation, initiating software downloads, etc. If a match is not found, the product handler returns an error code and message to the application 28 c. If a match is found, but the corresponding action is unavailable, incomplete, inactive or invalid, the product handler returns an error code and message to the calling application."

Element	'036 Claim Recitation	Exemplary Citations in Reference
		Rhoads at 84:11 – 17 • "B55. A method comprising: presenting an object to an optical sensor, the optical sensor producing image data, the object being steganographically encoded with plural bits of digital data; decoding said digital data from the image data; and in response to said digital data, displaying an image of a like object." To the extent it is found that Rhoads does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Rhoads does not anticipate this claim, Rhoads renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Rhoads with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1C]	a content service coupled with the identification platform, and that:	Rhoads discloses a content service coupled with the identification platform. Rhoads at Fig. 2 (and accompanying description)

Element	'036 Claim Recitation	Exemplary Citations in Reference
		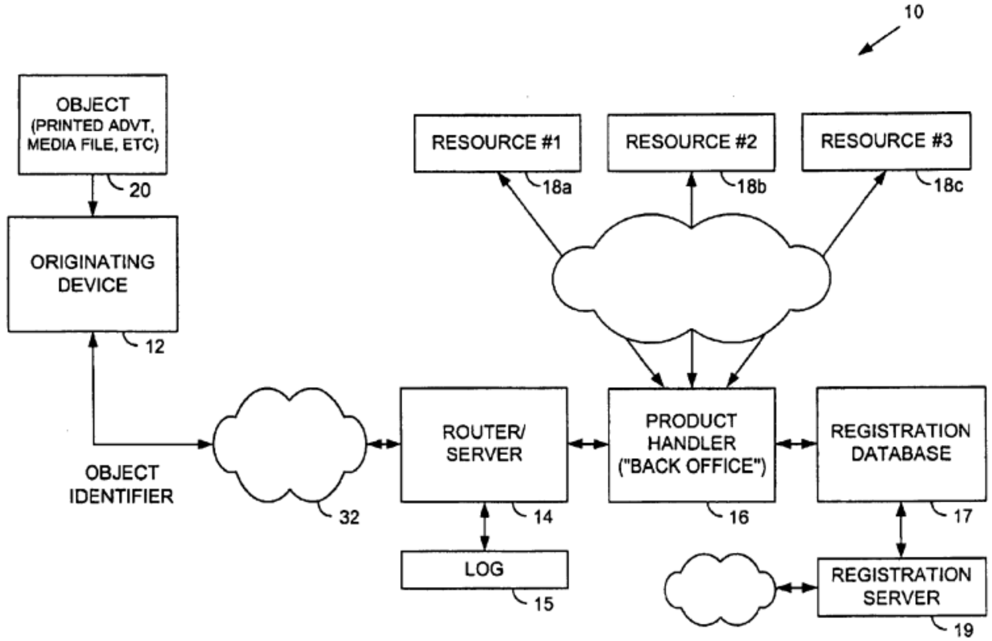 FIG. 2 Rhoads at 12:23 – 33 - "Sometimes the computer system may encounter a Bedoop object for which it does not have a registered application program. In such case, a default Bedoop application can be invoked. This default application can, e.g., establish an internet link to a remote server computer (or a network of such computers), and can transmit the Bedoop data (or a part of the Bedoop data) to that remote computer. The remote server can undertake the response itself, it can instruct the originating computer how to respond appropriately, or it can undertake some combination of these two responses. (Such arrangements are further considered below.)"

Element	'036 Claim Recitation	Exemplary Citations in Reference
		Rhoads at 16:66 – 17:10 • "Inside the building the candidate may encounter a security guard. Seeing an unfamiliar person, the guard may visually compare the photo on the badge with the candidate's face. Additionally, the guard can present the badge to a portable Bedoop device, or to one of many Bedoop systems scattered through the building (e g., at every telephone). The Bedoop system extracts the Bedoop data from the card (i.e., from the DMV photograph), interrogates the DMV's DNS server with this Bedoop data, and receives in reply the name of the person depicted in the photograph. (If the Bedoop system is a telephone, the name may be displayed on a small LCD display commonly provided on telephones.)" To the extent it is found that Rhoads does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Rhoads does not anticipate this claim, Rhoads renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Rhoads with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results. Rhoads at 49:19-47 • "Device 12 interacts with an object 20. The object can be electronic or not. Electronic objects can include computer files, representations of audio, video, or still imagery (e.g., files or in streaming form), etc. Non-electronic objects can include physical objects such as newspapers, magazine pages, posters, product packaging, event tickets, credit cards, paper currency, etc.

Element	'036 Claim Recitation	Exemplary Citations in Reference
		Non-electronic objects can also include sounds produced by loudspeakers. When used with nonelectronic objects, device 12 (FIG. 2) typically includes some form of sensor or transducer 22 to produce electronic signals or data corresponding to the object. Examples include CCD- or CMOS-based optical sensors (either as part of still- or video cameras, flatbed scanners, mice, or otherwise), microphones, barcode scanners, RF ID sensors, mag stripe readers, etc. In such cases, the sensor 22 may be coupled to associated interface electronics 24, which in turn may be coupled to device driver software 26, which in turn may be coupled to one or more application programs 28. Device driver software 26 serves as a software interface, communicating at a relatively high level with the application programs 28 (e.g., through API instructions whose content and format are standardized to facilitate application programming), and at a relatively low level with the interface electronics 24." Rhoads at 49:48-65 • "The detailed embodiment contemplates that the object 20 is a magazine advertisement encoded with a steganographic watermark conveying a plural bit object identifier. The watermark is hidden in the advertisement's image in a manner indiscernable to human observers, but detectable by computer analysis. That analysis is performed by a watermark detector 30. Watermark detector 30 can be implemented in various different locations in the system of FIG. 1. Typically, the detector is implemented in the originating device 12, e.g., in the driver software 26, or in application software 28 c that serves to link to external resources based on detected watermarks. But it may be implemented elsewhere, e.g., in hardware in the interface electronics 24, in an operating system associated with the device, or outside device 12 altogether. Some systems may have plural watermark detectors, implemented at different locations throughout the system."

Element	'036 Claim Recitation	Exemplary Citations in Reference
		Rhoads at 50:21–51:3 • "As more particularly detailed below, the handler 16 provides a response in accordance with the particular watermark payload. The response may be provided directly by the product handler to the device 12, or the handler may respond by communicating with a remote resource 18 (which may be, e.g., a data repository or service provider). In the former case, the handler 16 may identify a URL corresponding to the watermark (using the database 17), and return the URL to the application 28 c. Application 28 c can then pass the URL to a web browser 28 b in the device 12, and initiate a link to the internet site identified by the URL. Or the handler may have some locally stored data (e.g., audio or video, or software updates) and send it to the device 12 in response to the watermark. In the latter case, the handler 16 does not respond directly to the device 12. Instead, the handler responds by communicating with a remote resource 18. The communication can be as simple as logging receipt of the watermark message in a remote repository. Or it can be to authenticate device 12 (or a user thereof) to a remote resource in anticipation of a further transaction (e.g., the communication can form part of an on-line licensing or digital rights management transaction). Or the communication can request the remote resource to provide data or a service back to device 12 or to another destination (e.g., to initiate an FTP file transfer, or to request that a song selection identified by the watermark be downloaded to a user's personal music library, or to update software installed on device 12). In still other cases, hybrids of the two foregoing cases can be employed, e.g., handler 16 can send some data back to device 12, while also communicating with a remote resource 18. In some cases, the response returned to the device 12 by handler 6 (or a

Element	'036 Claim Recitation	Exemplary Citations in Reference
		remote resource 18) can serve to trigger some further action by the device 12. For example, the response returned to device 12 can include a WindowsMedia audio file, together with a request that the device 12 launch the WindowsMedia player installed on the device. (The launching of a browser pointed to a URL is another example of such triggering.) The illustrated product handler 16 comprises essentially the same hardware elements as the router 14, e.g., CPU, memory, etc. Although FIG. 2 shows just one product handler, several product handlers can be included in the system—either co-located of geographically distributed. Different handlers can be dedicated to different functions (e.g., serving URLs, serving music, etc.) or to different watermark sources (e.g., one responds to watermarks found in audio, another responds to watermarks found in print advertising, etc.)." Rhoads at Fig. 12 (and accompanying description)

Element	'036 Claim Recitation	Exemplary Citations in Reference
		FIG. 12 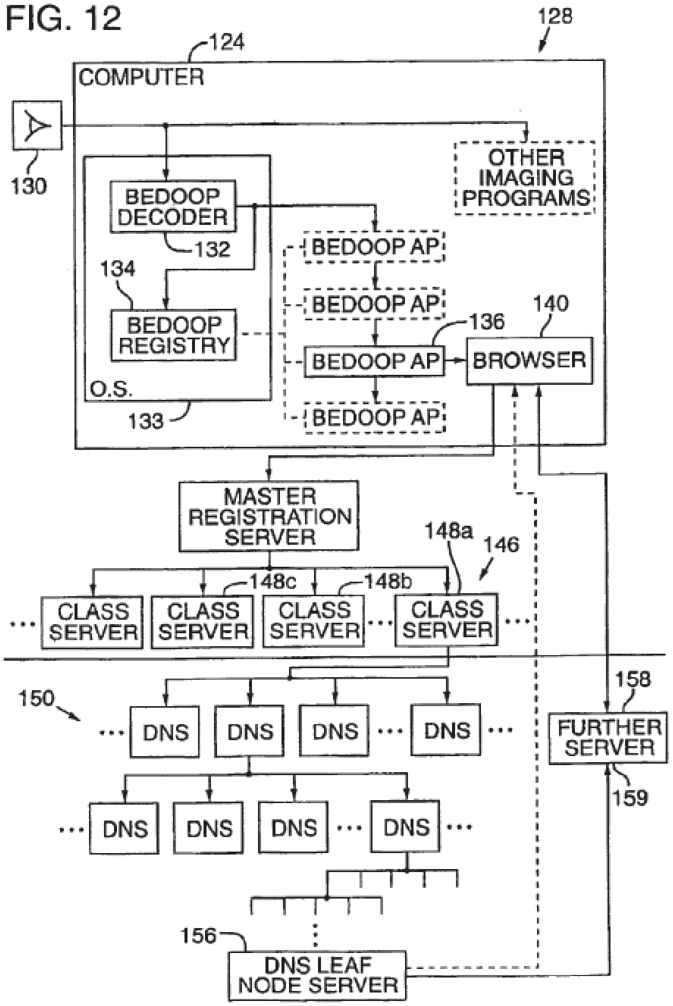 To the extent it is found that Rhoads does not expressly disclose this limitation, the

Element	'036 Claim Recitation	Exemplary Citations in Reference
		limitation is inherent. And to the extent it is found that Rhoads does not anticipate this claim, Rhoads renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Rhoads with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1C.i]	obtains content information related to the known target; and	Rhoads discloses obtains content information related to the known target. Rhoads at 22:28–33 • "Magazine (and newspaper) pages can be steganographically encoded with Bedoop data to provide another "paper as portal" experience. As with the earlier-described office document case, the encoded data yields an address to a computer location (e.g., a web page) having the same, or related, content." Rhoads at 48:32–53 • "As noted, the following disclosure focuses on one particular application—a system for linking print media to electronic content It should be reiterated, however, that the technology is not so limited, and may more generally be viewed as a system for linking any object (physical or electronic) to a corresponding networked or local resource. In accordance with an exemplary application of the below-detailed technology, digital watermarking is employed to convey a plural bit

Element	'036 Claim Recitation	Exemplary Citations in Reference
		identifier within print media, such as magazine advertisements or articles, direct mail coupons or catalogs, bank- or credit-cards, and business cards. This identifier is read by software in a user's computing device and forwarded to a remote database. The remote database identifies a URL corresponding to the identifier, and provides the URL back to the user's computer—permitting a browser on the user's computer to display the URL-identified web page. That web page can provide additional information or services—more timely and/or more extensive than that provided by the print material. By such arrangement, more efficient internet navigation and access is provided to consumers, and more effective means for linking readers to e-commerce points of sale is provided to advertisers." Rhoads at 50:35-60 • "In the latter case, the handler 16 does not respond directly to the device 12. Instead, the handler responds by communicating with a remote resource 18. The communication can be as simple as logging receipt of the watermark message in a remote repository. Or it can be to authenticate device 12 (or a user thereof) to a remote resource in anticipation of a further transaction (e.g., the communication can form part of an on-line licensing or digital rights management transaction). Or the communication can request the remote resource to provide data or a service back to device 12 or to another destination (e.g., to initiate an FTP file transfer, or to request that a song selection identified by the watermark be downloaded to a user's personal music library, or to update software installed on device 12). In still other cases, hybrids of the two foregoing cases can be employed, e.g., handler 16 can send some data back to device 12, while also communicating with a remote resource 18. In some cases, the response returned to the device 12 by handler 6 (or a remote resource 18) can serve to trigger some further action by the device 12. For example, the response returned to device 12 can include a WindowsMedia audio file, together with a request that the device 12 launch

Element	'036 Claim Recitation	Exemplary Citations in Reference
		the WindowsMedia player installed on the device. (The launching of a browser pointed to a URL is another example of such triggering.)" Rhoads at 51:18-37 • "When shown a watermarked image, the application 28 c analyzes the image and extracts the embedded watermark payload (more particularly detailed below) from the image. The application sends some or all of this information in a message format to the router 14. The router 14 decodes the received message, looking for vendor and product information. Based on this information, it passes the message to a corresponding product handler 16. The product handler receives the message and attempts to match the detected watermark serial number to a registered watermark serial number earlier stored in the database 17. If a match is found, the product handler performs the desired action. As noted, typical actions include returning a URL for web redirection, serving up an HTML page for initial user navigation, initiating software downloads, etc. If a match is not found, the product handler returns an error code and message to the application 28 c. If a match is found, but the corresponding action is unavailable, incomplete, inactive or invalid, the product handler returns an error code and message to the calling application." Rhoads at 52:3-29 • "[user] Sees the data/action associated with the object (e.g., views web page)" To the extent it is found that Rhoads does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Rhoads does not anticipate

Element	'036 Claim Recitation	Exemplary Citations in Reference
		this claim, Rhoads renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Rhoads with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1C.ii]	sends the content information to at least one of the identification platform and the mobile device.	Rhoads discloses sends the content information to at least one of the identification platform and the mobile device. Rhoads at 22:28–33 • "Magazine (and newspaper) pages can be steganographically encoded with Bedoop data to provide another "paper as portal" experience. As with the earlier-described office document case, the encoded data yields an address to a computer location (e.g., a web page) having the same, or related, content." Rhoads at 48:32–53 • "As noted, the following disclosure focuses on one particular application—a system for linking print media to electronic content It should be reiterated, however, that the technology is not so limited, and may more generally be viewed as a system for linking any object (physical or electronic) to a corresponding networked or local resource. In accordance with an exemplary application of the below-detailed technology, digital watermarking is employed to convey a plural bit

Element	'036 Claim Recitation	Exemplary Citations in Reference
		identifier within print media, such as magazine advertisements or articles, direct mail coupons or catalogs, bank- or credit-cards, and business cards. This identifier is read by software in a user's computing device and forwarded to a remote database. The remote database identifies a URL corresponding to the identifier, and provides the URL back to the user's computer—permitting a browser on the user's computer to display the URL-identified web page. That web page can provide additional information or services—more timely and/or more extensive than that provided by the print material. By such arrangement, more efficient internet navigation and access is provided to consumers, and more effective means for linking readers to e-commerce points of sale is provided to advertisers." Rhoads at 50:35-60 • "In the latter case, the handler 16 does not respond directly to the device 12. Instead, the handler responds by communicating with a remote resource 18. The communication can be as simple as logging receipt of the watermark message in a remote repository. Or it can be to authenticate device 12 (or a user thereof) to a remote resource in anticipation of a further transaction (e.g., the communication can form part of an on-line licensing or digital rights management transaction). Or the communication can request the remote resource to provide data or a service back to device 12 or to another destination (e.g., to initiate an FTP file transfer, or to request that a song selection identified by the watermark be downloaded to a user's personal music library, or to update software installed on device 12). In still other cases, hybrids of the two foregoing cases can be employed, e.g., handler 16 can send some data back to device 12, while also communicating with a remote resource 18. In some cases, the response returned to the device 12 by handler 6 (or a remote resource 18) can serve to trigger some further action by the device 12. For example, the response returned to device 12 can include a WindowsMedia audio file, together with a request that the device 12 launch

Element	'036 Claim Recitation	Exemplary Citations in Reference
		the WindowsMedia player installed on the device. (The launching of a browser pointed to a URL is another example of such triggering.)" Rhoads at 51:18-37 - "When shown a watermarked image, the application 28 c analyzes the image and extracts the embedded watermark payload (more particularly detailed below) from the image. The application sends some or all of this information in a message format to the router 14. The router 14 decodes the received message, looking for vendor and product information. Based on this information, it passes the message to a corresponding product handler 16. The product handler receives the message and attempts to match the detected watermark serial number to a registered watermark serial number earlier stored in the database 17. If a match is found, the product handler performs the desired action. As noted, typical actions include returning a URL for web redirection, serving up an HTML page for initial user navigation, initiating software downloads, etc. If a match is not found, the product handler returns an error code and message to the application 28 c. If a match is found, but the corresponding action is unavailable, incomplete, inactive or invalid, the product handler returns an error code and message to the calling application." Rhoads at 52:3-29 - "[user] Sees the data/action associated with the object (e.g., views web page)" To the extent it is found that Rhoads does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Rhoads does not

Element	'036 Claim Recitation	Exemplary Citations in Reference
		anticipate this claim, Rhoads renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Rhoads with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

Exhibit C-4

Bank of America Corporation ("BofA" or "Defendant") provides this chart subject to all reservations, objections, statements, and disclaimers set forth herein and in BofA's Final Invalidity Contentions Cover Pleading, as well as any amendment, supplement, or modification thereof, which are incorporated herein by reference in their entirety.

On information and belief, Mault was filed on September 7, 2001, claiming the priority date of September 7, 2000. Mault published at least as of April 25, 2002, and is available as prior art at least under 35 U.S.C. § 102(e).

As shown in the chart below, Mault anticipates asserted claim 1 of the '036 Patent. To the extent Mault is found not to expressly disclose certain limitations in the asserted claims, such limitations are inherent. To the extent Mault is found not to anticipate any asserted claims or claim elements of the '036 Patent, the reference nevertheless renders those claims or claim elements obvious under 35 U.S.C. § 103, either alone or in combination with other art identified herein or in the cover pleading of BofA's Final Invalidity Contentions, for the reasons identified in the cover pleading. These Final Invalidity Contentions are not an admission by BofA that the accused products, including any current or past versions of the accused products, are covered by, or infringe the asserted claims, but are based instead on the recognition that if the claims are interpreted to be broad enough to encompass the accused products, the claims must also be construed to have that same scope when considering whether they are invalid.

The following chart is based on the Court's claim construction rulings (Dkt. Nos. 145, 153, 236), and to the extent not contradictory, Plaintiffs' apparent interpretations of the claims in Plaintiffs' "Preliminary" Final Infringement Contentions. BofA notes that in many instances, Plaintiffs' "Preliminary" Final Infringement Contentions ignore and/or contradict the Court's claim construction rulings. BofA does not accept the assumptions concerning the scope and meaning implicit in Plaintiffs' Preliminary Final Infringement Contentions, to the extent those assumptions are discernible, and reserves the right to challenge Plaintiffs' proposed (or implied) constructions. Furthermore, Plaintiffs' "Preliminary" Final Contentions fail to comply with the level of specificity required by SPR 2.1.2 and 2.1.3, and include improper reservation of rights to supplement with "additional ways in which BoA Accused Products infringe." Any alleged deficient discovery is of Plaintiffs own making. Most if not all of the deficiencies in Plaintiffs' Preliminary Final Contentions have been documented by Defendants through multiple letters since at least May 18, 2021 and up to Bank of America's most recent letter of August 19, 2023. Plaintiffs' shifting infringement contentions, deficiencies in its infringement contentions, and changing claim scope have prejudiced BofA from being able to reasonably prepare its defenses, including its invalidity contentions. BofA also reserves the right to revise and supplement these charts if and when Plaintiffs are permitted to provide revised or supplemental Infringement Contentions. Citations given in the chart below are merely representative of the respective elements to provide notice to Plaintiffs of BofA's invalidity theories and are not meant to be exhaustive.

1. Claim 1

Claim	'036 Claim Recitation	Exemplary Citations in Reference
[1Pre]	A content provisioning system comprising:	Mault discloses a content provisioning system. Mault at Abstract • "A method for assisting a person create a record of food items consumed by the person comprises the capturing of food images corresponding to food items consumed, for example by creating an image using an optical image sensor; recording the food images in the memory of an electronic device; viewing the food images on the display of the electronic device, so as to identify food items consumed; selecting food item identifiers corresponding to the food images, and recording food item identifiers in a memory, So as to create a record of food items consumed." Mault at [0038] • "A person can be provided with a portable electronic device, having an image creation capability. For example, the portable electronic device may be a personal digital assistant (PDA) having an optical image sensor. The optical image sensor can be an accessory to the PDA, or can be within the housing of the PDA." To the extent it is found that Mault does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Mault does not anticipate this claim, Mault renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Mault with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to

Claim	'036 Claim Recitation	Exemplary Citations in Reference
		known methods which would yield predictable results.
[1A]	a target database storing known targets of different types and recognition parameters associated with the known targets;	Mault discloses a target database storing known targets of different types and recognition parameters associated with the known targets. Mault at [0016] • "The food image can be recorded in a memory of the electronic device, or any other memory accessible to the electronic device, such as a memory module, remote database, and the like. The food item name can be selected from one or more lists of food item names presented on the display of the electronic device, for example from one or more menus presented on the display." Mault at [0017] • "A method of creating a diet log comprises: recording an image of a label, the label corresponding to a food item consumed; analyzing the image of the label (for example using optical character recognition, image analysis, image recognition, or neural network techniques) so as to determine the identity of the food item consumed; and storing the determined identity of the food item consumed in an electronic memory, so as to create a diet log. The label can be a printed label, a barcode, an alphanumeric code, a nutritional data label, a list of ingredients, a product code, graphical designs (Such as trademarks), product names, and the like." Mault at [0018] • "A method by which a person can create a record of foods consumed comprises: recording a food image of a food consumed; transmitting the food image to a remote computer system, for example over a communications network Such as the internet, receiving food identification data from the remote computer System (which may be generated by a software running on

Claim	'036 Claim Recitation	Exemplary Citations in Reference
		the remote computer System, a computer expert System, an image analysis system, a human operative in communication with the remote computer system, or Some combination of techniques); and recording the food identification data So as to create a record of foods consumed." Mault at [0046] • "Other options which may be presented to the user include the entering of a food by name, for example into diet log software, or scanning the barcode of a food item. The electronic device may comprise or communicate with a barcode scanner, or the optical image sensor can be used to record an image of the barcode. The barcode can then be interpreted in terms of a UPC (universal product code) and/or product name using a database accessible to the PDA. A UPC database may be accessible over a communications network such as the Internet." Mault at [0048] • "FIG. 5 is a schematic of a system comprising an electronic device 100, an imaging System 102, a physiological Sensor 104, an activity Sensor 106, and a communications link through communications network 108 to a remote computer system (such as a server system) 110." Mault at Fig. 5

Claim	'036 Claim Recitation	Exemplary Citations in Reference
		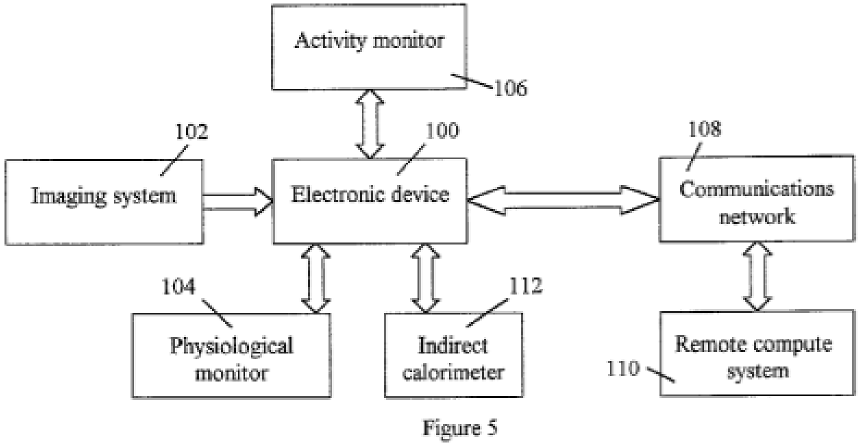 Figure 5 Mault at [0027] - " FIG. 5 illustrates a system allowing diet log creation;" Mault at [0048] - FIG. 5 is a schematic of a system comprising an electronic device 100, an imaging system 102, a physiological sensor 104, an activity sensor 106, and a communications link through communications network 108 to a remote computer system (such as a server system) 110. The electronic device can be a PDA or other device carried by the person. The imaging device can be included with the electronic device in a unitary device. The physiological sensor may detect eating, for example a blood glucose sensor can detect the rise in blood glucose after a meal, allowing the electronic device to alert the person to make a diet log entry. The activity sensor provides a signal correlated with a physical activity level of the person. Mault at [0049]

Claim	'036 Claim Recitation	Exemplary Citations in Reference
		• "Images recorded with the imaging device are transmitted over the communications network to the remote computer. The remote computer system can be a server system, and may comprise a computer expert system for image analysis. Image analysis can be performed by one or more persons, a computer expert system, or other mechanism." Mault at [0051] • "FIG. 6 is a flow chart of a method for diet logging. Box **120** corresponds to recording an image of food. Box **122** corresponds to transmission of the image to a remote location, Such as a remote computer System. . . . Box **124** corresponds to analysis of the image of the remote location." Mault at Fig. 6

Claim	'036 Claim Recitation	Exemplary Citations in Reference
		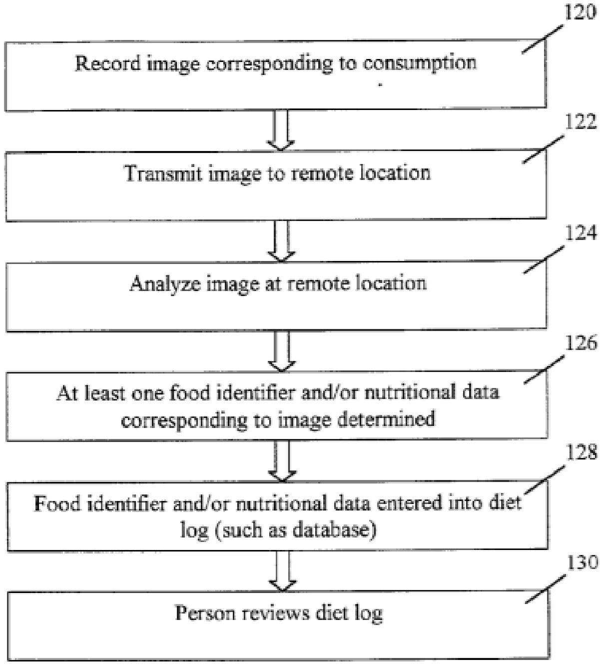 Figure 6 To the extent it is found that Mault does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Mault does not anticipate this claim, Mault renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Mault with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to

Claim	'036 Claim Recitation	Exemplary Citations in Reference
		known methods which would yield predictable results.
[1B]	an identification platform coupled with the target database, and that	Mault discloses an identification platform coupled with the target database. Mault at [0017] • "A method of creating a diet log comprises: recording an image of a label, the label corresponding to a food item consumed; analyzing the image of the label (for example using optical character recognition, image analysis, image recognition, or neural network techniques) so as to determine the identity of the food item consumed; and storing the determined identity of the food item consumed in an electronic memory, so as to create a diet log. The label can be a printed label, a barcode, an alphanumeric code, a nutritional data label, a list of ingredients, a product code, graphical designs (Such as trademarks), product names, and the like." Mault at [0018] • "A method by which a person can create a record of foods consumed comprises: recording a food image of a food consumed; transmitting the food image to a remote computer system, for example over a communications network Such as the internet, receiving food identification data from the remote computer System (which may be generated by a software running on the remote computer System, a computer expert System, an image analysis system, a human operative in communication with the remote computer system, or Some combination of techniques); and recording the food identification data So as to create a record of foods consumed." Mault at [0046] • "Other options which may be presented to the user include the entering of a food by name, for example into diet log software, or scanning the barcode of a food item. The electronic device may comprise or communicate with a

Claim	'036 Claim Recitation	Exemplary Citations in Reference
		barcode scanner, or the optical image sensor can be used to record an image of the barcode. The barcode can then be interpreted in terms of a UPC (universal product code) and/or product name using a database accessible to the PDA. A UPC database may be accessible over a communications network such as the Internet. Other options include receiving transmissions such as wireless data transmissions from food vendors or food vending devices, viewing the recorded diet log, editing the diet log, or analyzing previously recorded images." Mault at [0048] - "FIG. 5 is a schematic of a system comprising an electronic device 100, an imaging System 102, a physiological sensor 104, an activity sensor 106, and a communications link through communications network 108 to a remote computer system (such as a server system) 110." Mault at Fig. 5 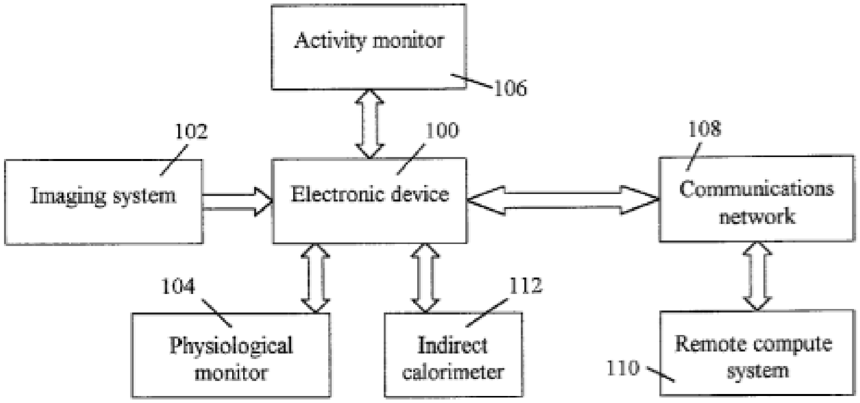 Figure 5

Claim	'036 Claim Recitation	Exemplary Citations in Reference
		Mault at [0027] • " FIG. 5 illustrates a system allowing diet log creation;" Mault at [0048] • FIG. 5 is a schematic of a system comprising an electronic device 100, an imaging system 102, a physiological sensor 104, an activity sensor 106, and a communications link through communications network 108 to a remote computer system (such as a server system) 110. The electronic device can be a PDA or other device carried by the person. The imaging device can be included with the electronic device in a unitary device. The physiological sensor may detect eating, for example a blood glucose sensor can detect the rise in blood glucose after a meal, allowing the electronic device to alert the person to make a diet log entry. The activity sensor provides a signal correlated with a physical activity level of the person. Mault at [0049] • "Images recorded with the imaging device are transmitted over the communications network to the remote computer. The remote computer system can be a server system, and may comprise a computer expert system for image analysis. Image analysis can be performed by one or more persons, a computer expert system, or other mechanism." Mault at [0051] • "Box **122** corresponds to transmission of the image to a remote location, such as a remote computer system. Other data can be transmitted with the image, for example a sales receipt, data provided by the food vendor, location the image was taken (which may be correlated with a restaurant identity), a menu image, and the weight of the food (which may be provided by a mat). Box **124** corresponds to analysis of the image of the remote location. Image

Claim	'036 Claim Recitation	Exemplary Citations in Reference
		recognition software may be used. A trained person, at a remote location, can select food names corresponding to the images, for example in conjunction with image recognition, providing names where image recognition has not been successful in identifying the food. A computer expert System, optical character recognition of a name on the package, barcode recognition, portion size estimate, and other food data can also be extracted from the image or images." Mault at Fig. 6

Claim	'036 Claim Recitation	Exemplary Citations in Reference
		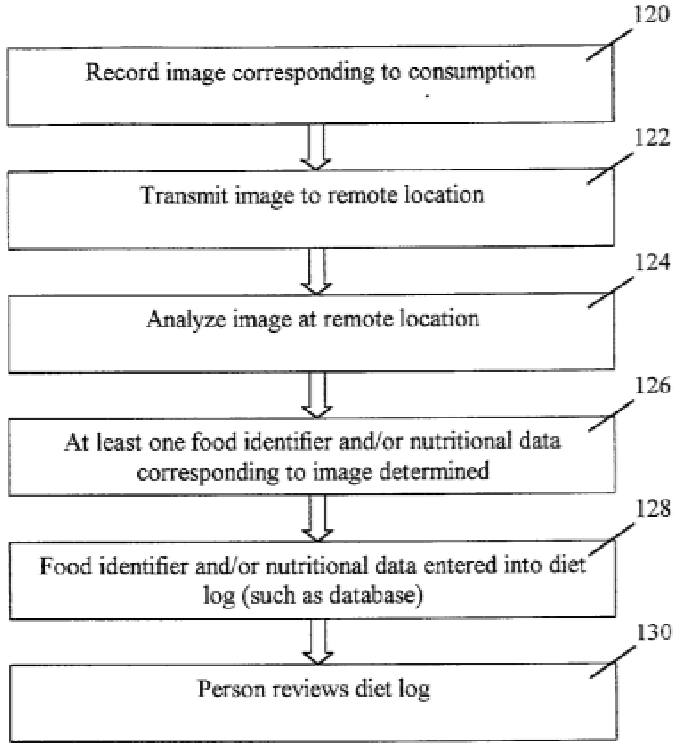 Figure 6 To the extent it is found that Mault does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Mault does not anticipate this claim, Mault renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Mault with any of the other prior art references

Claim	’036 Claim Recitation	Exemplary Citations in Reference
		identified in the cover pleading or herein with a reasonable expectation of success because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1B.i]	communicates with a mobile device capable of acquiring a digital representation of a scene containing at least a portion of a target;	Mault discloses communicates with a mobile device capable of acquiring a digital representation of a scene containing at least a portion of a target. Mault at Fig. 5 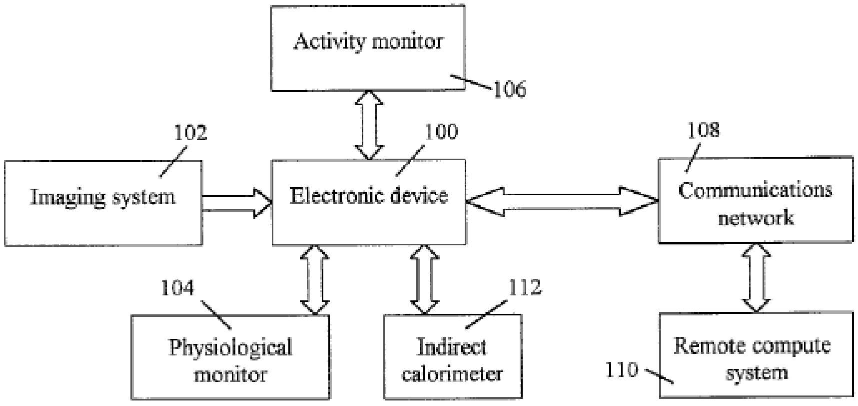 Figure 5 Mault at [0008] • “The image can be stored, for example, on the memory of an imaging device such as a camera, a memory of a device in communication with the imaging device, a memory module (such as a memory stick), a memory of a portable computing device, or on a memory associated with a remote computer system, such as one accessible over a communications network. In this context, the term image refers to an electronic file, or equivalent, containing information corresponding to an image. An example would be a conventional

Claim	'036 Claim Recitation	Exemplary Citations in Reference
		JPEG file." Mault at [0009] • "After recording the image, at a convenient later time the person can retrieve a stored image and view the image on a display. The display can be part of the imaging device, or other device such as a personal digital assistant, personal computer, Internet appliance, WebTV, e-book, tablet computer, pager, cell phone, interactive TV, spectacle mounted display, visor, and the like. An electronic device having a display used to view an image may be in communication with a separate device used to create the image, for example using wireless or cable-based links." Mault at [0011] • "The imaging device, used to create an image of foods consumed, can be a digital camera, device having the functionality of a camera, such as a personal digital assistant (PDA), cell phone, electronic book, pager, calculator, or other consumer electronics device. The imaging device can also be an ornamental component, clothing item, or other functional or decorative item having image creation capabilities, Such as spectacles, buttons, a pen, cutlery item, key ring, keys, belt, cigarette-like object, ring, body ornament, or jewelry." Mault at [0017] • "A method of creating a diet log comprises: recording an image of a label, the label corresponding to a food item consumed; analyzing the image of the label (for example using optical character recognition, image analysis, image recognition, or neural network techniques) so as to determine the identity of the food item consumed; and storing the determined identity of the food item consumed in an electronic memory, so as to create a diet log. The label can be a printed label, a barcode, an alphanumeric code, a nutritional data label, a

Claim	'036 Claim Recitation	Exemplary Citations in Reference
		list of ingredients, a product code, graphical designs (Such as trademarks), product names, and the like." Mault at [0018] • "A method by which a person can create a record of foods consumed comprises: recording a food image of a food consumed; transmitting the food image to a remote computer system, for example over a communications network Such as the internet, receiving food identification data from the remote computer System (which may be generated by a software running on the remote computer System, a computer expert System, an image analysis system, a human operative in communication with the remote computer system, or Some combination of techniques); and recording the food identification data So as to create a record of foods consumed." • Mault at [0027] • " FIG. 5 illustrates a system allowing diet log creation;" Mault at [0038–39] • "A person can be provided with a portable electronic device, having an image creation capability. For example, the portable electronic device may be a personal digital assistant (PDA) having an optical image sensor. The optical image sensor can be an accessory to the PDA, or can be within the housing of the PDA. The portable electronic device may also be a digital camera, wireless telephone, pager, organizer, calculator (numeric calculator, diet calculator, graphing calculator), spectacle mounted device, wrist mounted device, or other electronic device having image production capabilities. FIG. 1A shows a schematic of an electronic device which may be used in embodiments of the present invention, having a processor **10**, an image capturing system **12**, a data entry mechanism **14**, a memory **16**, a clock **18**,

Claim	'036 Claim Recitation	Exemplary Citations in Reference
		and display **18**." Mault at [0046] • "Other options which may be presented to the user include the entering of a food by name, for example into diet log software, or scanning the barcode of a food item. The electronic device may comprise or communicate with a barcode Scanner, or the optical image Sensor can be used to record an image of the barcode. The barcode can then be interpreted in terms of a UPC (universal product code) and/or product name using a database accessible to the PDA. A UPC database may be accessible over a communications network such as the Internet." Mault at [0048] • FIG. 5 is a schematic of a system comprising an electronic device 100, an imaging system 102, a physiological sensor 104, an activity sensor 106, and a communications link through communications network 108 to a remote computer system (such as a server system) 110. The electronic device can be a PDA or other device carried by the person. The imaging device can be included with the electronic device in a unitary device. The physiological sensor may detect eating, for example a blood glucose sensor can detect the rise in blood glucose after a meal, allowing the electronic device to alert the person to make a diet log entry. The activity sensor provides a signal correlated with a physical activity level of the person. Mault at [0049] • "Images recorded with the imaging device are transmitted over the communications network to the remote computer. The remote computer system can be a server system, and may comprise a computer expert system for image analysis. Image analysis can be performed by one or more persons, a computer expert system, or other mechanism."

Claim	'036 Claim Recitation	Exemplary Citations in Reference
		Mault at [0051] • "Box 124 corresponds to analysis of the image of the remote location. Image recognition software may be used. A trained person, at a remote location, can select food names corresponding to the images, for example in conjunction with image recognition, providing names where image recognition has not been successful in identifying the food. A computer expert System, optical character recognition of a name on the package, barcode recognition, portion size estimate, and other food data can also be extracted from the image or images." Mault at [0054] • "FIG. 9 shows a schematic of a system in which an imaging system **220**, such as a camera, transmits images over a communications network **222** to a remote computer system **224**." Mault at [0073] • "Fig. 11 shows a PDA 280 with display 284, buttons 282, and imaging device 286 directed at a can of pickles 288. The display 284 can be used to indicate the image to be captured. A button such as 282 may be pressed to record the image. The display 184 can be used to view captured images of foods stored in the memory of the PDA, for assisting the person in creating a diet log." Mault at [0080] • "Barcodes can also be used in food item identification. The PDA may also incorporate a bar-code scanner, or alternatively, barcodes may be identified from a recorded image of the bar-code using image analysis. The barcode (e.g. universal product code) can be used to retrieve nutritional data relating to the item from a database. The database may reside on a remote computer

Claim	**'036 Claim Recitation**	**Exemplary Citations in Reference**
		System (remote in this context meaning a computer system not carried by the person, for example a commercial server System, home computer, etc.). The person may scan the barcode of a box of cereal and the UPC code used to retrieve nutrition information specific to that brand of cereal from the database." Mault at [0081] • "FIG. 12 shows an image of a barcode presented on the display 304 of PDA 300 (the PDA having data entry buttons 302, and an image sensor 306). Image analysis software can be used to analyze the image of the barcode, So as to determine UPC, product name, nutritional data, and the like. Alternatively, a barcode Scanner can be used to read the barcode, or an accessory device which in communication with the PDA may be used as a barcode scanner." Mault at Fig. 12

Claim	**'036 Claim Recitation**	**Exemplary Citations in Reference**
		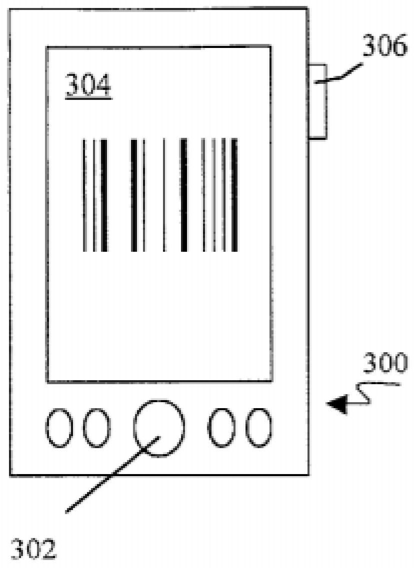 Figure 12 Mault at [0082] - "FIG. 13 shows an image of nutritional information on the display 324 of a PDA320, having data entry buttons 322 and optical image Sensor 326. Optical character recognition Software running on a processor of the PDA can be used to extract nutritional information from the image, and place it into memory for possible later transfer to a database." Mault at Fig. 13

Claim	'036 Claim Recitation	Exemplary Citations in Reference
		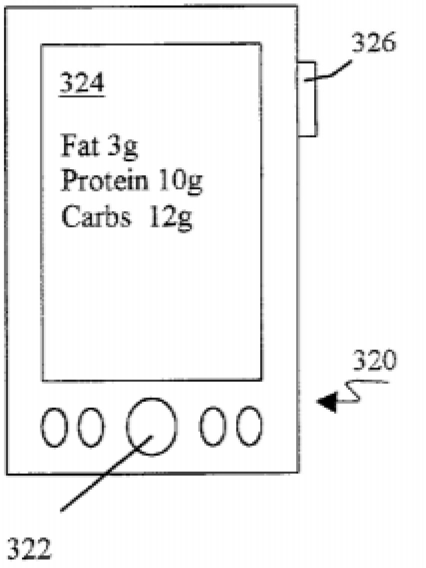 Figure 13 To the extent it is found that Mault does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Mault does not anticipate this claim, Mault renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Mault with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

Claim	'036 Claim Recitation	Exemplary Citations in Reference
[1B.ii]	receives the digital representation from the mobile device; and	Mault discloses receives the digital representation from the mobile device. Mault at [0017] • "A method of creating a diet log comprises: recording an image of a label, the label corresponding to a food item consumed; analyzing the image of the label (for example using optical character recognition, image analysis, image recognition, or neural network techniques) so as to determine the identity of the food item consumed; and storing the determined identity of the food item consumed in an electronic memory, so as to create a diet log. The label can be a printed label, a barcode, an alphanumeric code, a nutritional data label, a list of ingredients, a product code, graphical designs (Such as trademarks), product names, and the like." Mault at [0018] • "A method by which a person can create a record of foods consumed comprises: recording a food image of a food consumed; transmitting the food image to a remote computer system, for example over a communications network Such as the internet, receiving food identification data from the remote computer System (which may be generated by a software running on the remote computer System, a computer expert System, an image analysis system, a human operative in communication with the remote computer system, or Some combination of techniques); and recording the food identification data So as to create a record of foods consumed." Mault at [0048] • "FIG. 5 is a schematic of a system comprising an electronic device 100, an imaging System 102, a physiological Sensor 104, an activity Sensor 106, and a communications link through communications network 108 to a remote computer system (such as a server system) 110."

Claim	'036 Claim Recitation	Exemplary Citations in Reference
		Mault at Fig. 5 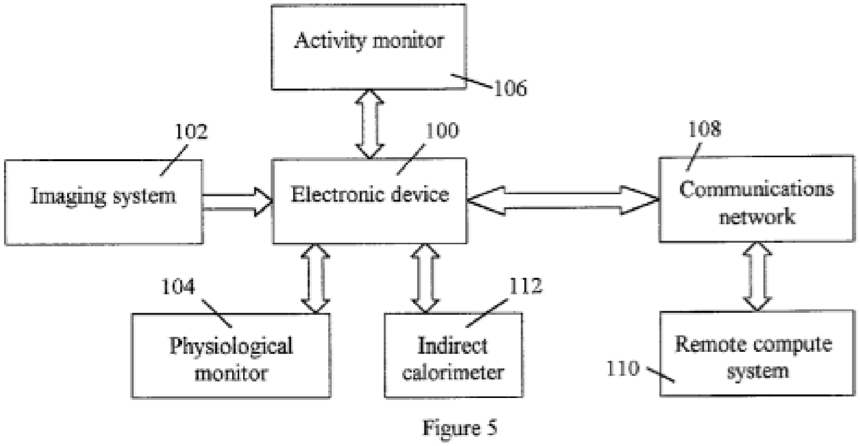 Figure 5 Mault at [0049] • "Images recorded with the imaging device are transmitted over the communications network to the remote computer. The remote computer system can be a server system, and may comprise a computer expert system for image analysis. Image analysis can be performed by one or more persons, a computer expert system, or other mechanism." Mault at [0051] • "FIG. 6 is a flow chart of a method for diet logging. Box 120 corresponds to recording an image of food. Box 122 corresponds to transmission of the image to a remote location, Such as a remote computer System. . . . Box 124 corresponds to analysis of the image of the remote location." Mault at Fig. 6

Claim	'036 Claim Recitation	Exemplary Citations in Reference
		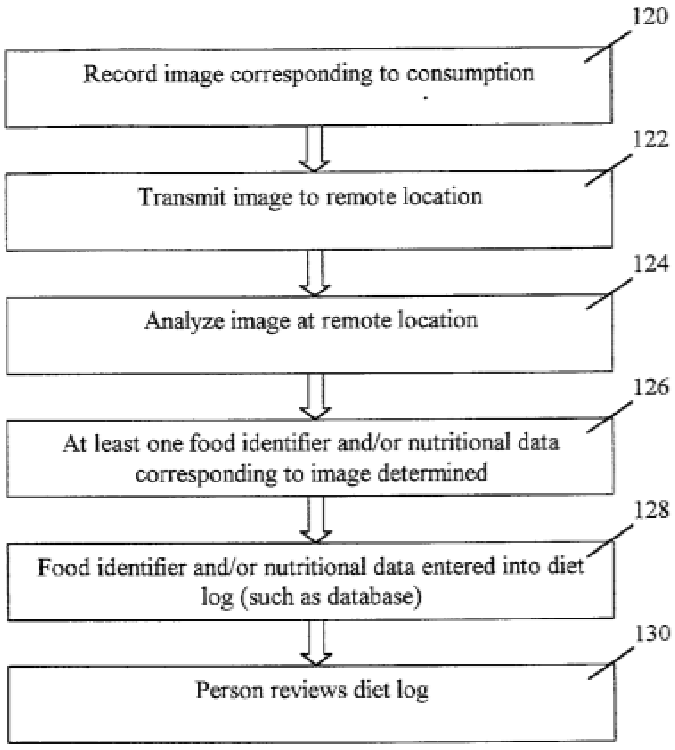 Figure 6 To the extent it is found that Mault does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Mault does not anticipate this claim, Mault renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have

Claim	'036 Claim Recitation	Exemplary Citations in Reference
		been motivated to combine Mault with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results. To the extent it is found that Mault does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Mault does not anticipate this claim, Mault renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Mault with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1B.iii]	recognizes the target as a known target from the target database based on comparing parameters derived from the digital representation to recognition parameters associated with the known targets; and	Mault discloses recognizes the target as a known target from the target database based on comparing parameters derived from the digital representation to recognition parameters associated with the known targets. Mault at [0017] • "A method of creating a diet log comprises: recording an image of a label, the label corresponding to a food item consumed; analyzing the image of the label (for example using optical character recognition, image analysis, image recognition, or neural network techniques) so as to determine the identity of the food item consumed; and storing the determined identity of the food item consumed in an electronic memory, so as to create a diet log. The label can be a printed label, a barcode, an alphanumeric code, a nutritional data label,

Claim	'036 Claim Recitation	Exemplary Citations in Reference
		a list of ingredients, a product code, graphical designs (Such as trademarks), product names, and the like." Mault at [0018]. • "A method by which a person can create a record of foods consumed comprises: recording a food image of a food consumed; transmitting the food image to a remote computer system, for example over a communications network such as the internet, receiving food identification data from the remote computer system (which may be generated by a software running on the remote computer system, a computer expert system, an image analysis system, a human operative in communication with the remote computer system, or some combination of techniques); and recording the food identification data so as to create a record of foods consumed." Mault at [0041] • "FIG. 1A, other components of the electronic device are omitted for clarity. On receiving a signal from the external control or from the processor, the controller **24** initiates the capture of an image by sensor **28**, and may store the image in memory **30** along with a time stamp provided by clock **32**." Mault at [0046] • "Other options which may be presented to the user include the entering of a food by name, for example into diet log software, or scanning the barcode of a food item. The electronic device may comprise or communicate with a barcode scanner, or the optical image sensor can be used to record an image of the barcode. The barcode can then be interpreted in terms of a UPC (universal product code) and/or product name using a database accessible to the PDA. A UPC database may be accessible over a communications network such as the Internet. Other options include receiving transmissions such as wireless data transmissions from food vendors or food vending devices,

Claim	'036 Claim Recitation	Exemplary Citations in Reference
		viewing the recorded diet log, editing the diet log, or analyzing previously recorded images." Mault at [0047] • "The time and location that the image was recorded can also be recorded in memory and displayed, so as to assist diet logging." Mault at [0051] • "Box **124** corresponds to analysis of the image of the remote location. Image recognition software may be used. A trained person, at a remote location, can select food names corresponding to the images, for example in conjunction with image recognition, providing names where image recognition has not been successful in identifying the food. A computer expert system, optical character recognition of a name on the package, barcode recognition, portion size estimate, and other food data can also be extracted from the image or images. Box **126** corresponds to determining food identity and/or nutritional content of the food within the image." Mault at [0075–77] • "Image processing techniques can be used to identify food items from images. For packaged foods, optical character recognition may be used to record and identify nutritional information and/or identify the item consumed. Computer analysis of images may be carried out on the PDA or using another computer system in communication with the PDA. For example, an image of food packaging may be used to identify the food contained. An image of the nutritional content information panel provided by the manufacturer on the package may also be recorded. Optical character recognition may be used to obtain information from the image for storage in a diet log. An image of a box of corn flakes may be recognized by a computer as such, and used to generate a corn flake serving record in a diet log.

Claim	'036 Claim Recitation	Exemplary Citations in Reference
		For non-packaged foods, e.g. restaurant meals, image analysis may also be used in identifying the food items. Imaging at a number of wavelengths, followed by false color image generation, can be used to help identify image components. Spectroscopic imaging can be used in computer-assisted food recognition. Image processing, image recognition, and pattern recognition algorithms are useful in analyzing stored images so as to create a diet log. Algorithms may be applied to color images, or images recorded at a number of different wavelengths (e.g. in the IR, optical, and UV) which may assist identification. Stored images can be identified by comparison with previously stored images that the person has identified, using a software learning mechanism such as a neural network." Mault at [0080] • "Barcodes can also be used in food item identification. The PDA may also incorporate a bar-code Scanner, or alternatively, barcodes may be identified from a recorded image of the bar-code using image analysis. The barcode (e.g. universal product code) can be used to retrieve nutritional data relating to the item from a database. The database may reside on a remote computer System (remote in this context meaning a computer System not carried by the person, for example a commercial Server System, home computer, etc.). The person may scan the barcode of a box of cereal and the UPC code used to retrieve nutrition information specific to that brand of cereal from the database." Mault at [0081] • "FIG. 12 shows an image of a barcode presented on the display **304** of PDA **300** (the PDA having data entry buttons **302**, and an image sensor **306**).

Claim	'036 Claim Recitation	Exemplary Citations in Reference
		Image analysis software can be used to analyze the image of the barcode, so as to determine UPC, product name, nutritional data, and the like." Mault at [0082] - "FIG. 13 shows an image of nutritional information on the display **324** of a PDA **320**, having data entry buttons **322** and optical image Sensor **326**. Optical character recognition Software running on a processor of the PDA can be used to extract nutritional information from the image, and place it into memory for possible later transfer to a database." Mault at [0090] - "Colors, shapes, printed codes, and fluorescent markers of tablets can be used to assist identification, either by a person or an [sic] software image analysis program." Mault at [0091] - "Images can be recorded at intervals by a device carried by the person, and the images used to create a lifestyle log for the person. This can include analysis of images of meals consumed, so as to create a diet log. Images can be recorded automatically, for example at time intervals, so as to spare the person the effort of remembering to record images of foods consumed. Image creation can be triggered by events, such as physiological parameters which change during eating, position (such as using a global positioning system), voice commands, wireless transmission from a remote device, such as a food vending machine, and the like." To the extent it is found that Mault does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Mault does not anticipate this claim, Mault renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with

Claim	**'036 Claim Recitation**	**Exemplary Citations in Reference**
		other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Mault with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1C]	a content service coupled with the identification platform, and that:	Mault discloses a content service coupled with the identification platform, Mault at Fig. 4

Claim	'036 Claim Recitation	Exemplary Citations in Reference
		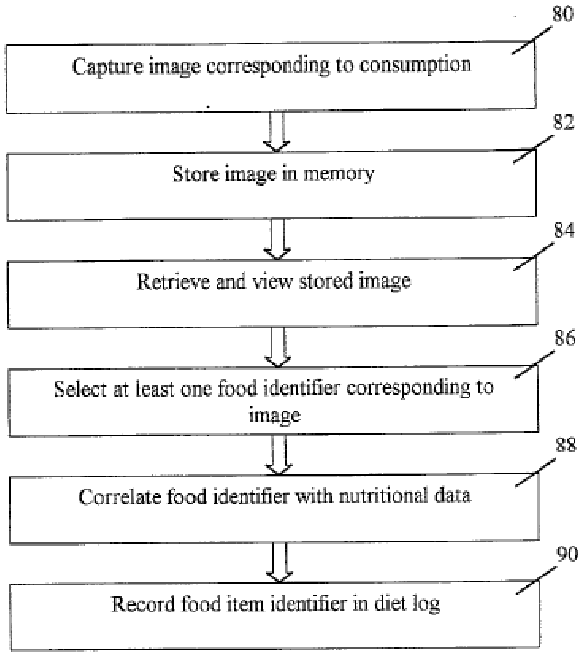 Figure 4 Mault at [0017] • "The identity of food items consumed can be correlated with nutritional data, for example using a nutritional database, and the nutritional data can stored in an electronic memory so as to create a nutritional record, for example as part of a diet log." Mault at [0047] • "Box **88** corresponds to selecting a food item identifier, such as a food name, corresponding to the displayed image. The selection can be made from a menu displayed to the person. Box **88** corresponds to correlating a food item identifier with nutritional data. Box **90** corresponds to recording the food item

Claim	'036 Claim Recitation	Exemplary Citations in Reference
		identifier and corresponding nutritional data in memory, so as to create a diet log." Mault at [0059] • "For example, a person can sit down to a meal and record an image of the meal using the electronic device. The image is transmitted to a remote location where it is analyzed and converted into a caloric value. The caloric value can be transmitted back to the person, allowing the person to decide whether to eat the meal, choose another meal, or eat some portion of the meal consistent with reaching a calorie balance goal." Mault at [0082–83] • "FIG. 13 shows an image of nutritional information on the display **324** of a PDA **320**, having data entry buttons **322** and optical image sensor **326**. Optical character recognition software running on a processor of the PDA can be used to extract nutritional information from the image, and place it into memory for possible later transfer to a database. Food item identities and nutritional information can also be received from food vendors, e.g. grocery stores, on-line retailers, restaurants, vending machines, diet food retailers, etc. The PDA may prompt the person to identify the source of the food being imaged, and this information can be used to help in identification." To the extent it is found that Mault does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Mault does not anticipate this claim, Mault renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have

Claim	'036 Claim Recitation	Exemplary Citations in Reference
		been motivated to combine Mault with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1C.i]	obtains content information related to the known target; and	Mault discloses obtains content information related to the known target. *See* [1C] above To the extent it is found that Mault does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Mault does not anticipate this claim, Mault renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Mault with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1C.ii]	sends the content information to at least one of the identification platform and the mobile device.	Mault discloses sends the content information to at least one of the identification platform and the mobile device. Mault at [0052] • "A diet log channel, generated by remote server (or other remote computer systems, or other content provider) **176**, and received by the set top box over communications network **178**, can be selected for display on the interactive

Claim	'036 Claim Recitation	Exemplary Citations in Reference
		television **160**. A display of food item identifiers, such as food names in a menu format, can be displayed on the display **162** of the interactive TV **160**. An image stored in the local memory **172** can be added to the diet log channel signal, using techniques known in the art, so that the person views the diet log channel with an embedded image of the food consumed. The person then selects items corresponding to the food consumed from the displayed menu. Selection data is generated with user input **146** (for example a keyboard) and transmitted to the set top box and to the remote server. The display is changed appropriately on receipt of the selection data by the content provider and the selection is added to the diet log of the person, which is stored in a database associated with or otherwise accessible by the remote server **176**." Mault at [0051] • "FIG. 6 is a flow chart of a method for diet logging. Box **120** corresponds to recording an image of food. Box **122** corresponds to transmission of the image to a remote location, such as a remote computer system. Other data can be transmitted with the image, for example a sales receipt, data provided by the food vendor, location the image was taken (which may be correlated with a restaurant identity), a menu image, and the weight of the food (which may be provided by a mat). Box **124** corresponds to analysis of the image of the remote location. Image recognition software may be used. A trained person, at a remote location, can select food names corresponding to the images, for example in conjunction with image recognition, providing names where image recognition has not been successful in identifying the food. A computer expert system, optical character recognition of a name on the package, barcode recognition, portion size estimate, and other food data can also be extracted from the image or images. Box **126** corresponds to determining food identity and/or nutritional content of the food within the image. Box **128** corresponds to storing the food identity and nutritional content into a database, so as to create a diet log. Box **130** corresponds to the person reviewing the diet log for accuracy at a convenient later time."

Claim	'036 Claim Recitation	Exemplary Citations in Reference
		Mault at [0059] • "For example, a person can sit down to a meal and record an image of the meal using the electronic device. The image is transmitted to a remote location where it is analyzed and converted into a caloric value. The caloric value can be transmitted back to the person, allowing the person to decide whether to eat the meal, choose another meal, or eat some portion of the meal consistent with reaching a calorie balance goal." Mault at [0074] • "FIG. 11 shows a PDA **280** with display **284**, buttons **282**, and imaging device **286** directed at a can of pickles **288**. The display **284** can be used to indicate the image to be captured. A button such as **282** may be pressed to record the image. The display **184** can be used to view captured images of foods stored in the memory of the PDA, for assisting the person in creating a diet log." Mault at [0075] • "Optical character recognition may be used to obtain information from the image for storage in a diet log. An image of a box of corn flakes may be recognized by a computer as such, and used to generate a corn flake serving record in a diet log. Computer analysis may be used as a first attempt in producing a diet log, and if unsuccessful, the image can be passed to a human for analysis." To the extent it is found that Mault does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Mault does not anticipate this claim, Mault renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's

Claim	'036 Claim Recitation	Exemplary Citations in Reference
		Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Mault with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

Exhibit C-5

Bank of America Corporation ("BofA" or "Defendant") provides this chart subject to all reservations, objections, statements, and disclaimers set forth herein and in BofA's Final Invalidity Contentions Cover Pleading, as well as any amendment, supplement, or modification thereof, which are incorporated herein by reference in their entirety.

On information and belief, the application that matured into Sizer was filed on March 28, 1997. Sizer published at least as of March 14, 2000, and is available as prior art at least under 35 U.S.C. § 102 (a), (b) and (e).

As shown in the chart below, Sizer anticipates asserted claim 1 of the '036 Patent. To the extent Sizer is found not to expressly disclose certain limitations in the asserted claims, such limitations are inherent. To the extent Sizer is found not to anticipate any asserted claims or claim elements of the '036 Patent, the reference nevertheless renders those claims or claim elements obvious under 35 U.S.C. § 103, either alone or in combination with other art identified herein or in the cover pleading of BofA's Final Invalidity Contentions, for the reasons identified in the cover pleading. These Final Invalidity Contentions are not an admission by BofA that the accused products, including any current or past versions of the accused products, are covered by, or infringe the asserted claims, but are based instead on the recognition that if the claims are interpreted to be broad enough to encompass the accused products, the claims must also be construed to have that same scope when considering whether they are invalid.

The following chart is based on the Court's claim construction rulings (Dkt. Nos. 145, 153, 236), and to the extent not contradictory, Plaintiffs' apparent interpretations of the claims in Plaintiffs' "Preliminary" Final Infringement Contentions. BofA notes that in many instances, Plaintiffs' "Preliminary" Final Infringement Contentions ignore and/or contradict the Court's claim construction rulings. BofA does not accept the assumptions concerning the scope and meaning implicit in Plaintiffs' Preliminary Final Infringement Contentions, to the extent those assumptions are discernible, and reserves the right to challenge Plaintiffs' proposed (or implied) constructions. Furthermore, Plaintiffs' "Preliminary" Final Contentions fail to comply with the level of specificity required by SPR 2.1.2 and 2.1.3, and include improper reservation of rights to supplement with "additional ways in which BoA Accused Products infringe." Any alleged deficient discovery is of Plaintiffs own making. Most if not all of the deficiencies in Plaintiffs' Preliminary Final Contentions have been documented by Defendants through multiple letters since at least May 18, 2021 and up to Bank of America's most recent letter of August 19, 2023. Plaintiffs' shifting infringement contentions, deficiencies in its infringement contentions, and changing claim scope have prejudiced BofA from being able to reasonably prepare its defenses, including its invalidity contentions. BofA also reserves the right to revise and supplement these charts if and when Plaintiffs are permitted to provide revised or supplemental Infringement Contentions. Citations given in the chart below are merely representative of the respective elements to provide notice to Plaintiffs of BofA's invalidity theories and are not meant to be exhaustive.

1. Claim 1

Element	’036 Claim Recitation	Exemplary Citations in Reference
[1Pre]	A content provisioning system comprising:	Sizer discloses a system. Sizer at 12:62-13:13 • In accordance with the present invention, a capture device now can initiate a transaction by capturing transaction data from marks contained on an object, such as a magazine advertisement, where the marks have embedded therein a code corresponding to transaction data for initiating a transaction. The capture device includes a scanner, operable by a user, for reading marks contained on the object and a controller for interpreting the marks and retrieving transaction data embedded in the marks. Through associated apparatus as described above, a telephone call can be originated by information contained in the transaction data in at least a portion of the transaction data transferred to a desired destination for initiating a transaction. As a result, it is now possible now for magazines, newspapers, printed media and many other objects to contain marks relating to an advertisement or solicitation which can be easily scanned and a transaction completed to the vendor of the advertisement or solicitation. This saves much time and allows easier access to services and products.
[1A]	a target database storing known targets of different types and recognition parameters associated with the known targets;	Sizer discloses a target database storing known targets of barcodes. Sizer at FIG. 1

Element	'036 Claim Recitation	Exemplary Citations in Reference
		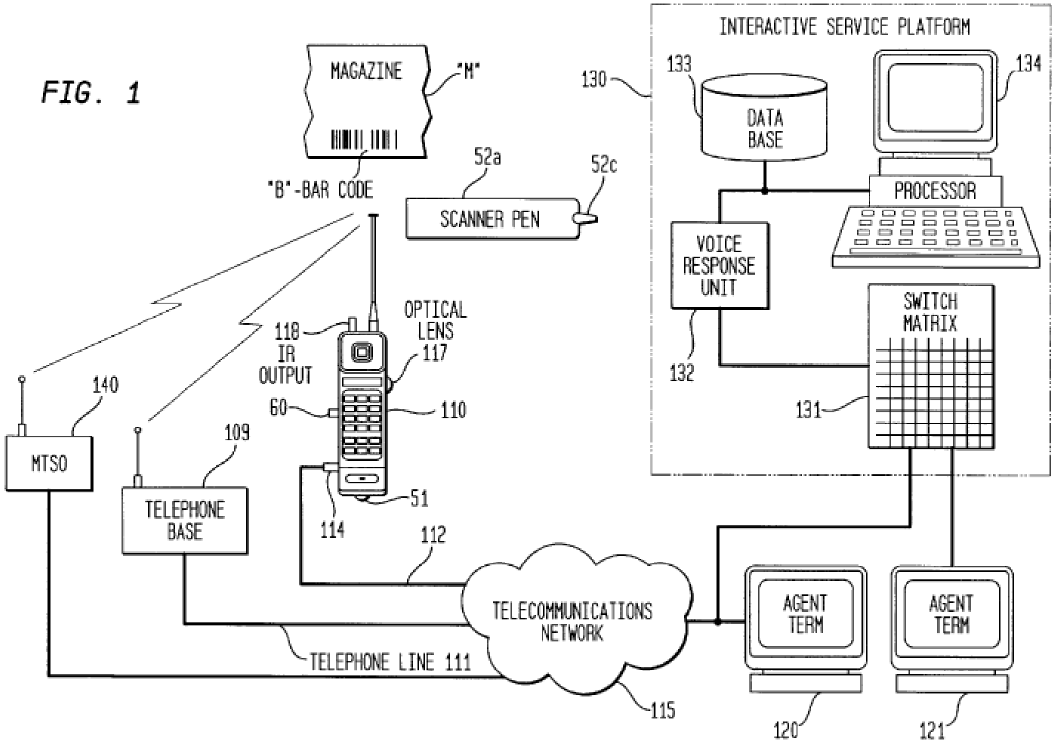 FIG. 1 Sizer at 6:18-38 • "If an outgoing telephone call in initiated by capture device 110, the call can illustratively be routed through telecommunications network 115 directly to an agent terminal 120, or to an interactive service platform 130, which is connected to another agent terminal 121. When a connection is completed to platform 130, the call may be routed by a switch matrix 131 to a voice response unit 132 that operates under the control of a processor

Element	'036 Claim Recitation	Exemplary Citations in Reference
		134. Information included in the captured encoded can be provided to the attendants, or processor 134 can perform a look-up and retrieval operation in database 131, using encoded data as a key, in order to effectuate a desired transaction. For certain transactions, information connected with a transaction may also be transmitted from service platform 130 to capture device 110, either for display to a user, storage for later use, or output to a point of sale system or device. This information may, for example, be confirmation information, indicating that a transaction has been effected, or details regarding a just completed transaction, such as ticket information relating to a travel reservation." Sizer at 10:45-49 • "A call is then originated launch, such by transmitting the dialing information from memory 221 through base station interface 250 to base station 109 and thence to telephone line 111 and telecommunications network 115 to interactive service platform 130." Sizer at 12:41-43 • "The call may be placed to an interactive service platform, and additional captured data, such as coupon or similar information, can then be used to effectuate a transaction." To the extent it is found that Sizer does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Sizer does not anticipate this claim, Sizer renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Sizer with any of the other prior art references identified in

Element	'036 Claim Recitation	Exemplary Citations in Reference
		the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1B]	an identification platform coupled with the target database, and that	Sizer discloses an identification platform coupled with the target database. Sizer at FIG. 2 FIG. 2 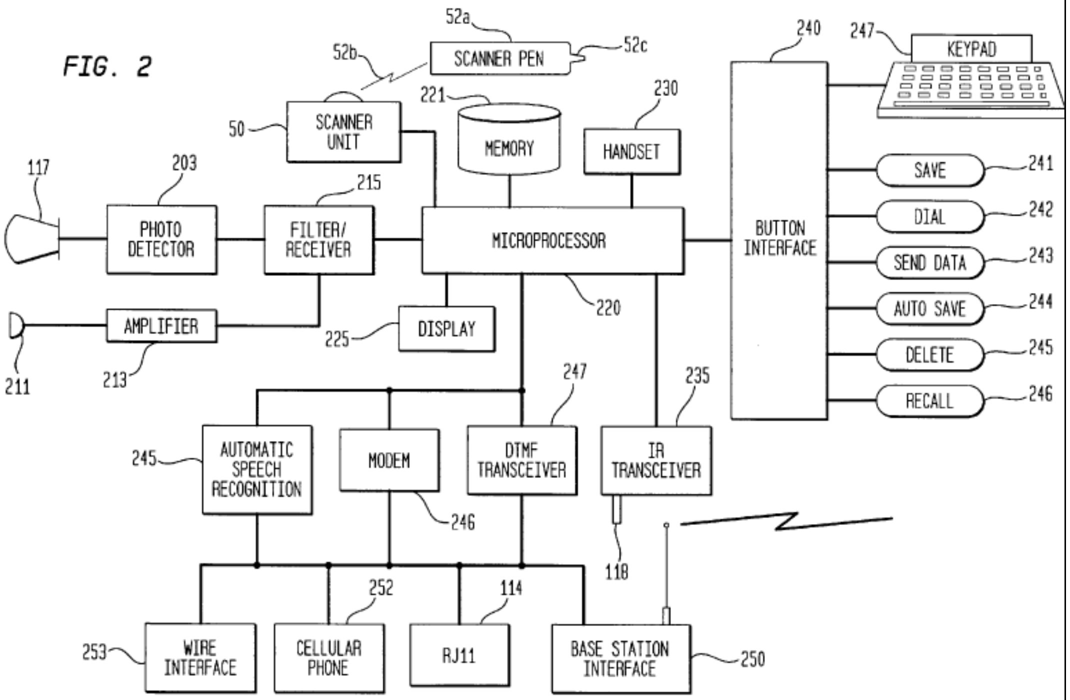 Sizer at 5:3-20 • "Briefly, it was found using a spectrum analyzer on a typical video signal,

Element	'036 Claim Recitation	Exemplary Citations in Reference
		that there are comparatively large frequency components at the line rate and at the frame rate and its harmonics, but that between these frequencies, there are other frequency bands in which little information is carried. One such frequency band is between 15 and 30 kHz. By adding a low level carrier signal or tone at a frequency in this band, say 25 kHz, the video image is not degraded, but a properly tuned decoder can receive and decode the encoded information. In this way, digital information can be subliminally inserted in a video signal by adding to the video signal an amplitude shift keyed (ASK) or frequency shift keyed (FSK) carrier signal, and the digital information can later be recovered using ASK or FSK decoding. Encoded data can also be inserted in a television signal and recovered by a received that responds to the picture displayed on the television, in the manner described in U.S. Pat. No. 4,807,031 cited above." Sizer at 6:41-62 • "If the encoded non-perceptible data is inserted in the video portion of a television signal, the images displayed on a visual display area a monitor or television are captured by optical lens 117 and an associated photo-detector 203, which is arranged to supply an electrical signal to a filter/receiver 215 representing the image. On the other hand, if the encoded non-perceptible data is inserted in the audio portion of a television signal, e.g. the output from a speaker is captured by a microphone 211 and an associated amplifier 213, which is arranged to supply an electrical signal to filter/receiver 215 representing the sound energy. In either event, the encoded non-perceptible data is decoded in the filter/receiver 215 in a manner consistent with the manner in which the original data was encoded. Thus, filter/receiver 215 can be arranged to perform the same functionality as the elements illustrated in FIG. 4 of the above referenced Broughton et al. patent, or, alternatively, the operation of filter/receiver 215 can be as described in conjunction with FIGS. 4 or 7 in the copending application of T. Sizer."

Element	**'036 Claim Recitation**	**Exemplary Citations in Reference**
		Sizer at 7:1-6 • "The output of filter/receiver 215, representing the data captured by capture device 110, is applied to a microprocessor 220, which is arranged to perform various data processing and control functions in coordination with programs stored in an associated memory 221 as well as inputs received from a user of capture device 110." To the extent it is found that Sizer does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Sizer does not anticipate this claim, Sizer renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Sizer with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1B.i]	communicates with a mobile device capable of acquiring a digital representation of a scene containing at least a portion of a target;	Sizer discloses an identification platform coupled with the target database that communicates with a mobile device capable of acquiring a digital representation of a scene containing at least a portion of a target. Sizer at FIG. 2

Element	'036 Claim Recitation	Exemplary Citations in Reference
		FIG. 2 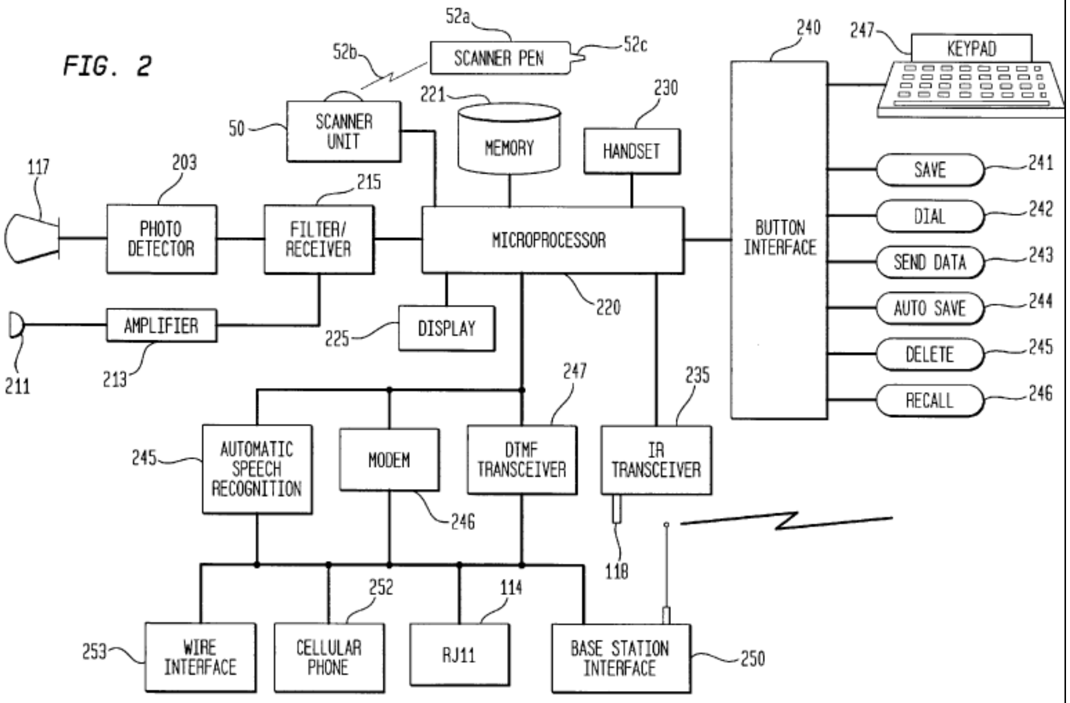 Sizer at 3:34-38 • The scanner 50 unit includes a scanning head 51 positioned on the bottom of the wireless capture device 110. In the particular embodiment illustrated with the scanner head 51 positioned on the bottom of the wireless capture device 110, the device is manually oriented in use so that the scanner head 51 moves over the marks to be read. The scanner unit 50 is wired directly to a microprocessor 220 (FIG. 2) of the device 110 and works in a manner that will be explained below.

Element	'036 Claim Recitation	Exemplary Citations in Reference
		Sizer at 6:41-62 • "If the encoded non-perceptible data is inserted in the video portion of a television signal, the images displayed on a visual display area a monitor or television are captured by optical lens 117 and an associated photo-detector 203, which is arranged to supply an electrical signal to a filter/receiver 215 representing the image. On the other hand, if the encoded non-perceptible data is inserted in the audio portion of a television signal, e.g. the output from a speaker is captured by a microphone 211 and an associated amplifier 213, which is arranged to supply an electrical signal to filter/receiver 215 representing the sound energy. In either event, the encoded non-perceptible data is decoded in the filter/receiver 215 in a manner consistent with the manner in which the original data was encoded. Thus, filter/receiver 215 can be arranged to perform the same functionality as the elements illustrated in FIG. 4 of the above referenced Broughton et al. patent, or, alternatively, the operation of filter/receiver 215 can be as described in conjunction with FIGS. 4 or 7 in the copending application of T. Sizer." Sizer at 12:34-35 • It is further to be noted that the present invention can be used in the context of automobiles and mobile telephones. To the extent it is found that Sizer does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Sizer does not anticipate this claim, Sizer renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Sizer with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the

Element	'036 Claim Recitation	Exemplary Citations in Reference
		reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1B.ii]	receives the digital representation from the mobile device; and	Sizer discloses an identification platform coupled with the target database that receives the digital representation from the mobile device. Sizer at FIG. 2 FIG. 2 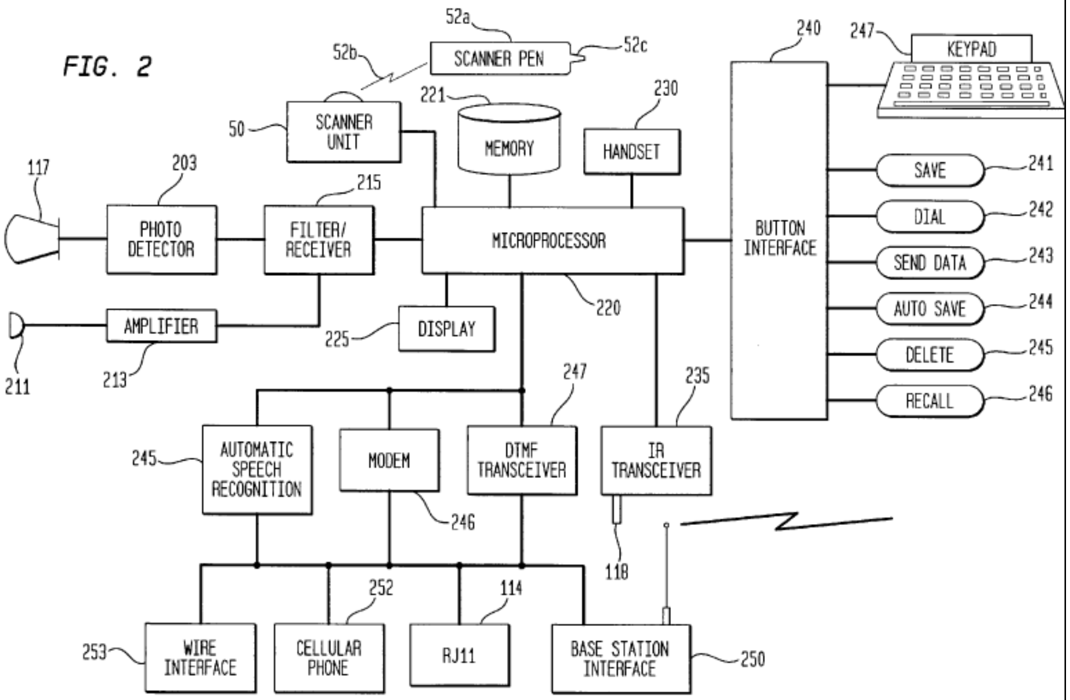 Sizer at 6:41-62 • "If the encoded non-perceptible data is inserted in the video portion of a

Element	'036 Claim Recitation	Exemplary Citations in Reference
		television signal, the images displayed on a visual display area a monitor or television are captured by optical lens 117 and an associated photo-detector 203, which is arranged to supply an electrical signal to a filter/receiver 215 representing the image. On the other hand, if the encoded non-perceptible data is inserted in the audio portion of a television signal, e.g. the output from a speaker is captured by a microphone 211 and an associated amplifier 213, which is arranged to supply an electrical signal to filter/receiver 215 representing the sound energy. In either event, the encoded non-perceptible data is decoded in the filter/receiver 215 in a manner consistent with the manner in which the original data was encoded. Thus, filter/receiver 215 can be arranged to perform the same functionality as the elements illustrated in FIG. 4 of the above referenced Broughton et al. patent, or, alternatively, the operation of filter/receiver 215 can be as described in conjunction with FIGS. 4 or 7 in the copending application of T. Sizer." To the extent it is found that Sizer does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Sizer does not anticipate this claim, Sizer renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Sizer with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1B.iii]	recognizes the target as a known target from the target database based on comparing parameters	Sizer discloses recognizes the target as a known target from the target database based on comparing parameters derived from the digital representation to recognition parameters associated with the known targets.

Element	'036 Claim Recitation	Exemplary Citations in Reference
	derived from the digital representation to recognition parameters associated with the known targets; and	Sizer at FIG. 2 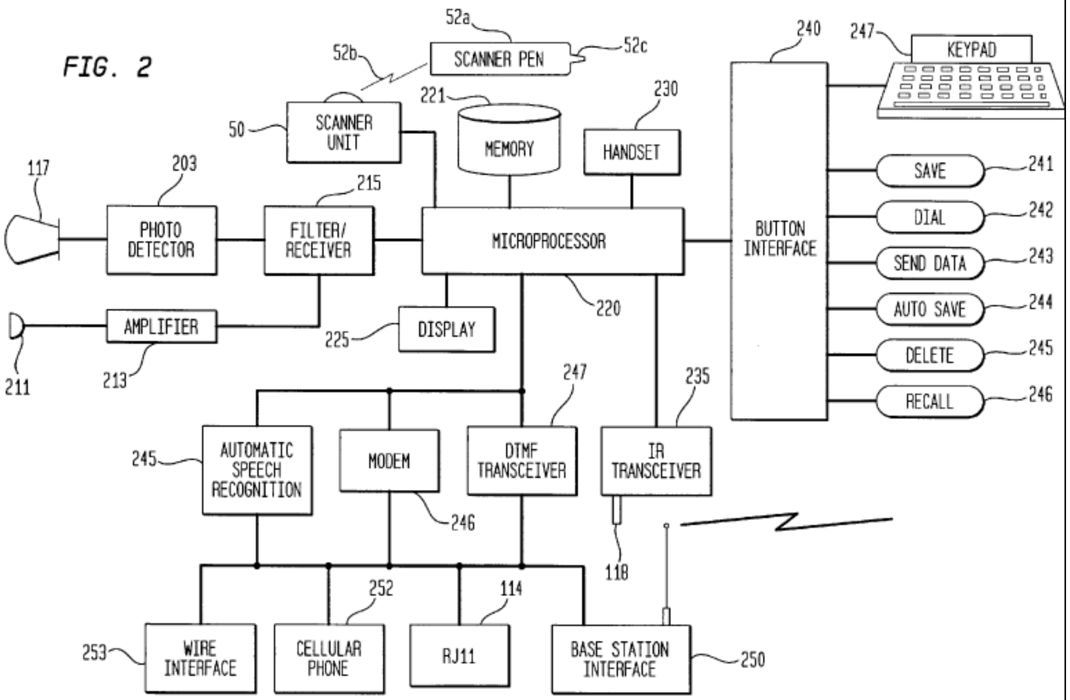 FIG. 2 Sizer at 6:41-62 • "If the encoded non-perceptible data is inserted in the video portion of a television signal, the images displayed on a visual display area a monitor or television are captured by optical lens 117 and an associated photo-detector 203, which is arranged to supply an electrical signal to a filter/receiver 215 representing the image. On the other hand, if the encoded non-perceptible data is inserted in the audio portion of a television signal, e.g. the output from a speaker is

Element	'036 Claim Recitation	Exemplary Citations in Reference
		captured by a microphone 211 and an associated amplifier 213, which is arranged to supply an electrical signal to filter/receiver 215 representing the sound energy. In either event, the encoded non-perceptible data is decoded in the filter/receiver 215 in a manner consistent with the manner in which the original data was encoded. Thus, filter/receiver 215 can be arranged to perform the same functionality as the elements illustrated in FIG. 4 of the above referenced Broughton et al. patent, or, alternatively, the operation of filter/receiver 215 can be as described in conjunction with FIGS. 4 or 7 in the copending application of T. Sizer." To the extent it is found that Sizer does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Sizer does not anticipate this claim, Sizer renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Sizer with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1C]	a content service coupled with the identification platform, and that:	Sizer discloses a content service coupled with the identification platform. Sizer at FIG. 1

Element	’036 Claim Recitation	Exemplary Citations in Reference
		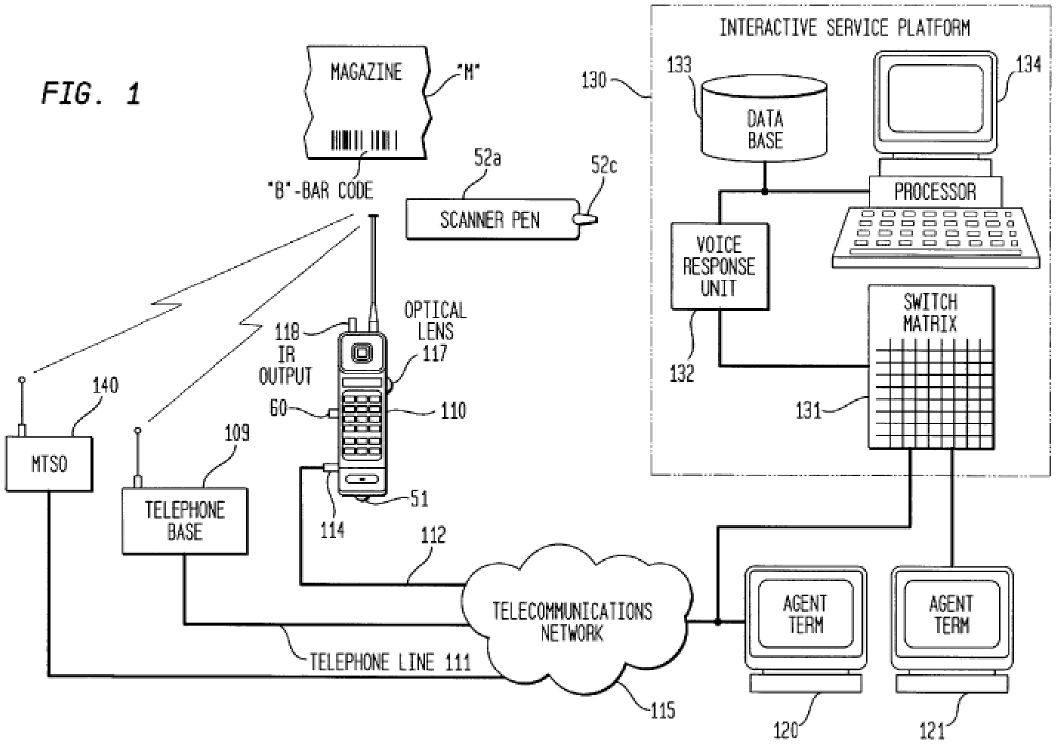 Sizer at 6:18-38 • “If an outgoing telephone call in initiated by capture device 110, the call can illustratively be routed through telecommunications network 115 directly to an agent terminal 120, or to an interactive service platform 130, which is connected to another agent terminal 121. When a connection is completed to platform 130, the call may be routed by a switch matrix 131 to a voice response unit 132 that operates under the control of a processor

Element	'036 Claim Recitation	Exemplary Citations in Reference
		134. Information included in the captured encoded can be provided to the attendants, or processor 134 can perform a look-up and retrieval operation in database 131, using encoded data as a key, in order to effectuate a desired transaction. For certain transactions, information connected with a transaction may also be transmitted from service platform 130 to capture device 110, either for display to a user, storage for later use, or output to a point of sale system or device. This information may, for example, be confirmation information, indicating that a transaction has been effected, or details regarding a just completed transaction, such as ticket information relating to a travel reservation." Sizer at 10:45-49 • "A call is then originated launch, such by transmitting the dialing information from memory 221 through base station interface 250 to base station 109 and thence to telephone line 111 and telecommunications network 115 to interactive service platform 130." Sizer at 12:41-43 • "The call may be placed to an interactive service platform, and additional captured data, such as coupon or similar information, can then be used to effectuate a transaction." To the extent it is found that Sizer does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Sizer does not anticipate this claim, Sizer renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Sizer with any of the other prior art references identified in

Element	**'036 Claim Recitation**	**Exemplary Citations in Reference**
		the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1C.i]	obtains content information related to the known target; and	Sizer discloses [a content service coupled with the identification platform] that obtains content information related to the known target. Sizer at FIG. 1

Element	'036 Claim Recitation	Exemplary Citations in Reference
		FIG. 1 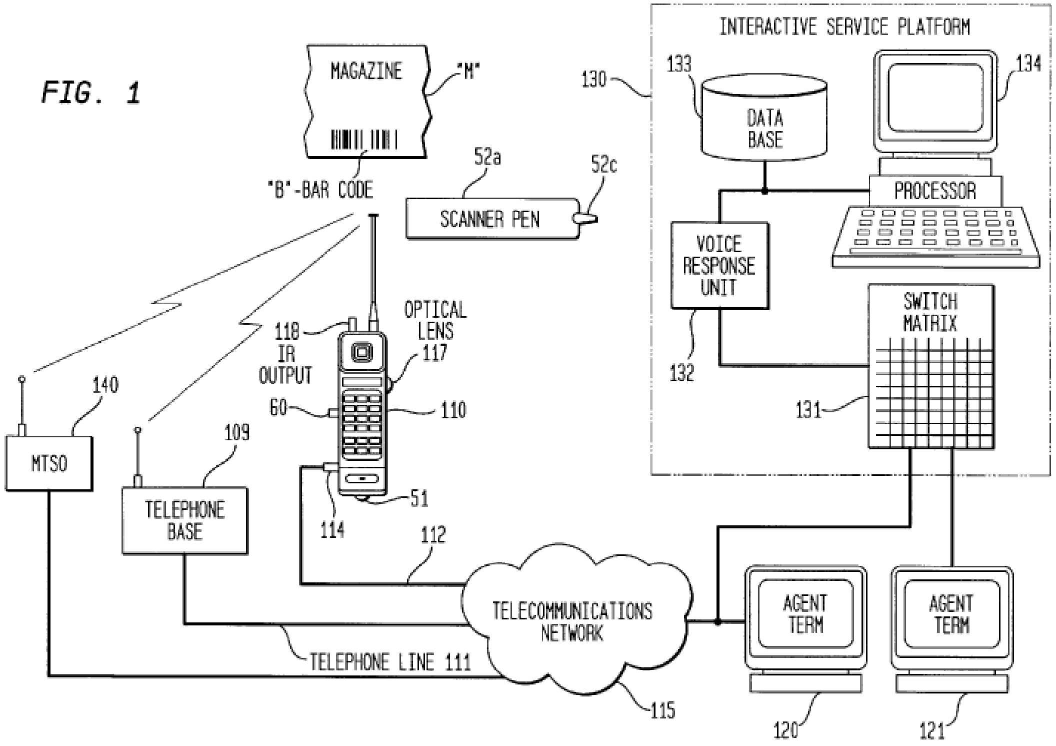 Sizer at 6:18-38 • "If an outgoing telephone call in initiated by capture device 110, the call can illustratively be routed through telecommunications network 115 directly to an agent terminal 120, or to an interactive service platform 130, which is connected to another agent terminal 121. When a connection is completed to platform 130, the call may be routed by a switch matrix 131 to a voice response unit 132 that operates under the control of a processor

Element	'036 Claim Recitation	Exemplary Citations in Reference
		134. Information included in the captured encoded can be provided to the attendants, or processor 134 can perform a look-up and retrieval operation in database 131, using encoded data as a key, in order to effectuate a desired transaction. For certain transactions, information connected with a transaction may also be transmitted from service platform 130 to capture device 110, either for display to a user, storage for later use, or output to a point of sale system or device. This information may, for example, be confirmation information, indicating that a transaction has been effected, or details regarding a just completed transaction, such as ticket information relating to a travel reservation." Sizer at 10:45-49 • "A call is then originated launch, such by transmitting the dialing information from memory 221 through base station interface 250 to base station 109 and thence to telephone line 111 and telecommunications network 115 to interactive service platform 130." Sizer at 12:41-43 • "The call may be placed to an interactive service platform, and additional captured data, such as coupon or similar information, can then be used to effectuate a transaction." To the extent it is found that Sizer does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Sizer does not anticipate this claim, Sizer renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Sizer with any of the other prior art references identified in

Element	**'036 Claim Recitation**	**Exemplary Citations in Reference**
		the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1C.ii]	sends the content information to at least one of the identification platform and the mobile device.	Sizer discloses [a content service coupled with the identification platform] that sends the content information to at least one of the identification platform and the mobile device. Sizer at FIG. 1

Element	'036 Claim Recitation	Exemplary Citations in Reference
		FIG. 1 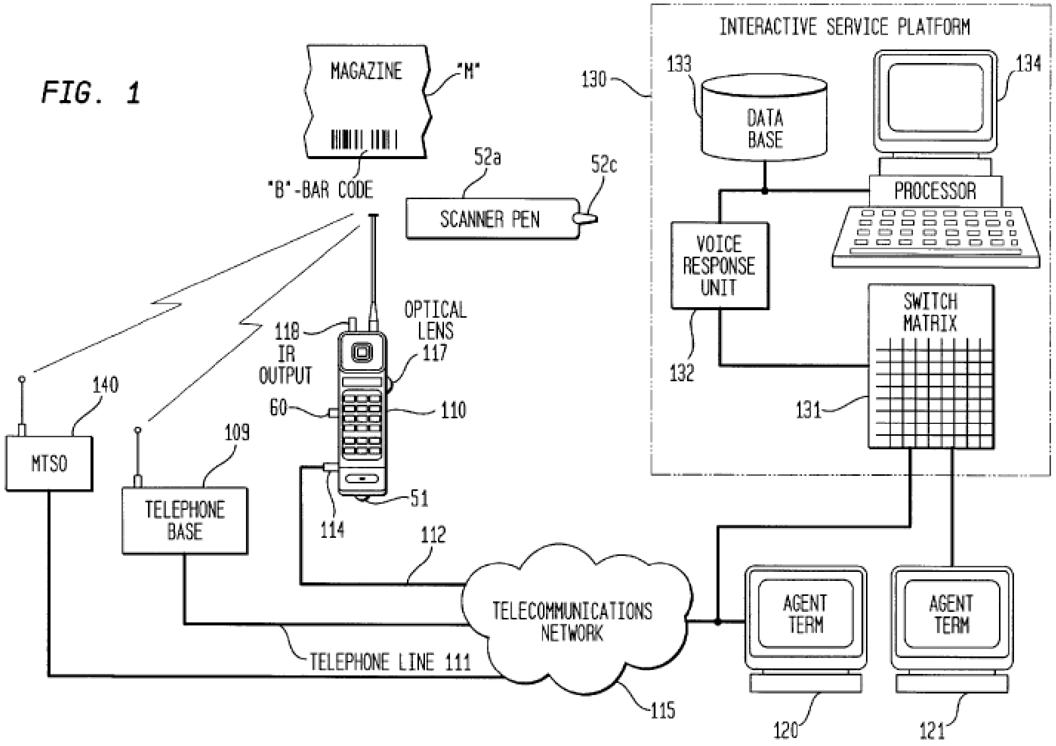 Sizer at 6:18-38 • "If an outgoing telephone call in initiated by capture device 110, the call can illustratively be routed through telecommunications network 115 directly to an agent terminal 120, or to an interactive service platform 130, which is connected to another agent terminal 121. When a connection is completed to platform 130, the call may be routed by a switch matrix 131 to a voice response unit 132 that operates under the control of a processor

Element	'036 Claim Recitation	Exemplary Citations in Reference
		134. Information included in the captured encoded can be provided to the attendants, or processor 134 can perform a look-up and retrieval operation in database 131, using encoded data as a key, in order to effectuate a desired transaction. For certain transactions, information connected with a transaction may also be transmitted from service platform 130 to capture device 110, either for display to a user, storage for later use, or output to a point of sale system or device. This information may, for example, be confirmation information, indicating that a transaction has been effected, or details regarding a just completed transaction, such as ticket information relating to a travel reservation." Sizer at 10:45-49 • "A call is then originated launch, such by transmitting the dialing information from memory 221 through base station interface 250 to base station 109 and thence to telephone line 111 and telecommunications network 115 to interactive service platform 130." Sizer at 12:41-43 • "The call may be placed to an interactive service platform, and additional captured data, such as coupon or similar information, can then be used to effectuate a transaction." To the extent it is found that Sizer does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Sizer does not anticipate this claim, Sizer renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Sizer with any of the other prior art references identified in

Element	'036 Claim Recitation	Exemplary Citations in Reference
		the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

Nantworks, LLC v. Bank of America Corporation
Central District of California, No. 2:20-cv-7872
Invalidity Chart for U.S. Patent 8,478,036 based on
U.S. Patent No. 6,666,377 to Harris ("Harris")

Exhibit C-6

Bank of America Corporation ("BofA" or "Defendant") provides this chart subject to all reservations, objections, statements, and disclaimers set forth herein and in BofA's Final Invalidity Contentions Cover Pleading, as well as any amendment, supplement, or modification thereof, which are incorporated herein by reference in their entirety.

On information and belief, the application that matured into Harris was filed on July 18, 2000. Harris published at least as of December 23, 2003, and is available as prior art at least under 35 U.S.C. § 102(e).

As shown in the chart below, Harris anticipates asserted claim 1 of the '036 Patent. To the extent Harris is found not to expressly disclose certain limitations in the asserted claims, such limitations are inherent. To the extent Harris is found not to anticipate any asserted claims or claim elements of the '036 Patent, the reference nevertheless renders those claims or claim elements obvious under 35 U.S.C. § 103, either alone or in combination with other art identified herein or in the cover pleading of BofA's Final Invalidity Contentions, for the reasons identified in the cover pleading. These Final Invalidity Contentions are not an admission by BofA that the accused products, including any current or past versions of the accused products, are covered by, or infringe the asserted claims, but are based instead on the recognition that if the claims are interpreted to be broad enough to encompass the accused products, the claims must also be construed to have that same scope when considering whether they are invalid.

The following chart is based on the Court's claim construction rulings (Dkt. Nos. 145, 153, 236), and to the extent not contradictory, Plaintiffs' apparent interpretations of the claims in Plaintiffs' "Preliminary" Final Infringement Contentions. BofA notes that in many instances, Plaintiffs' "Preliminary" Final Infringement Contentions ignore and/or contradict the Court's claim construction rulings. BofA does not accept the assumptions concerning the scope and meaning implicit in Plaintiffs' Preliminary Final Infringement Contentions, to the extent those assumptions are discernible, and reserves the right to challenge Plaintiffs' proposed (or implied) constructions. Furthermore, Plaintiffs' "Preliminary" Final Contentions fail to comply with the level of specificity required by SPR 2.1.2 and 2.1.3, and include improper reservation of rights to supplement with "additional ways in which BoA Accused Products infringe." Any alleged deficient discovery is of Plaintiffs own making. Most if not all of the deficiencies in Plaintiffs' Preliminary Final Contentions have been documented by Defendants through multiple letters since at least May 18, 2021 and up to Bank of America's most recent letter of August 19, 2023. Plaintiffs' shifting infringement contentions, deficiencies in its infringement contentions, and changing claim scope have prejudiced BofA from being able to reasonably prepare its defenses, including its invalidity contentions. BofA also reserves the right to revise and supplement these charts if and when Plaintiffs are permitted to provide revised or supplemental Infringement Contentions. Citations given in the chart below are merely representative of the respective elements to provide notice to Plaintiffs of BofA's invalidity theories and are not meant to be exhaustive.

1. Claim 1

Element	'036 Claim Recitation	Exemplary Citations in Reference
[1Pre]	A content provisioning system comprising:	Harris discloses a system. Harris at 4:41-48 "In this system, the information can represent a cue, e.g. a number, that is associated with a special function. The cue is used during a hot sync to access a publicly available information network, e.g. by direct connection or by the Internet. The cue 431 is sent over the Internet 432 to the server 450, and addresses more detailed information in a memory of the server. The server returns that information as 455, and the client receives the more detailed information."
[1A]	a target database storing known targets of different types and recognition parameters associated with the known targets;	Harris discloses a target database storing known targets of barcodes. Harris at FIG. 4

Element	'036 Claim Recitation	Exemplary Citations in Reference
		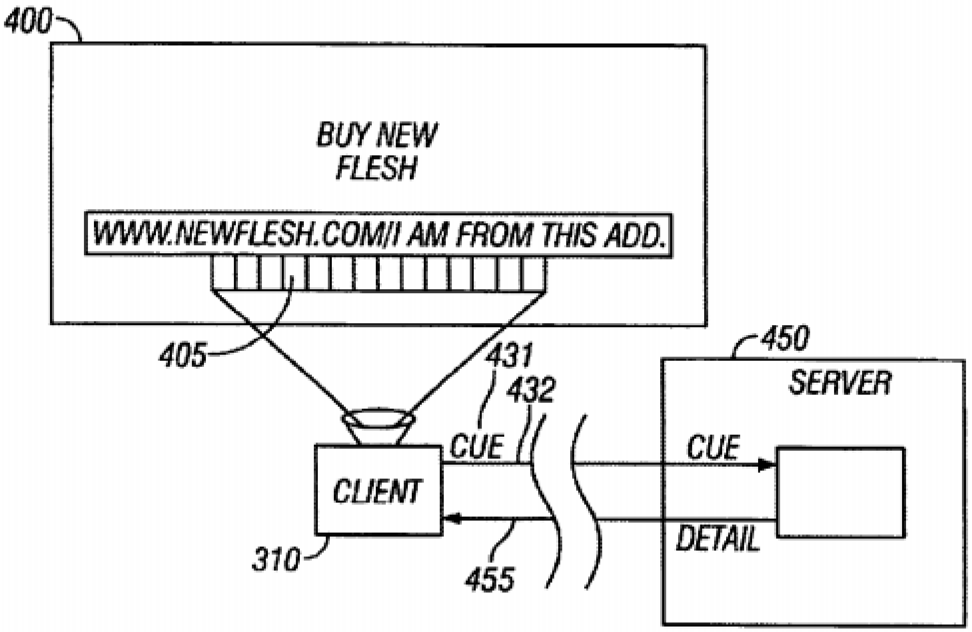 FIG. 4 Harris at 4:41-48 • "In this system, the information can represent a cue, e.g. a number, that is associated with a special function. The cue is used during a hot sync to access a publicly available information network, e.g. by direct connection or by the Internet. The cue 431 is sent over the Internet 432 to the server 450, and addresses more detailed information in a memory of the server. The server returns that information as 455, and the client receives the more detailed information." To the extent it is found that Harris does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Harris does not anticipate

Element	'036 Claim Recitation	Exemplary Citations in Reference
		this claim, Harris renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Harris with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1B]	an identification platform coupled with the target database, and that	Harris discloses an identification platform coupled with the target database. Harris at FIG. 1A and FIG. 1B

Element	’036 Claim Recitation	Exemplary Citations in Reference
		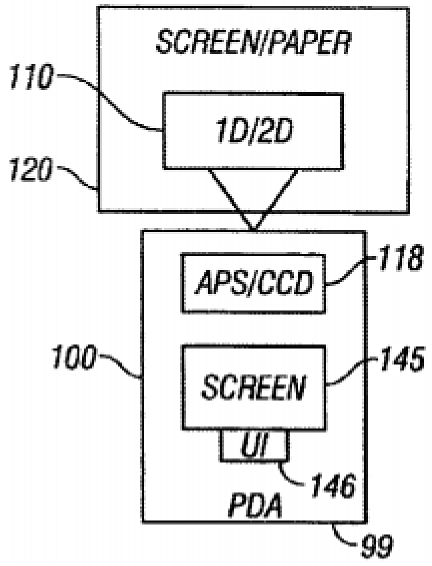 FIG. 1A 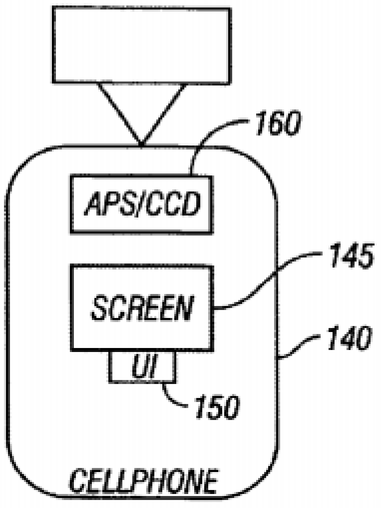 FIG. 1B Harris at 2: 3-19 • “The bar codes can be imaged/scanned in a number of different ways. One embodiment discloses using a camera to input and decode these bar codes. The embodiments are shown in FIGS. 1A and 1B. FIG. 1A describes using **a personal digital assistant 100** as the input device. One preferred input device of this type is the Palm V(TM) type hand held computer. A bar code scanner can be used, such as the commercially available Symbol Technology SPT 1700. Alternatively, a camera add-on unit can be added to the Palm V and used as described herein. The PDA includes a screen 145 and user interface 146 all within the same housing 99. FIG. 1B shows the client being **a cellular telephone** which also may include a screen 145 and user interface 150. The cellular telephone can be

Element	'036 Claim Recitation	Exemplary Citations in Reference
		associated with either a dedicated bar code scanner or an image sensor 160 of the type used to obtain photographs for video telephony." Harris at FIG. 4 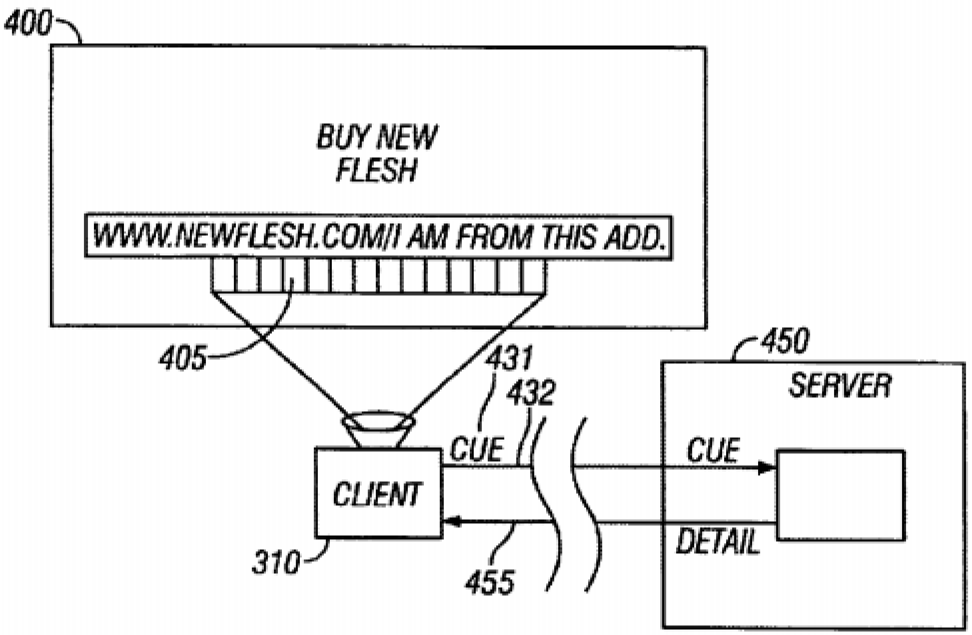 FIG. 4 Harris at 4:15-30 - "The advertisement 400 includes a bar code 405 therein. The bar code 405 is associated with the advertisement, and includes some additional information about the advertisement. For example, the bar code may include the web site address of the company preparing the advertisement, or appointment information about the advertisement, or a "vcf file" or the

Element	'036 Claim Recitation	Exemplary Citations in Reference
		like which is an importable file with address information about the company or author sponsoring the advertisement. In operation, the user brings one of the clients, either the cellular telephone 310 or PDA 315, into range of the bar code 405. The client reads the bar code and decodes it as noted above. The decoded information can represent ASCII information, a compressed file such as a zip file, G code information, or any other compressed or non-compressed information." To the extent it is found that Harris does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Harris does not anticipate this claim, Harris renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Harris with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1B.i]	communicates with a mobile device capable of acquiring a digital representation of a scene containing at least a portion of a target;	Harris discloses an identification platform coupled with the target database that communicates with a mobile device capable of acquiring a digital representation of a scene containing at least a portion of a target. Harris at FIG. 1A and FIG. 1B

Element	'036 Claim Recitation	Exemplary Citations in Reference
		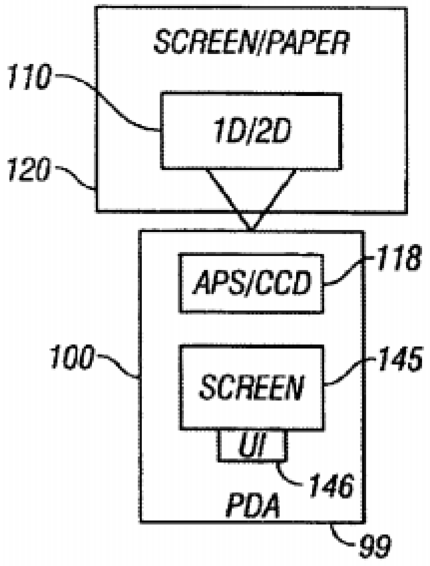 FIG. 1A 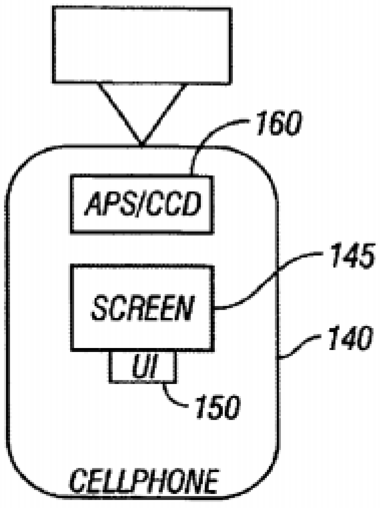 FIG. 1B Harris at 2: 3-19 • "The bar codes can be imaged/scanned in a number of different ways. One embodiment discloses using a camera to input and decode these bar codes. The embodiments are shown in FIGS. 1A and 1B. FIG. 1A describes using **a personal digital assistant 100** as the input device. One preferred input device of this type is the Palm V(TM) type hand held computer. A bar code scanner can be used, such as the commercially available Symbol Technology SPT 1700. Alternatively, a camera add-on unit can be added to the Palm V and used as described herein. The PDA includes a screen 145 and user interface 146 all within the same housing 99. FIG. 1B shows the client being **a cellular telephone** which also may include a screen 145 and user interface 150. The cellular telephone can be

Element	’036 Claim Recitation	Exemplary Citations in Reference
		associated with either a dedicated bar code scanner or an image sensor 160 of the type used to obtain photographs for video telephony.” Harris at FIG. 4 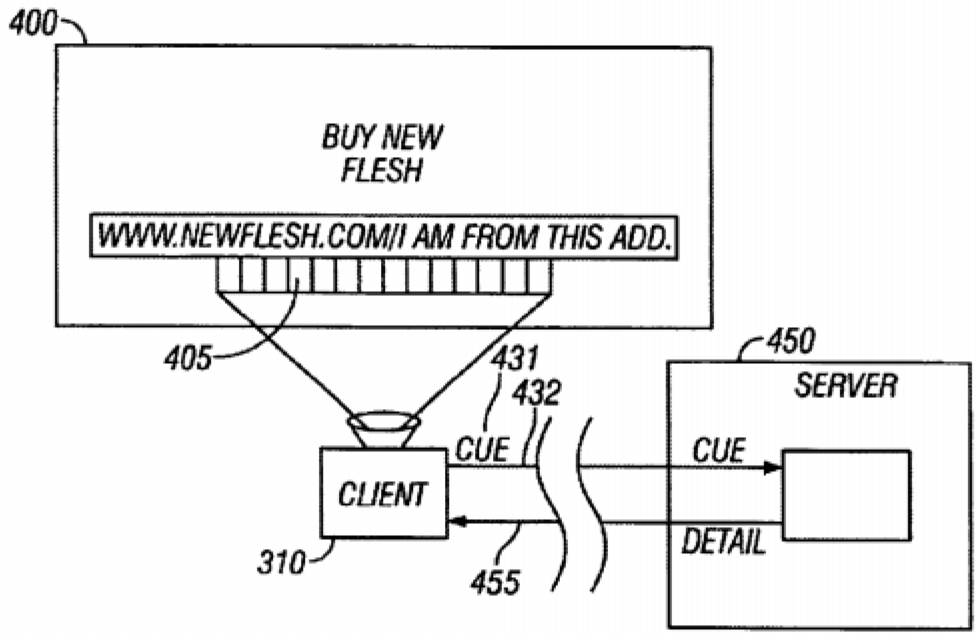 FIG. 4 Harris at 4:15-30 • “The advertisement 400 includes a bar code 405 therein. The bar code 405 is associated with the advertisement, and includes some additional information about the advertisement. For example, the bar code may include the web site address of the company preparing the advertisement, or appointment information about the advertisement, or a “vcf file” or the

Element	**'036 Claim Recitation**	**Exemplary Citations in Reference**
		like which is an importable file with address information about the company or author sponsoring the advertisement. In operation, the user brings one of the clients, either the cellular telephone 310 or PDA 315, into range of the bar code 405. The client reads the bar code and decodes it as noted above. The decoded information can represent ASCII information, a compressed file such as a zip file, G code information, or any other compressed or non-compressed information." To the extent it is found that Harris does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Harris does not anticipate this claim, Harris renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Harris with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1B.ii]	receives the digital representation from the mobile device; and	Harris discloses an identification platform coupled with the target database that receives the digital representation from the mobile device. Harris at FIG. 1A and FIG. 1B

Element	'036 Claim Recitation	Exemplary Citations in Reference
		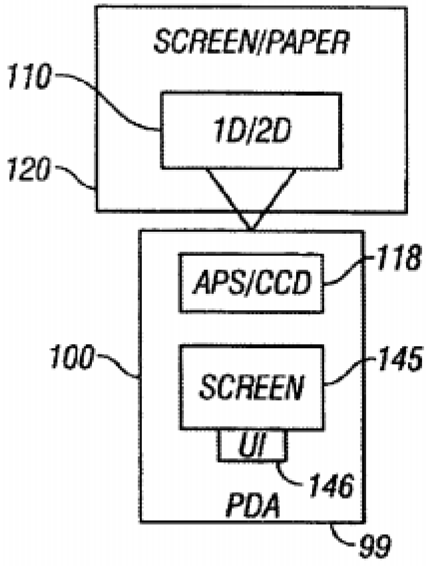 FIG. 1A 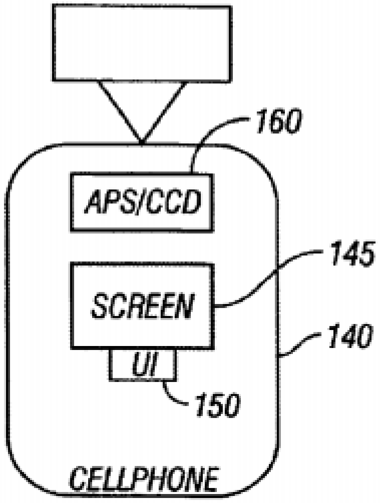 FIG. 1B Harris at 2: 3-19 • "The bar codes can be imaged/scanned in a number of different ways. One embodiment discloses using a camera to input and decode these bar codes. The embodiments are shown in FIGS. 1A and 1B. FIG. 1A describes using **a personal digital assistant 100** as the input device. One preferred input device of this type is the Palm V(TM) type hand held computer. A bar code scanner can be used, such as the commercially available Symbol Technology SPT 1700. Alternatively, a camera add-on unit can be added to the Palm V and used as described herein. The PDA includes a screen 145 and user interface 146 all within the same housing 99. FIG. 1B shows the client being **a cellular telephone** which also may include a screen 145 and user interface 150. The cellular telephone can be

Element	'036 Claim Recitation	Exemplary Citations in Reference
		associated with either a dedicated bar code scanner or an image sensor 160 of the type used to obtain photographs for video telephony." Harris at FIG. 4 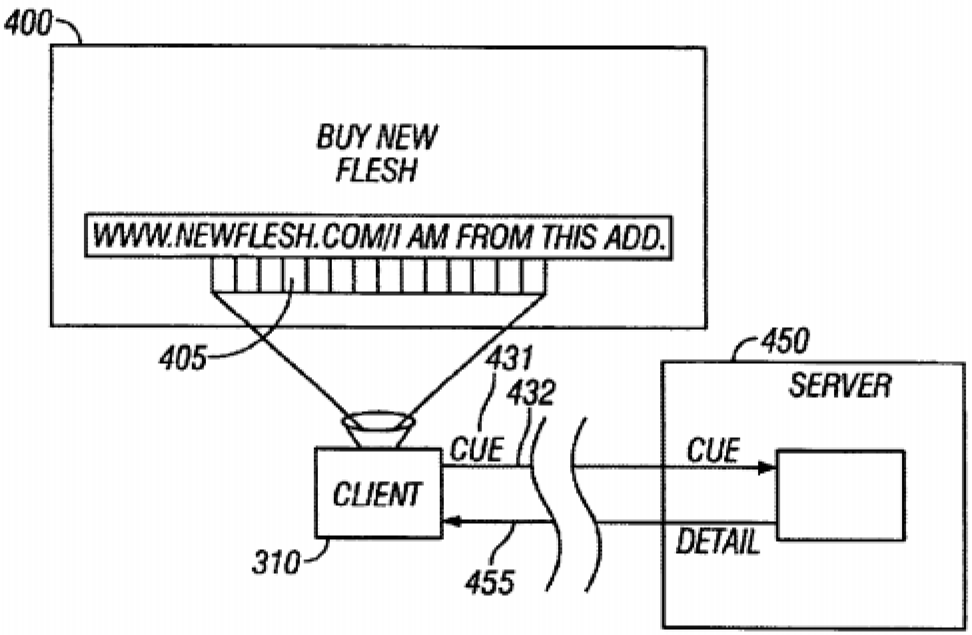 FIG. 4 Harris at 4:15-30 - "The advertisement 400 includes a bar code 405 therein. The bar code 405 is associated with the advertisement, and includes some additional information about the advertisement. For example, the bar code may include the web site address of the company preparing the advertisement, or appointment information about the advertisement, or a "vcf file" or the

Element	'036 Claim Recitation	Exemplary Citations in Reference
		like which is an importable file with address information about the company or author sponsoring the advertisement. In operation, the user brings one of the clients, either the cellular telephone 310 or PDA 315, into range of the bar code 405. The client reads the bar code and decodes it as noted above. The decoded information can represent ASCII information, a compressed file such as a zip file, G code information, or any other compressed or non-compressed information." To the extent it is found that Harris does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Harris does not anticipate this claim, Harris renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Harris with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1B.iii]	recognizes the target as a known target from the target database based on comparing parameters derived from the digital representation to recognition parameters associated with the known targets; and	Harris discloses recognizes the target as a known target from the target database based on comparing parameters derived from the digital representation to recognition parameters associated with the known targets. Harris at FIG. 1A and FIG. 1B

Element	'036 Claim Recitation	Exemplary Citations in Reference
		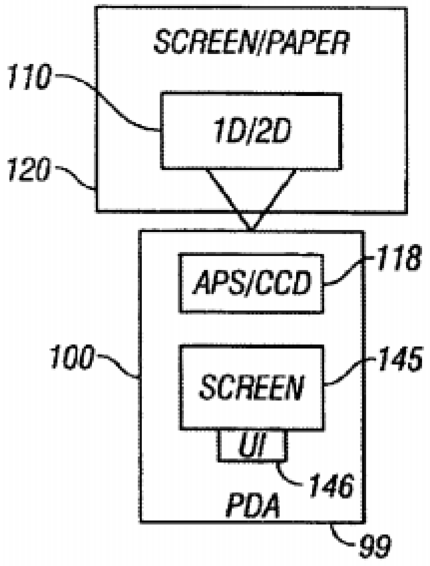 FIG. 1A 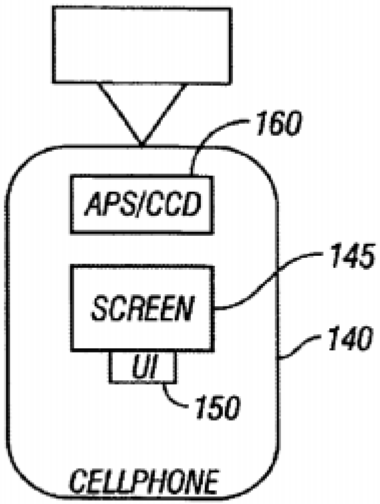 FIG. 1B Harris at 2: 3-19 • "The bar codes can be imaged/scanned in a number of different ways. One embodiment discloses using a camera to input and decode these bar codes. The embodiments are shown in FIGS. 1A and 1B. FIG. 1A describes using **a personal digital assistant 100** as the input device. One preferred input device of this type is the Palm V(TM) type hand held computer. A bar code scanner can be used, such as the commercially available Symbol Technology SPT 1700. Alternatively, a camera add-on unit can be added to the Palm V and used as described herein. The PDA includes a screen 145 and user interface 146 all within the same housing 99. FIG. 1B shows the client being **a cellular telephone** which also may include a screen 145 and user interface 150. The cellular telephone can be

Element	'036 Claim Recitation	Exemplary Citations in Reference
		associated with either a dedicated bar code scanner or an image sensor 160 of the type used to obtain photographs for video telephony." Harris at FIG. 4 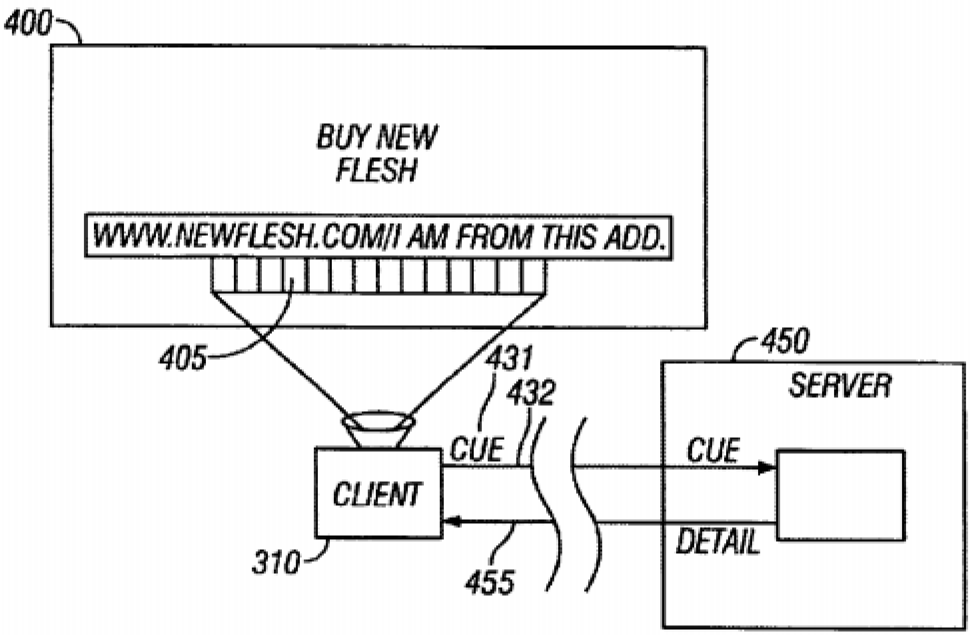 FIG. 4 Harris at 4:15-30 - "The advertisement 400 includes a bar code 405 therein. The bar code 405 is associated with the advertisement, and includes some additional information about the advertisement. For example, the bar code may include the web site address of the company preparing the advertisement, or appointment information about the advertisement, or a "vcf file" or the

Element	'036 Claim Recitation	Exemplary Citations in Reference
		like which is an importable file with address information about the company or author sponsoring the advertisement. In operation, the user brings one of the clients, either the cellular telephone 310 or PDA 315, into range of the bar code 405. The client reads the bar code and decodes it as noted above. The decoded information can represent ASCII information, a compressed file such as a zip file, G code information, or any other compressed or non-compressed information." To the extent it is found that Harris does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Harris does not anticipate this claim, Harris renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Harris with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1C]	a content service coupled with the identification platform, and that:	Harris discloses a content service coupled with the identification platform. Harris at FIG. 4

Element	'036 Claim Recitation	Exemplary Citations in Reference
		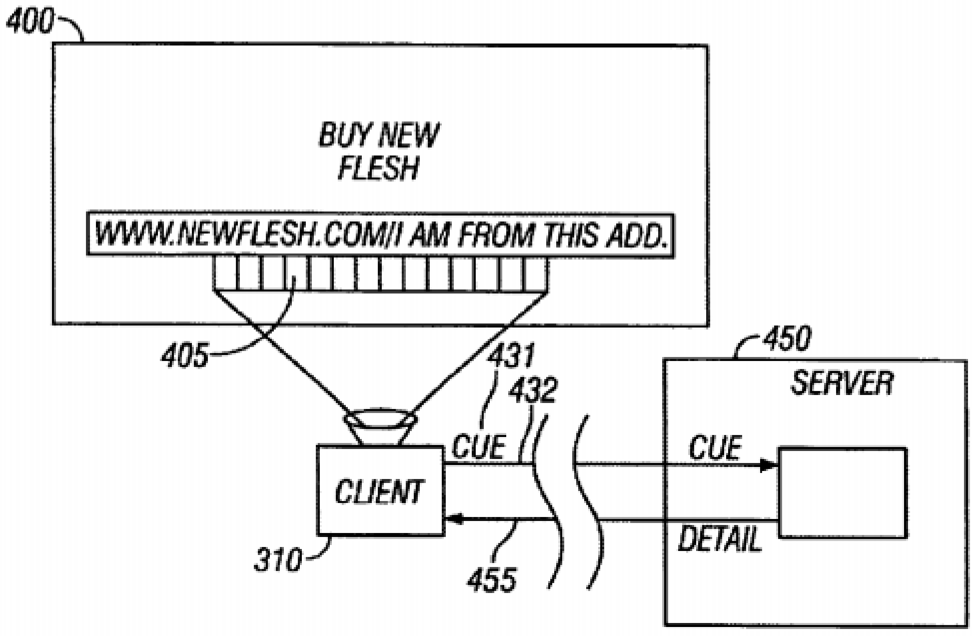 FIG. 4 Harris at 4:41-48 • "In this system, the information can represent a cue, e.g. a number, that is associated with a special function. The cue is used during a hot sync to access a publicly available information network, e.g. by direct connection or by the Internet. The cue 431 is sent over the Internet 432 to the server 450, and addresses more detailed information in a memory of the server. The server returns that information as 455, and the client receives the more detailed information." To the extent it is found that Harris does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Harris does not anticipate

Element	'036 Claim Recitation	Exemplary Citations in Reference
		this claim, Harris renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Harris with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1C.i]	obtains content information related to the known target; and	Harris discloses [a content service coupled with the identification platform] that obtains content information related to the known target. Harris at FIG. 4

Element	'036 Claim Recitation	Exemplary Citations in Reference
		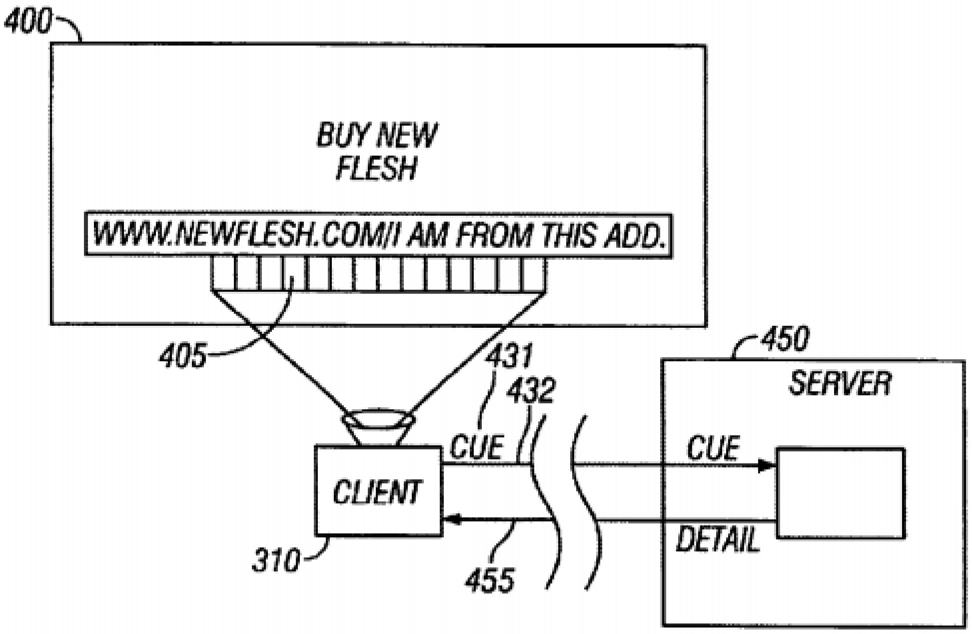 FIG. 4 Harris at 4:41-48 • "In this system, the information can represent a cue, e.g. a number, that is associated with a special function. The cue is used during a hot sync to access a publicly available information network, e.g. by direct connection or by the Internet. The cue 431 is sent over the Internet 432 to the server 450, and addresses more detailed information in a memory of the server. The server returns that information as 455, and the client receives the more detailed information." To the extent it is found that Harris does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Harris does not anticipate

Element	**'036 Claim Recitation**	**Exemplary Citations in Reference**
		this claim, Harris renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Harris with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1C.ii]	sends the content information to at least one of the identification platform and the mobile device.	Harris discloses [a content service coupled with the identification platform] that sends the content information to at least one of the identification platform and the mobile device. Harris at FIG. 4

Element	'036 Claim Recitation	Exemplary Citations in Reference
		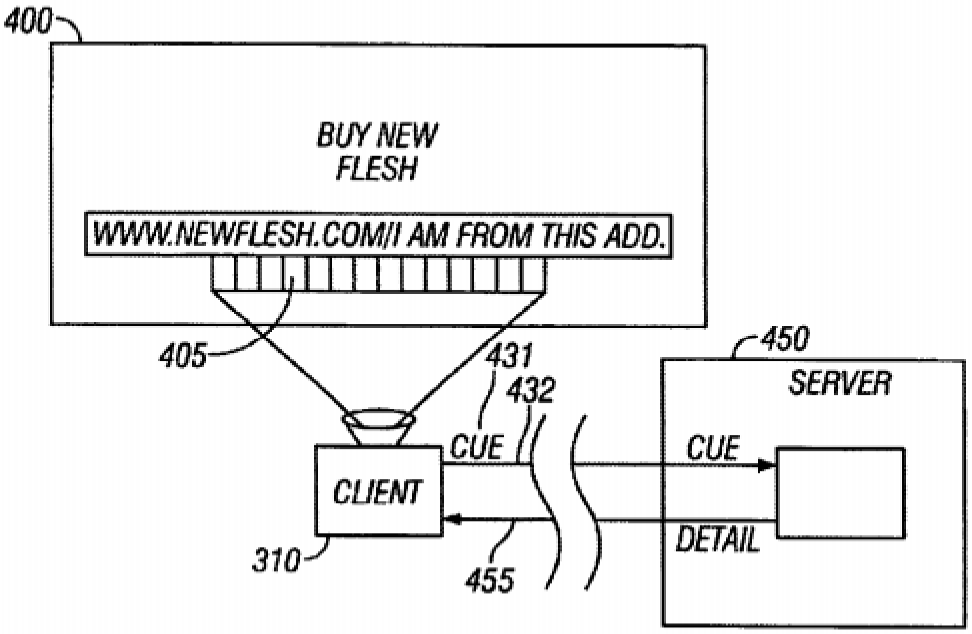 FIG. 4 Harris at 4:41-48 - "In this system, the information can represent a cue, e.g. a number, that is associated with a special function. The cue is used during a hot sync to access a publicly available information network, e.g. by direct connection or by the Internet. The cue 431 is sent over the Internet 432 to the server 450, and addresses more detailed information in a memory of the server. The server returns that information as 455, and the client receives the more detailed information." To the extent it is found that Harris does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Harris does not anticipate

Element	'036 Claim Recitation	Exemplary Citations in Reference
		this claim, Harris renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Harris with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

Nantworks, LLC v. Bank of America Corporation
Central District of California, No. 2:20-cv-7872
Invalidity Chart for U.S. Patent 8,478,036 based on
U.S. Patent No. 6,722,569 to Ehrhart et al. ("Ehrhart")

Exhibit C-11

Bank of America Corporation ("BofA" or "Defendant") provides this chart subject to all reservations, objections, statements, and disclaimers set forth herein and in BofA's Final Invalidity Contentions Cover Pleading, as well as any amendment, supplement, or modification thereof, which are incorporated herein by reference in their entirety.

On information and belief, the application that matured into Ehrhart was filed on July 13, 2001, claiming the priority date of July 13, 2001. Ehrhart published at least as of April 20, 2004, and is available as prior art at least under 35 U.S.C. § 102(e).

As shown in the chart below, Ehrhart anticipates asserted claim 1 of the '036 Patent. To the extent Ehrhart is found not to expressly disclose certain limitations in the asserted claims, such limitations are inherent. To the extent Ehrhart is found not to anticipate any asserted claims or claim elements of the '036 Patent, the reference nevertheless renders those claims or claim elements obvious under 35 U.S.C. § 103, either alone or in combination with other art identified herein or in the cover pleading of BofA's Invalidity Contentions, for the reasons identified in the cover pleading. These Invalidity Contentions are not an admission by BofA that the accused products, including any current or past versions of the accused products, are covered by, or infringe the asserted claims, but are based instead on the recognition that if the claims are interpreted to be broad enough to encompass the accused products, the claims must also be construed to have that same scope when considering whether they are invalid.

The following chart is based on the Court's claim construction rulings (Dkt. Nos. 145, 153, 236), and to the extent not contradictory, Plaintiffs' apparent interpretations of the claims in Plaintiffs' "Preliminary" Final Infringement Contentions. BofA notes that in many instances, Plaintiffs' "Preliminary" Final Infringement Contentions ignore and/or contradict the Court's claim construction rulings. BofA does not accept the assumptions concerning the scope and meaning implicit in Plaintiffs' Preliminary Final Infringement Contentions, to the extent those assumptions are discernible, and reserves the right to challenge Plaintiffs' proposed (or implied) constructions. Furthermore, Plaintiffs' "Preliminary" Final Contentions fail to comply with the level of specificity required by SPR 2.1.2 and 2.1.3, and include improper reservation of rights to supplement with "additional ways in which BoA Accused Products infringe." Any alleged deficient discovery is of Plaintiffs own making. Most if not all of the deficiencies in Plaintiffs' Preliminary Final Contentions have been documented by Defendants through multiple letters since at least May 18, 2021 and up to Bank of America's most recent letter of August 19, 2023. Plaintiffs' shifting infringement contentions, deficiencies in its infringement contentions, and changing claim scope have prejudiced BofA from being able to reasonably prepare its defenses, including its invalidity contentions. BofA also reserves the right to revise and supplement these charts if and when Plaintiffs are permitted to provide revised or supplemental Infringement Contentions. Citations given in the chart below are merely representative of the respective elements to provide notice to Plaintiffs of BofA's invalidity theories and are not meant to be exhaustive.

1. Claim 1

Element	'036 Claim Recitation	Exemplary Citations in Reference
[1Pre]	A content provisioning system comprising:	Ehrhart discloses a content provisioning system. "The present invention relates to an optical reader that includes a color imaging assembly that generates color imaging data. An image analysis circuit determines if the acquired image should be characterized as a color photograph or as including a graphical symbol. A processing circuit processes the imaging data based on the image analysis circuit's determination of whether the image is a graphical symbol or a color photograph. The present invention allows a user to acquire and process both color images and graphical symbols, such as bar codes, text, OCR symbols or signatures. The optical reader of the present invention is also configured to associate an acquired image with at least one other acquired image." Ehrhart at Abstract.
[1A]	a target database storing known targets of different types and recognition parameters associated with the known targets;	Ehrhart discloses a target database storing known targets of barcodes. *Id.* at [1Pre]. "In one embodiment, this mode includes a Signature verification program wherein the user may select between a Static verification or a dynamic verification. In the Static mode, the user captures the image of a Signature. The captured image is compared with a reference image Stored in a remote database. In the dynamic mode, optical reader 10 uses the Stylus and Signature block to capture the Signature. In this mode, Signature block 62 measures unique dynamic parameters, Such as applied pressure, direction and timing of movements, or a combination of these parameters. One of ordinary skill in the art will recognize that this list is not meant to be all-

Element	'036 Claim Recitation	Exemplary Citations in Reference
		inclusive, but rather, is a representative example. The captured dynamic parameters are compared with a reference data Stored in a remote database." Ehrhart at 8:64-9:10. "LAN 1820 includes server 1822, computer 1824, database 1826, and a plurality of optical readers 10. Database 1826 is used to store associated images along with other data fields. For example, it would be rather useful to store additional information with the associated images shown in FIG. 13. One may want to associate the delivery means, route, driver, and other related information for subsequent analysis. Network 1810 allows reader 10, PAN 1850, and wireless system 1400 a way to store such data in database 1826. System analysts can access this information via personal computer 1802 connected to network 1810. In one embodiment, LAN 1820 includes an Internet website. In this embodiment, users are authenticated before gaining access to database 1826." Ehrhart at 16:1-12. To the extent it is found that Ehrhart does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Ehrhart does not anticipate this claim, Ehrhart renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Ehrhart with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1B]	an identification platform coupled with the target database, and that	Ehrhart discloses an identification platform coupled with the target database. *Id.* at [1Pre-A].

Element	'036 Claim Recitation	Exemplary Citations in Reference
		"As embodied herein, and depicted in FIGS. 1A-1D, perspective views of the optical reader in accordance with various embodiments of the present invention are disclosed. FIG. 1A shows the underside of hand held wireless optical reader 10. FIG. 1B shows the top of the optical reader depicted in FIG. 1A. Optical reader 10 includes housing 100, antenna 102, window 104 and trigger 12. Window 104 accommodates illumination assembly 20 and imaging assembly 30. As shown in FIG. 1B, the top side of reader 10 includes function keys 14, alphanumeric key pad 16, and display 60. In one embodiment, function keys 14 include an enter key and up and down cursor keys. FIG. 1C is also a hand held wireless optical reader 10. Reader 10 includes function keys 14, alphanumeric key pad 16, writing stylus 18, display 60, and signature block 62. Stylus 18 is employed by a user to write his signature in signature block 62. FIG. 1D shows yet another embodiment of optical reader 10 of the present invention. In this embodiment, reader 10 includes a gun-shaped housing 100. Display 60 and keypad 16 are disposed on a top portion of gun-shaped housing 100, whereas trigger 12 is disposed on the underside of the top portion of housing 100. Housing 100 also includes window 104 that accommodates illumination assembly 20 and imaging assembly 30. Wire 106 is disposed at the butt-end of housing 100. Wire 106 provides optical reader 10 with a hard wired communication link for external devices such as a host processor or other data collection devices." Ehrhart at 5:51-6:10. "In one embodiment, this mode includes a Signature verification program wherein the user may select between a Static verification or a dynamic verification. In the Static mode, the user captures the image of a Signature. The captured image is compared with a reference image Stored in a remote database. In the dynamic mode, optical reader 10 uses the Stylus and Signature block to capture the Signature. In this mode, Signature block 62 measures unique dynamic parameters, Such as applied pressure, direction and timing of movements, or a combination of these parameters. One of ordinary skill in the art will recognize that this list is not meant to be all-

Element	'036 Claim Recitation	Exemplary Citations in Reference
		inclusive, but rather, is a representative example. The captured dynamic parameters are compared with a reference data Stored in a remote database." Ehrhart at 8:64-9:10. To the extent it is found that Ehrhart does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Ehrhart does not anticipate this claim, Ehrhart renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Ehrhart with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1B.i]	communicates with a mobile device capable of acquiring a digital representation of a scene containing at least a portion of a target;	Ehrhart discloses an identification platform coupled with the target database that communicates with a mobile device capable of acquiring a digital representation of a scene containing at least a portion of a target. *Id.* at [1Pre-B]. "As embodied herein, and depicted in FIGS. 1A-1D, perspective views of the optical reader in accordance with various embodiments of the present invention are disclosed. FIG. 1A shows the underside of hand held wireless optical reader 10. FIG. 1B shows the top of the optical reader depicted in FIG. 1A. Optical reader 10 includes housing 100, antenna 102, window 104 and trigger 12. Window 104 accommodates illumination assembly 20 and imaging assembly 30. As shown in FIG. 1B, the top side of reader 10 includes function keys 14, alphanumeric key pad 16, and display 60. In one embodiment, function keys 14 include an enter key and up and down cursor keys. FIG. 1C is also a hand

Element	**'036 Claim Recitation**	**Exemplary Citations in Reference**
		held wireless optical reader 10. Reader 10 includes function keys 14, alphanumeric key pad 16, writing stylus 18, display 60, and signature block 62. Stylus 18 is employed by a user to write his signature in signature block 62. FIG. 1D shows yet another embodiment of optical reader 10 of the present invention. In this embodiment, reader 10 includes a gun-shaped housing 100. Display 60 and keypad 16 are disposed on a top portion of gun-shaped housing 100, whereas trigger 12 is disposed on the underside of the top portion of housing 100. Housing 100 also includes window 104 that accommodates illumination assembly 20 and imaging assembly 30. Wire 106 is disposed at the butt-end of housing 100. Wire 106 provides optical reader 10 with a hard wired communication link for external devices such as a host processor or other data collection devices." Ehrhart at 5:51-6:10. To the extent it is found that Ehrhart does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Ehrhart does not anticipate this claim, Ehrhart renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Ehrhart with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1B.ii]	receives the digital representation from the mobile device; and	Ehrhart discloses an identification platform coupled with the target database that receives the digital representation from the mobile device. *Id.* at [1Pre-B].

Element	'036 Claim Recitation	Exemplary Citations in Reference
		"In accordance with the invention, the present invention for an optical reader includes a color imaging assembly for acquiring color imaging data. An image analysis circuit determines if the acquired image includes at least one graphical symbol. A processing circuit processes the imaging data based on the determination of whether the image includes at least one graphical symbol. The present invention allows a user to read graphical symbols, such as bar codes, text, OCR characters or signatures using a color imager. The color optical reader of the present invention is configured to automatically determine whether a color image includes a graphical symbol, or is merely a color photographic image. The optical reader of the present invention also is operative to associate one acquired image with at least one subsequently acquired image." Ehrhart at 5:36-50. "In the embodiment depicted in FIG. 2, ASIC 44 is implemented using a programmable logic array (PLA) device. In a similar embodiment, ASIC 44 is implemented using a field programmable gate array (FPGA) device. ASIC 44 is tasked with controlling the image acquisition process, and the storage of image data. As part of the image acquisition process, ASIC 44 performs various timing and control functions including control of light source 24, control of color imager 34, and control of external interface 56. It will be apparent to those of ordinary skill in the pertinent art that modifications and variations can be made to processor 40 of the present invention depending on the cost, availability, and performance of off-the-shelf microprocessors, and the type of color imager used. In one embodiment, microprocessor 42 and ASIC 44 are replaced by a single microprocessor 40. In one embodiment, microprocessor 40 is implemented using a single RISC processor. In yet another embodiment, microprocessor 40 is implemented using a RISC and DSP hybrid processor." Ehrhart at 7:32-50. To the extent it is found that Ehrhart does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Ehrhart does not

Element	'036 Claim Recitation	Exemplary Citations in Reference
		anticipate this claim, Ehrhart renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Ehrhart with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1B.iii]	recognizes the target as a known target from the target database based on comparing parameters derived from the digital representation to recognition parameters associated with the known targets; and	Ehrhart discloses recognizes the target as a known target from the target database based on comparing parameters derived from the digital representation to recognition parameters associated with the known targets. *Id.* at [1Pre-B]. "In the Automatic imaging mode, processor 40 is programmed to analyze the color imaging data to determine if an acquired image includes a graphical symbol or is merely a color photographic image. If it makes the determination that the color image includes a graphical symbol, it further analyzes the acquired image and classifies it as a bar code, OCR symbol, text, or a signature. Based on the classification, optical reader 10 jumps to the appropriate routine in EROM 48. The semi-automatic mode is similar. Thus, in the automatic or semi-automatic modes, the bar code scanning mode, the OCR/text mode, the signature capture mode, the color photography mode, and the association mode are controlled by the application program, not by the user." Ehrhart at 8:36-49. "In one embodiment, this mode includes a Signature verification program wherein the user may select between a Static verification or a dynamic verification. In the Static mode, the user captures the image of a Signature. The captured image is compared with a reference image Stored in a remote database. In the dynamic mode,

Element	'036 Claim Recitation	Exemplary Citations in Reference
		optical reader 10 uses the Stylus and Signature block to capture the Signature. In this mode, Signature block 62 measures unique dynamic parameters, Such as applied pressure, direction and timing of movements, or a combination of these parameters. One of ordinary skill in the art will recognize that this list is not meant to be all-inclusive, but rather, is a representative example. The captured dynamic parameters are compared with a reference data Stored in a remote database." Ehrhart at 8:64-9:10. "If processor 40 does not locate a bar code symbol it moves onto step 420 and looks for OCR-A or OCR-B characters. If it finds these characters it performs optical character recognition in step 422. If it does not, processor evaluates the image for the presence of text. If text is located, the image is cropped, and the text is compressed and stored in steps 428 and 430. If the image does not include text, processor 40 evaluates the image for the presence of a signature. If one is present, the image is cropped, and the data is compressed and stored in steps 428 and 430. In another embodiment, optical reader 10 is networked, and processor 40 communicates with remote network resources to provide signature verification services. If processor 40 cannot detect a bar code symbol, OCR symbols, text, or a signature, the user is asked in step 432 if he desires to store the image. If he does, the color photographic image is stored in memory in step 434." Ehrhart at 10:5-20. "As embodied herein and depicted in FIG. 11, a flow chart showing a method for performing OCR in accordance with yet another embodiment of the present invention is disclosed. In step 1100, reader 10 produces a bit-mapped image of the page. Subsequently, processor 40 finds lines of text in the image, locates the white spaces in each line, and isolates the characters. In step 1108, processor 40 performs character recognition, either OCR-A or OCR-B, as desired. The decoded characters are stored in memory." Ehrhart at 13:1-9.

		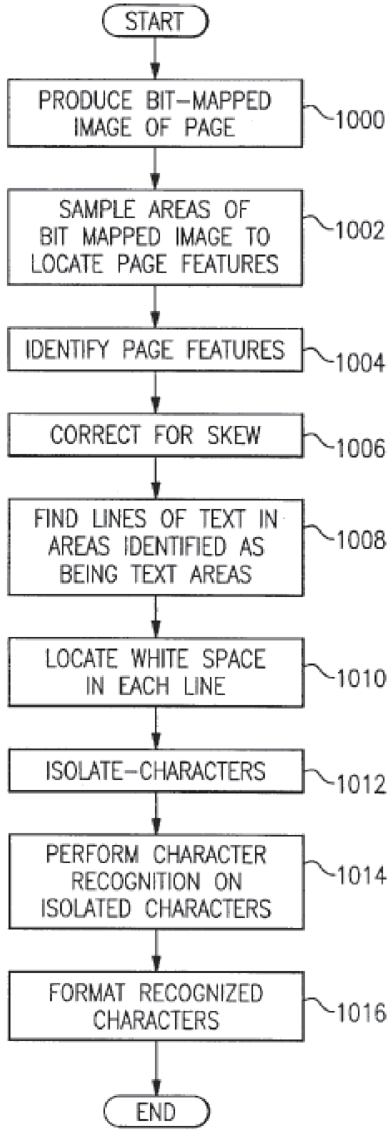START PRODUCE BIT-MAPPED IMAGE OF PAGE 1000 SAMPLE AREAS OF BIT MAPPED IMAGE TO LOCATE PAGE FEATURES 1002 IDENTIFY PAGE FEATURES 1004 CORRECT FOR SKEW 1006 FIND LINES OF TEXT IN AREAS IDENTIFIED AS BEING TEXT AREAS 1008 LOCATE WHITE SPACE IN EACH LINE 1010 ISOLATE-CHARACTERS 1012 PERFORM CHARACTER RECOGNITION ON ISOLATED CHARACTERS 1014 FORMAT RECOGNIZED CHARACTERS 1016 END FIG.10

Element	'036 Claim Recitation	Exemplary Citations in Reference
		Ehrhart at Fig. 10. To the extent it is found that Ehrhart does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Ehrhart does not anticipate this claim, Ehrhart renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Ehrhart with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1C]	a content service coupled with the identification platform, and that:	Ehrhart discloses a content service coupled with the identification platform. *Id.* at [1Pre-B]. "LAN 1820 includes server 1822, computer 1824, database 1826, and a plurality of optical readers 10. Database 1826 is used to store associated images along with other data fields. For example, it would be rather useful to store additional information with the associated images shown in FIG. 13. One may want to associate the delivery means, route, driver, and other related information for subsequent analysis. Network 1810 allows reader 10, PAN 1850, and wireless system 1400 a way to store such data in database 1826. System analysts can access this information via personal computer 1802 connected to network 1810. In one embodiment, LAN 1820 includes an Internet website. In this embodiment, users are authenticated before gaining access to database 1826." Ehrhart at 16:1-12.

Element	'036 Claim Recitation	Exemplary Citations in Reference
		To the extent it is found that Ehrhart does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Ehrhart does not anticipate this claim, Ehrhart renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Ehrhart with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1C.i]	obtains content information related to the known target; and	Ehrhart discloses [a content service coupled with the identification platform] that obtains content information related to the known target. *Id.* at [1Pre-C]. "LAN 1820 includes server 1822, computer 1824, database 1826, and a plurality of optical readers 10. Database 1826 is used to store associated images along with other data fields. For example, it would be rather useful to store additional information with the associated images shown in FIG. 13. One may want to associate the delivery means, route, driver, and other related information for subsequent analysis. Network 1810 allows reader 10, PAN 1850, and wireless system 1400 a way to store such data in database 1826. System analysts can access this information via personal computer 1802 connected to network 1810. In one embodiment, LAN 1820 includes an Internet website. In this embodiment, users are authenticated before gaining access to database 1826." Ehrhart at 16:1-12.

Element	'036 Claim Recitation	Exemplary Citations in Reference
		To the extent it is found that Ehrhart does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Ehrhart does not anticipate this claim, Ehrhart renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Ehrhart with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1C.ii]	sends the content information to at least one of the identification platform and the mobile device.	Ehrhart discloses [a content service coupled with the identification platform] that sends the content information to at least one of the identification platform and the mobile device. *Id.* at [1Pre-C]. "LAN 1820 includes server 1822, computer 1824, database 1826, and a plurality of optical readers 10. Database 1826 is used to store associated images along with other data fields. For example, it would be rather useful to store additional information with the associated images shown in FIG. 13. One may want to associate the delivery means, route, driver, and other related information for subsequent analysis. Network 1810 allows reader 10, PAN 1850, and wireless system 1400 a way to store such data in database 1826. System analysts can access this information via personal computer 1802 connected to network 1810. In one embodiment, LAN 1820 includes an Internet website. In this embodiment, users are authenticated before gaining access to database 1826." Ehrhart at 16:1-12.

Element	'036 Claim Recitation	Exemplary Citations in Reference
		To the extent it is found that Ehrhart does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Ehrhart does not anticipate this claim, Ehrhart renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Ehrhart with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

Exhibit C-23

Bank of America Corporation ("BofA" or "Defendant") provides this chart subject to all reservations, objections, statements, and disclaimers set forth herein and in BofA's Final Invalidity Contentions Cover Pleading, as well as any amendment, supplement, or modification thereof, which are incorporated herein by reference in their entirety. A statement of reasons and good cause for each supplement in this chart is identified in the cover pleading.

QBIC was known or used by others in the United States and publicly used or on sale in the United States by November 5, 2000. *See generally* W. Niblack et al., The QBIC Project: Querying Images By Content Using Color, Texture, and Shape, SPIE Vol. 1908 (1993), pp. 173–187 (IBM0002390–404); M. Flickner et al., Query by Image and Video Content: The QBIC System ("QBIC") (IBM0000893)."), IEEE Sept. 1995, pp. 23–32 (IBM0000893–902); W. Niblack et al, Updates to the QBIC System, SPIE Vol. 3312 (1997), pp. 150–161 (IBM 0002419–430); and Query By Image Content QBIC Demonstration Program, IBM Research, Sept. 1998 (IBM000747). QBIC is available as prior art at least under 35 U.S.C. § 102(a), (b), and (g), having been known or used by others in the United States and publicly used or on sale in the United States by at least 1995 and certainly by November 5, 2000, including, for example, at http://libra.uc.davis.edu through the Art History Department at U.C. Davis; at www..thinker.org/imagebase/indox-2 .htm3 through the Fine Arts Museum of San Francisco; through demos at http://wwwqbic.almaden.ibm.com; and through IBM's commercial products, including in IBM's Ultimedia Manager, the DB2 series of products such as the Image Extender in DB2 Universal Database, and the Enterprise Information Portal. *See generally* Deposition of IBM Corporate Representative Myron Flickner, *Nantworks, LLC and Nant Holdings IP, LLC v. Bank of America Corporation and Bank of America, N.A.*, 2:20-cv-07872-GW-PVC (C.D.C.A. 2020) ("Flickner 8/29/23 Depo") at e.g., 32–33, 39–40, 44, 62–63, 66–67,72–74, 103–107, 110–112, 127, 163–174, 183–189, 191–195, 215, 219–220, 241, 307–309, 311–317, Exs. 3–13; *see also* Deposition of IBM Corporate Representative Myron Flickner, *Network-1 Technologies, Inc., v. Google LLC and Youtube LLC*, 1:14-cv-09558-PGG (S.D.N.Y. 2019) ("Flickner Depo"), including, e.g., at 13, 14, 15–18, 29–30, 34, 36, 37, 43–46, 48, 58, 60, 61–62, 67–69, 72, 74–76, 84–86, 88–89, 98–100, 102, 106, 113, 125, 135, 152–154, 160–161, 170–172, 180–181, 194, 198–201, 203–205, 215–216, 220–221, and Exs. 1–16.

As shown in the chart below, QBIC anticipates asserted claim 1 of the '036 Patent. The disclosures/exemplary citations for each element incorporate the disclosures citations of the proceeding limitations. To the extent QBIC is found not to expressly disclose certain limitations in the asserted claims, such limitations are inherent. To the extent QBIC is found not to anticipate any asserted claims or claim elements of the '036 Patent, the reference nevertheless renders those claims or claim elements obvious under 35 U.S.C. § 103, either alone or in combination with other art identified herein or in the cover pleading of BofA's Final Invalidity Contentions, for the reasons identified in the cited references and/or in the cover pleading. These Invalidity Contentions are not an admission by BofA that the accused products, including any current or past versions of the accused products, are covered by, or infringe the asserted claims, but are based instead on the recognition that if the claims are interpreted to be broad enough to encompass the accused products, the claims must also be construed to have that same scope when considering whether they are invalid. The following chart incorporates by reference the analysis from Exhibit C-21 for each limitation.

The following chart is based on the Court's claim construction rulings (Dkt. Nos. 145, 153, 236), and to the extent not contradictory, Plaintiffs' apparent interpretations of the claims in Plaintiffs' "Preliminary" Final Infringement Contentions. BofA notes that in many instances, Plaintiffs' "Preliminary" Final Infringement Contentions ignore and/or contradict the Court's claim construction rulings. BofA does not accept the assumptions concerning the scope and meaning implicit in Plaintiffs' Preliminary Final Infringement Contentions, to the extent those assumptions are discernible, and reserves the right to challenge Plaintiffs' proposed (or implied) constructions. Furthermore, Plaintiffs' "Preliminary" Final Contentions fail to comply with the level of specificity required by SPR 2.1.2 and 2.1.3, and include improper reservation of rights to supplement with "additional ways in which BoA Accused Products infringe." Any alleged deficient discovery is of Plaintiffs own making. Most if not all of the deficiencies in Plaintiffs' Preliminary Final Contentions have been documented by Defendants through multiple letters since at least May 18, 2021 and up to Bank of America's most recent letter of August 19, 2023. Plaintiffs' shifting infringement contentions, deficiencies in its infringement contentions, and changing claim scope have prejudiced BofA from being able to reasonably prepare its defenses, including its invalidity contentions. BofA also reserves the right to revise and supplement these charts if and when Plaintiffs are permitted to provide revised or supplemental Infringement Contentions. Citations given in the chart below are merely representative of the respective elements to provide notice to Plaintiffs of BofA's invalidity theories and are not meant to be exhaustive.

1. Claim 1

Element	'036 Claim Recitation	Exemplary Citations in Reference
[1Pre]	A content provisioning system comprising:	QBIC discloses a content provisioning system. The QBIC Project: Querying Images by Content Using Color, Texture, and Shape (IBM 0002390–404) at 1 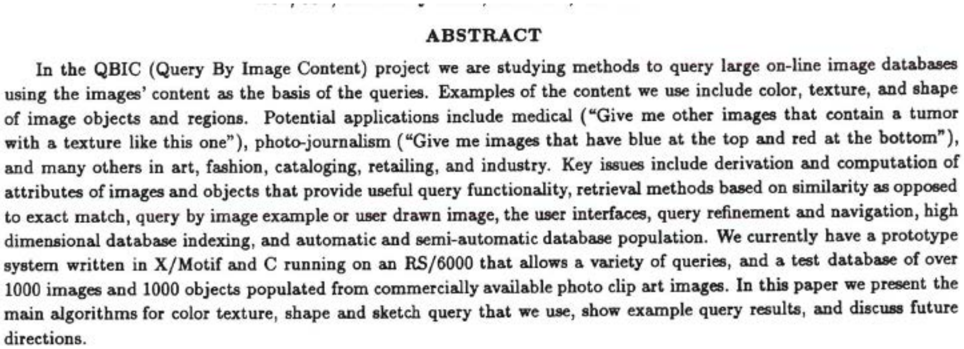ABSTRACT In the QBIC (Query By Image Content) project we are studying methods to query large on-line image databases using the images' content as the basis of the queries. Examples of the content we use include color, texture, and shape of image objects and regions. Potential applications include medical ("Give me other images that contain a tumor with a texture like this one"), photo-journalism ("Give me images that have blue at the top and red at the bottom"), and many others in art, fashion, cataloging, retailing, and industry. Key issues include derivation and computation of attributes of images and objects that provide useful query functionality, retrieval methods based on similarity as opposed to exact match, query by image example or user drawn image, the user interfaces, query refinement and navigation, high dimensional database indexing, and automatic and semi-automatic database population. We currently have a prototype system written in X/Motif and C running on an RS/6000 that allows a variety of queries, and a test database of over 1000 images and 1000 objects populated from commercially available photo clip art images. In this paper we present the main algorithms for color texture, shape and sketch query that we use, show example query results, and discuss future directions. Query by Image and Video Content: The QBIC System (IBM 0000893–902): "We have developed the QBIC (Query by Image Content) system to explore content-based retrieval methods. QBIC allows queries on large image and video databases based on • example images, • user-constructed sketches and drawings, • selected color and texture patterns, • camera and object motion, and • other graphical information." *Id.* at pp. 1–3. "To achieve this functionality, QBIC has two main components: database population (the process of creating an image database) and database query. During the population, images and videos are processed to extract features describing their content-colors, textures, shapes, and camera and object motion-and the features are stored in a database. During the query, the user composes a query graphically. Features are generated from the graphical query and then input to a matching engine that finds images or videos from the database with similar features. Figure

<table>
<tr><th>Element</th><th>'036 Claim Recitation</th><th>Exemplary Citations in Reference</th></tr>
<tr><td></td><td></td><td>

2 shows the system architecture." *Id.* at p. 3.

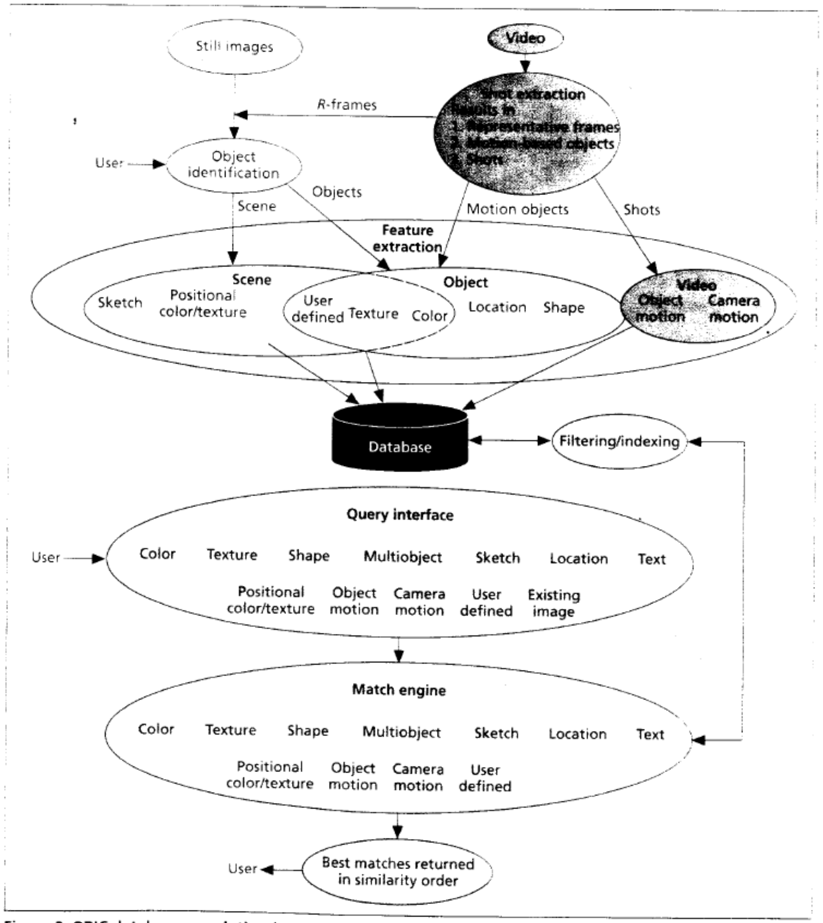

Figure 2. QBIC database population (top) and query (bottom) architecture.

Id. at Fig. 2.

</td></tr>
</table>

Element	'036 Claim Recitation	Exemplary Citations in Reference
		QBIC Demo (IBM000747):  QBIC(TM) -- IBM's Query By Image Content Goto: QBIC allows you to make queries of image databases large and small based on visual characteristics -- properties such as color percentages, color layout and textures occurring in the images. The queries match image colors, textures and their positions without having to describe them with in words. QBIC's content-based queries can also be combined with keywords to provide a powerful multimedia search and retrieval tool. As image collections grow larger and become more prolific, QBIC is the best tool to more efficiently search and navigate through them. In this demo , we use QBIC to index about 470 color images based on their color and texture information. You can evaluate the QBIC Search capability by following the link located at left called IBM Image Search Demo. For license and business related questins, please contact Ted Loewenberg by phone 408-927-1202, or by email QBIC Demo (IBM000747): 

Element	'036 Claim Recitation	Exemplary Citations in Reference
		Flicker Depo at 13: 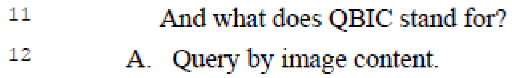 And what does QBIC stand for? A. Query by image content. Flickner Depo at 18 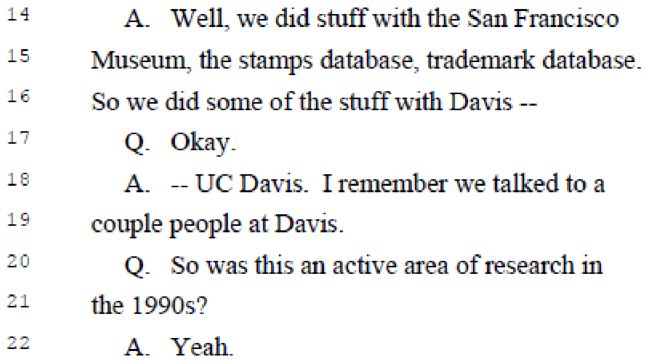 A. Well, we did stuff with the San Francisco Museum, the stamps database, trademark database. So we did some of the stuff with Davis -- Q. Okay. A. -- UC Davis. I remember we talked to a couple people at Davis. Q. So was this an active area of research in the 1990s? A. Yeah. Flickner Depo at 40:

Element	'036 Claim Recitation	Exemplary Citations in Reference
		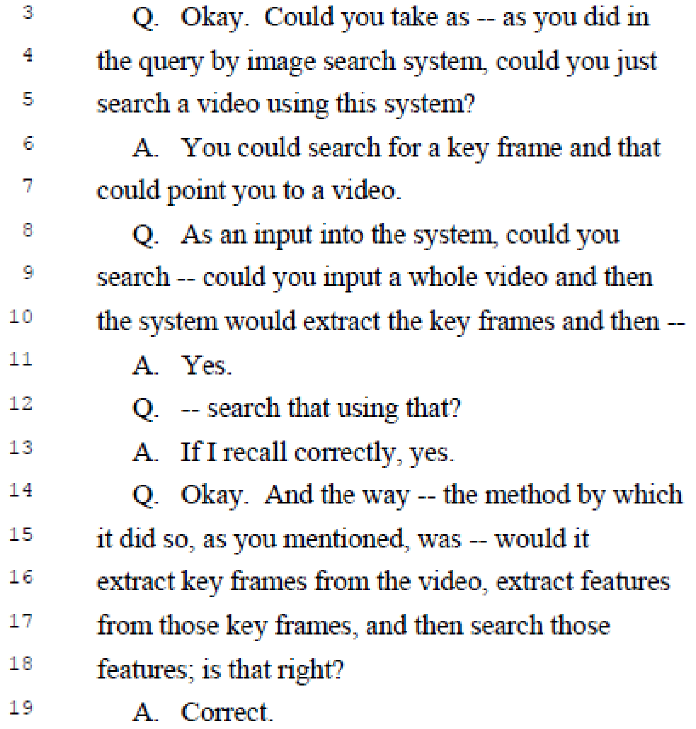 Q. Okay. Could you take as -- as you did in the query by image search system, could you just search a video using this system? A. You could search for a key frame and that could point you to a video. Q. As an input into the system, could you search -- could you input a whole video and then the system would extract the key frames and then -- A. Yes. Q. -- search that using that? A. If I recall correctly, yes. Q. Okay. And the way -- the method by which it did so, as you mentioned, was -- would it extract key frames from the video, extract features from those key frames, and then search those features; is that right? A. Correct.

Element	'036 Claim Recitation	Exemplary Citations in Reference
		Flickner Depo at 72: 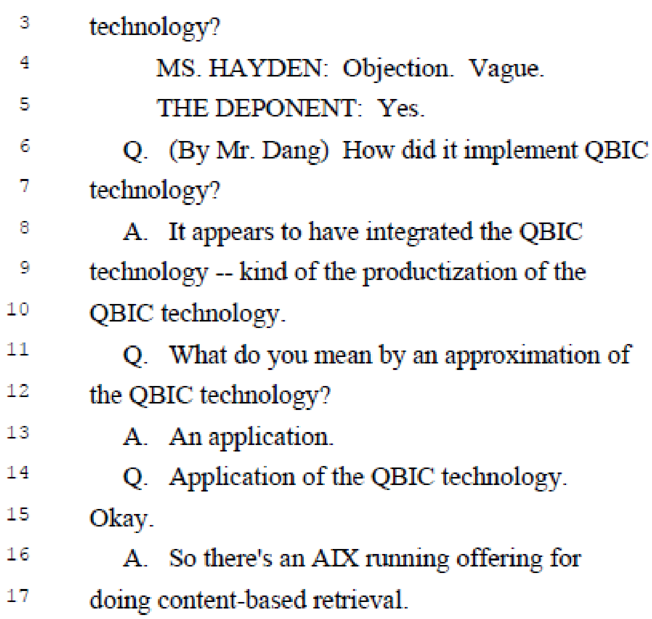 Q. And did the MediaMiner implement QBIC technology? MS. HAYDEN: Objection. Vague. THE DEPONENT: Yes. Q. (By Mr. Dang) How did it implement QBIC technology? A. It appears to have integrated the QBIC technology -- kind of the productization of the QBIC technology. Q. What do you mean by an approximation of the QBIC technology? A. An application. Q. Application of the QBIC technology. Okay. A. So there's an AIX running offering for doing content-based retrieval. Flickner Depo at 76: 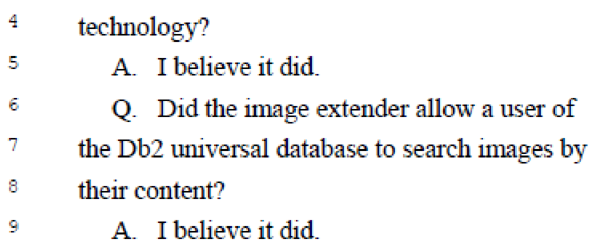 Q. Did the image extender incorporate QBIC technology? A. I believe it did. Q. Did the image extender allow a user of the Db2 universal database to search images by their content? A. I believe it did. Flickner Depo at 91:

Element	'036 Claim Recitation	Exemplary Citations in Reference
		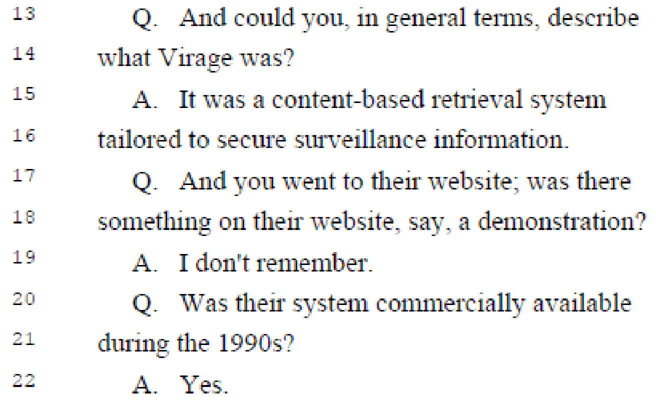 Q. And could you, in general terms, describe what Virage was? A. It was a content-based retrieval system tailored to secure surveillance information. Q. And you went to their website; was there something on their website, say, a demonstration? A. I don't remember. Q. Was their system commercially available during the 1990s? A. Yes. Flickner Depo at 95: 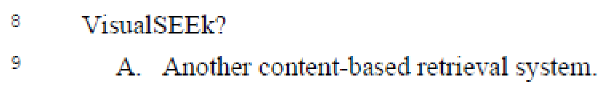 Q. And in general terms, what was VisualSEEk? A. Another content-based retrieval system. Flickner Depo at 113: 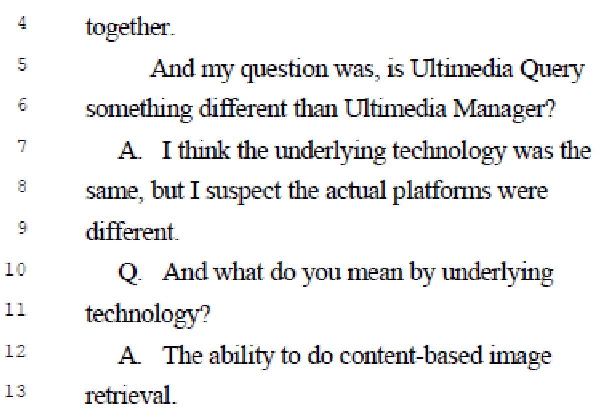 Q. I'm just trying to understand how Ultimedia Query, Ultimedia Manager and QBIC all fit together. And my question was, is Ultimedia Query something different than Ultimedia Manager? A. I think the underlying technology was the same, but I suspect the actual platforms were different. Q. And what do you mean by underlying technology? A. The ability to do content-based image retrieval.

Element	'036 Claim Recitation	Exemplary Citations in Reference
		Flickner Depo at 50: 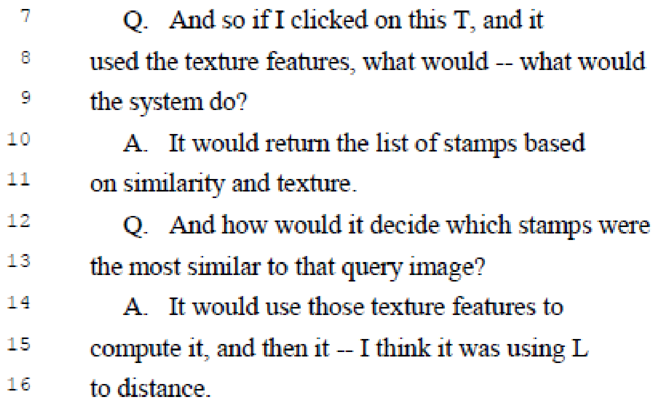 Q. And so if I clicked on this T, and it used the texture features, what would -- what would the system do? A. It would return the list of stamps based on similarity and texture. Q. And how would it decide which stamps were the most similar to that query image? A. It would use those texture features to compute it, and then it -- I think it was using L to distance. Flickner Depo at 54-55:

Element	'036 Claim Recitation	Exemplary Citations in Reference
		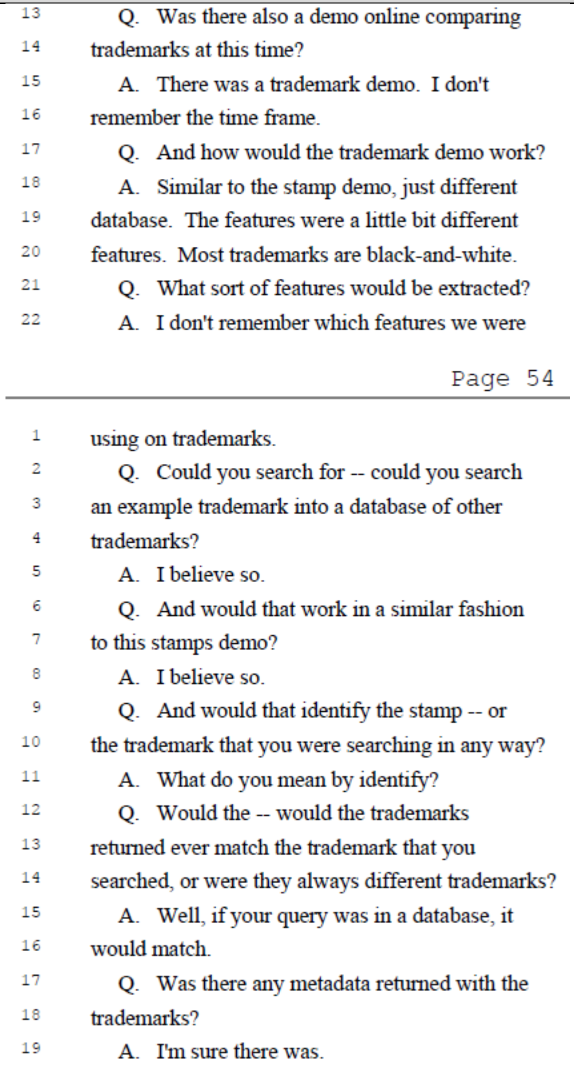 Q. Was there also a demo online comparing trademarks at this time? A. There was a trademark demo. I don't remember the time frame. Q. And how would the trademark demo work? A. Similar to the stamp demo, just different database. The features were a little bit different features. Most trademarks are black-and-white. Q. What sort of features would be extracted? A. I don't remember which features we were Page 54 using on trademarks. Q. Could you search for -- could you search an example trademark into a database of other trademarks? A. I believe so. Q. And would that work in a similar fashion to this stamps demo? A. I believe so. Q. And would that identify the stamp -- or the trademark that you were searching in any way? A. What do you mean by identify? Q. Would the -- would the trademarks returned ever match the trademark that you searched, or were they always different trademarks? A. Well, if your query was in a database, it would match. Q. Was there any metadata returned with the trademarks? A. I'm sure there was.

Element	'036 Claim Recitation	Exemplary Citations in Reference
		Flickner Depo at 89-90: 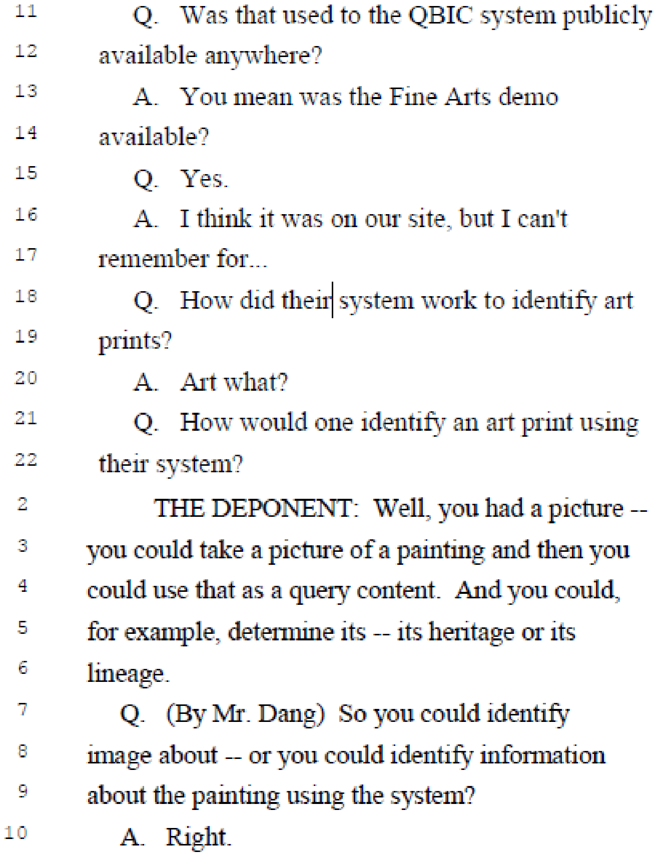 Q. Was that used to the QBIC system publicly available anywhere? A. You mean was the Fine Arts demo available? Q. Yes. A. I think it was on our site, but I can't remember for... Q. How did their system work to identify art prints? A. Art what? Q. How would one identify an art print using their system? THE DEPONENT: Well, you had a picture -- you could take a picture of a painting and then you could use that as a query content. And you could, for example, determine its -- its heritage or its lineage. Q. (By Mr. Dang) So you could identify image about -- or you could identify information about the painting using the system? A. Right. Flickner Depo at, e.g., 24, 41-42, 62-63 (user could take a picture or video and submit it to the QBIC system):

Element	'036 Claim Recitation	Exemplary Citations in Reference
		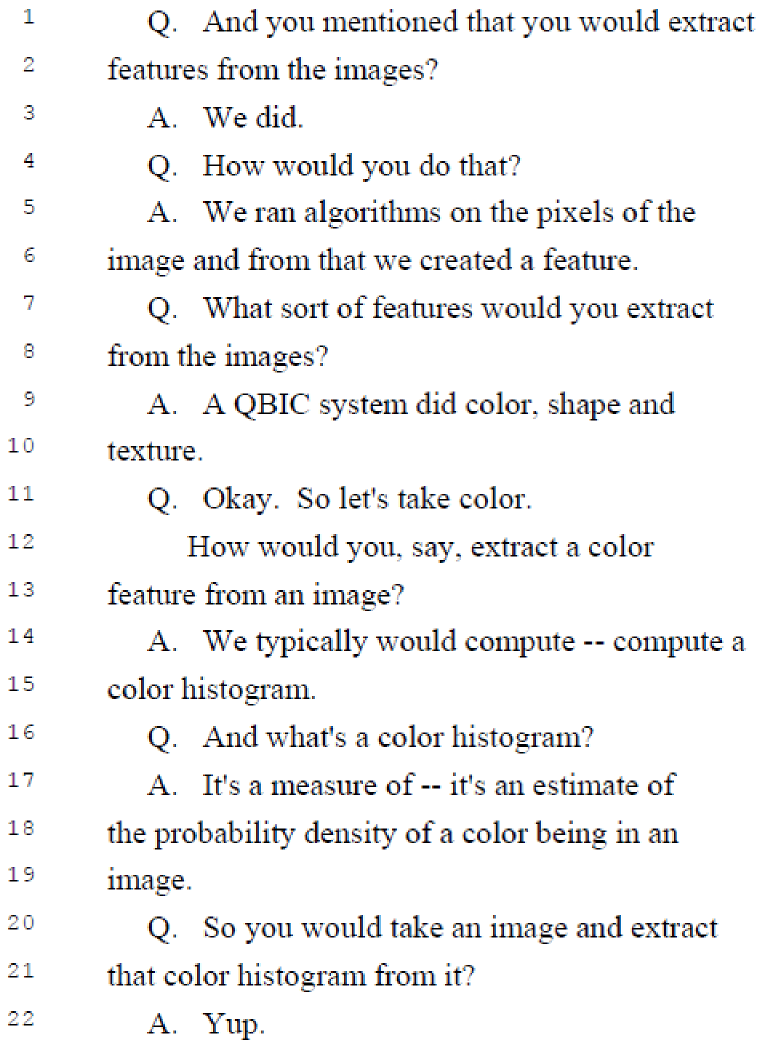 Q. And you mentioned that you would extract features from the images? A. We did. Q. How would you do that? A. We ran algorithms on the pixels of the image and from that we created a feature. Q. What sort of features would you extract from the images? A. A QBIC system did color, shape and texture. Q. Okay. So let's take color. How would you, say, extract a color feature from an image? A. We typically would compute -- compute a color histogram. Q. And what's a color histogram? A. It's a measure of -- it's an estimate of the probability density of a color being in an image. Q. So you would take an image and extract that color histogram from it? A. Yup.

		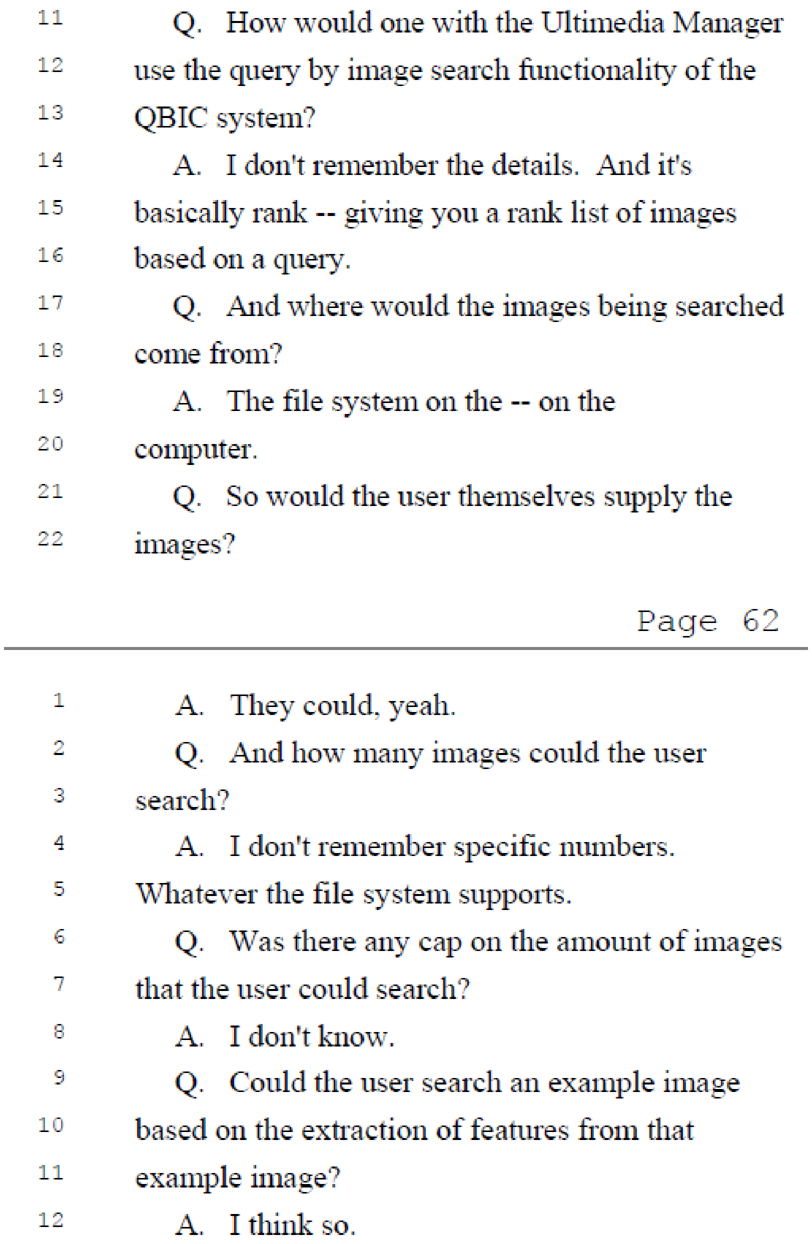 Q. How would one with the Ultimedia Manager use the query by image search functionality of the QBIC system? A. I don't remember the details. And it's basically rank -- giving you a rank list of images based on a query. Q. And where would the images being searched come from? A. The file system on the -- on the computer. Q. So would the user themselves supply the images? Page 62 A. They could, yeah. Q. And how many images could the user search? A. I don't remember specific numbers. Whatever the file system supports. Q. Was there any cap on the amount of images that the user could search? A. I don't know. Q. Could the user search an example image based on the extraction of features from that example image? A. I think so.

Element	'036 Claim Recitation	Exemplary Citations in Reference
		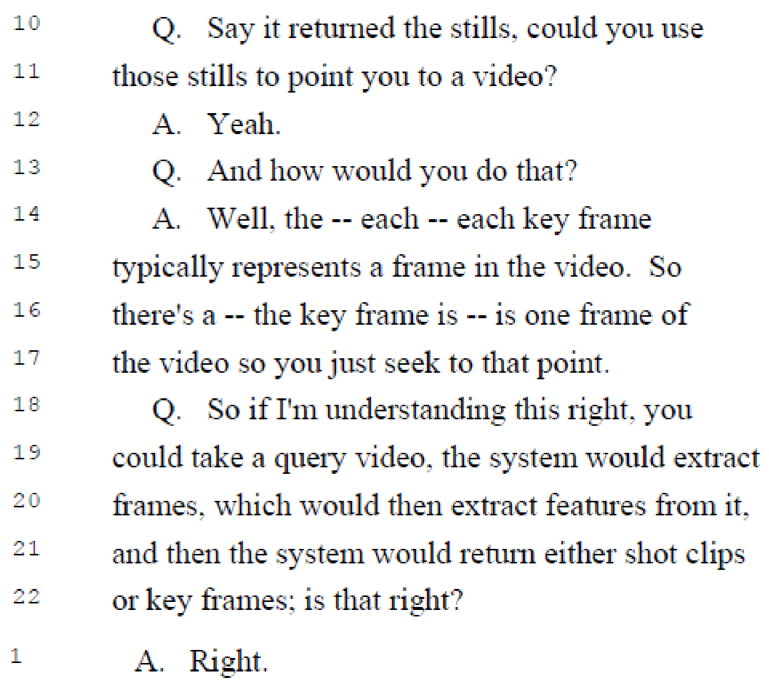 Q. Say it returned the stills, could you use those stills to point you to a video? A. Yeah. Q. And how would you do that? A. Well, the -- each -- each key frame typically represents a frame in the video. So there's a -- the key frame is -- is one frame of the video so you just seek to that point. Q. So if I'm understanding this right, you could take a query video, the system would extract frames, which would then extract features from it, and then the system would return either shot clips or key frames; is that right? A. Right. Flickner 8/29/23 Depo at 42:

Element	'036 Claim Recitation	Exemplary Citations in Reference
		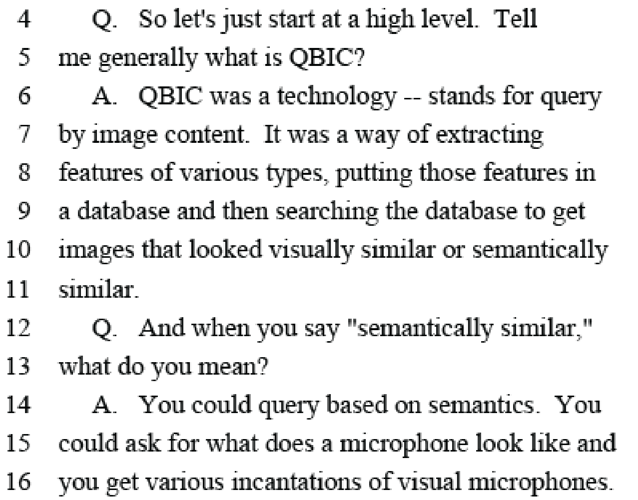 Q. So let's just start at a high level. Tell me generally what is QBIC? A. QBIC was a technology -- stands for query by image content. It was a way of extracting features of various types, putting those features in a database and then searching the database to get images that looked visually similar or semantically similar. Q. And when you say "semantically similar," what do you mean? A. You could query based on semantics. You could ask for what does a microphone look like and you get various incantations of visual microphones. *Id.* at 205–206:

Element	'036 Claim Recitation	Exemplary Citations in Reference
		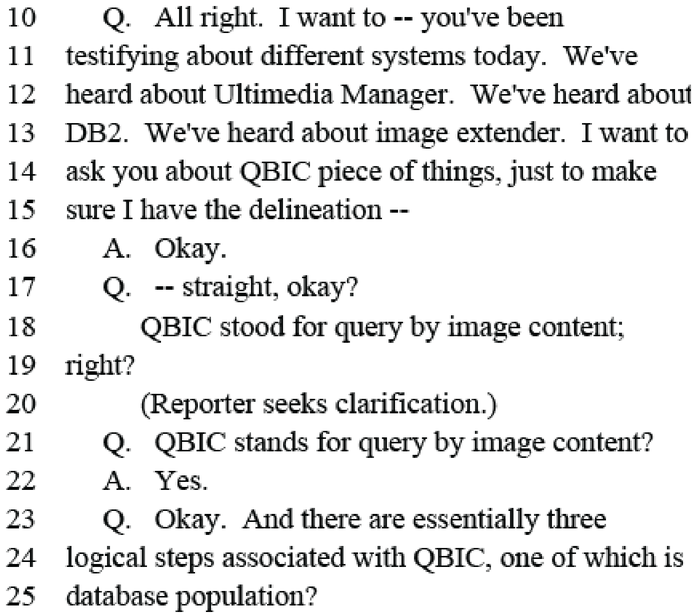 Q. All right. I want to -- you've been testifying about different systems today. We've heard about Ultimedia Manager. We've heard about DB2. We've heard about image extender. I want to ask you about QBIC piece of things, just to make sure I have the delineation -- A. Okay. Q. -- straight, okay? QBIC stood for query by image content; right? (Reporter seeks clarification.) Q. QBIC stands for query by image content? A. Yes. Q. Okay. And there are essentially three logical steps associated with QBIC, one of which is database population? 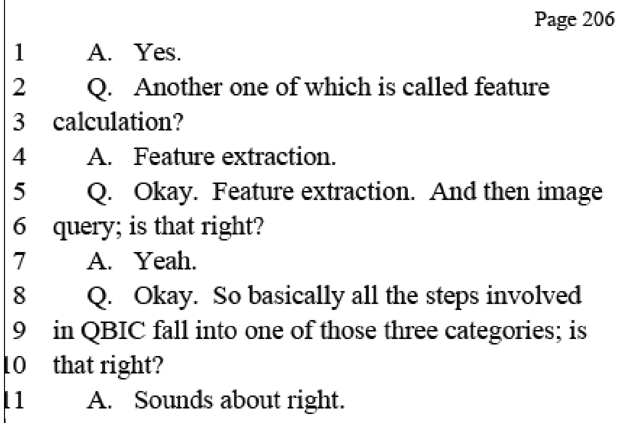 Page 206 A. Yes. Q. Another one of which is called feature calculation? A. Feature extraction. Q. Okay. Feature extraction. And then image query; is that right? A. Yeah. Q. Okay. So basically all the steps involved in QBIC fall into one of those three categories; is that right? A. Sounds about right. Flickner 8/29/23 Depo at Ex. 5:

Element	'036 Claim Recitation	Exemplary Citations in Reference
		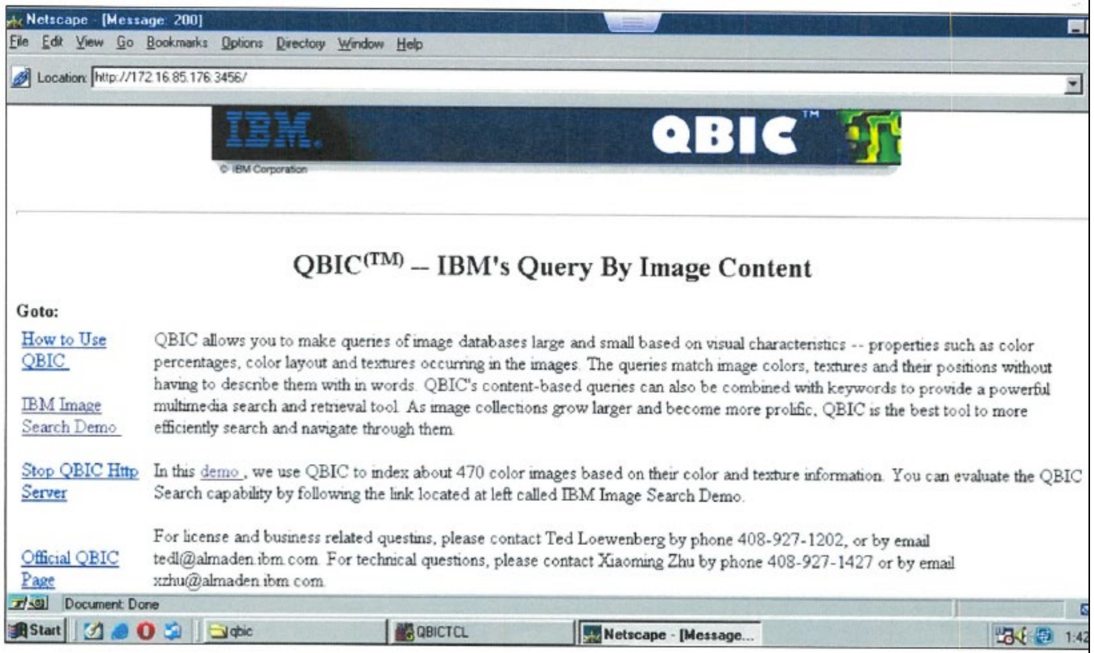 QBIC(TM) -- IBM's Query By Image Content Goto: How to Use QBIC IBM Image Search Demo Stop QBIC Http Server Official QBIC Page QBIC allows you to make queries of image databases large and small based on visual characteristics -- properties such as color percentages, color layout and textures occurring in the images. The queries match image colors, textures and their positions without having to describe them with in words. QBIC's content-based queries can also be combined with keywords to provide a powerful multimedia search and retrieval tool. As image collections grow larger and become more prolific, QBIC is the best tool to more efficiently search and navigate through them. In this demo, we use QBIC to index about 470 color images based on their color and texture information. You can evaluate the QBIC Search capability by following the link located at left called IBM Image Search Demo. For license and business related questins, please contact Ted Loewenberg by phone 408-927-1202, or by email tedl@almaden.ibm.com. For technical questions, please contact Xiaoming Zhu by phone 408-927-1427 or by email xzhu@almaden.ibm.com. The QBIC system was implemented in IBM commercial products, including in IBM's Ultimedia Manager, the DB2 series of products such as the Image Extender in DB2 Universal Database, and the Enterprise Information Portal. All of the features referenced in this chart were incorporated into these commercial implementations before November 6, 2000, as reflected in a combination of scientific papers, IBM's internal documentation, and the QBIC demo software applications. *See, e.g.*, Flickner 8/29/23 Depo Tr. at 33:9-24, 71:21-72:10, 164:21-166:2, 166:11-175:24, 178:5-186:17, 186:21-192:4, 192:14-194:1, 195:16-196:4, 194:1-308:1; IBM 000001-617; IBM 000747; IBM 000777-880; IBM 000893-902; IBM 0002315-320; IBM 0002390-404; IBM 0002418-430.
[1A]	a target database storing known targets of different types and recognition parameters associated with the known targets;	QBIC discloses a target database storing known targets and their associated recognition parameters. The QBIC Project: Querysing Images by Content Using Color, Texture, and Shape

Element	'036 Claim Recitation	Exemplary Citations in Reference
		(IBM 0002390–404) at 2390 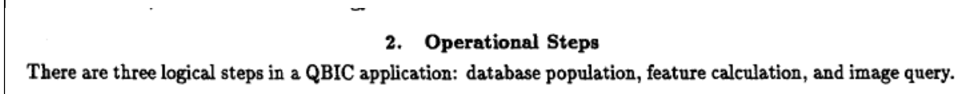 **2. Operational Steps** There are three logical steps in a QBIC application: database population, feature calculation, and image query. IBM0002390 at 2391: 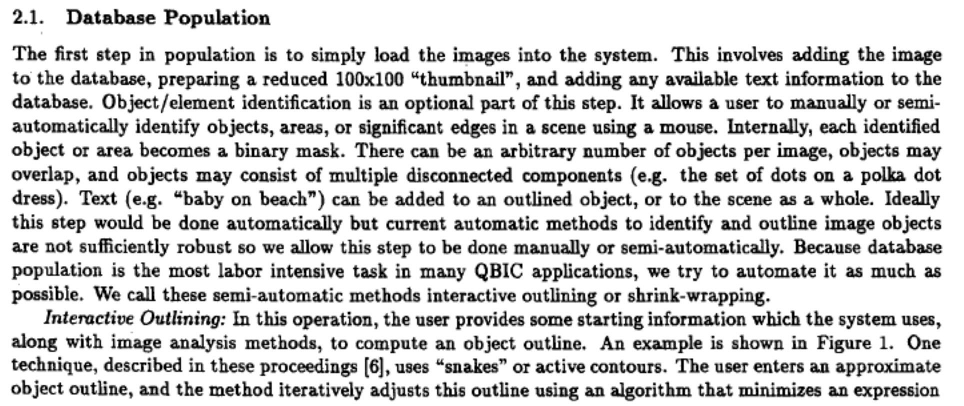 **2.1. Database Population** The first step in population is to simply load the images into the system. This involves adding the image to the database, preparing a reduced 100x100 "thumbnail", and adding any available text information to the database. Object/element identification is an optional part of this step. It allows a user to manually or semi-automatically identify objects, areas, or significant edges in a scene using a mouse. Internally, each identified object or area becomes a binary mask. There can be an arbitrary number of objects per image, objects may overlap, and objects may consist of multiple disconnected components (e.g. the set of dots on a polka dot dress). Text (e.g. "baby on beach") can be added to an outlined object, or to the scene as a whole. Ideally this step would be done automatically but current automatic methods to identify and outline image objects are not sufficiently robust so we allow this step to be done manually or semi-automatically. Because database population is the most labor intensive task in many QBIC applications, we try to automate it as much as possible. We call these semi-automatic methods interactive outlining or shrink-wrapping. *Interactive Outlining:* In this operation, the user provides some starting information which the system uses, along with image analysis methods, to compute an object outline. An example is shown in Figure 1. One technique, described in these proceedings [6], uses "snakes" or active contours. The user enters an approximate object outline, and the method iteratively adjusts this outline using an algorithm that minimizes an expression

Element	'036 Claim Recitation	Exemplary Citations in Reference
		containing two main terms: total image edge strength along the outline, and the outline curvature. Another method, flood fill, can take as starting information a single mouse-clicked image point. Starting from this object point, all adjacent pixels are included in the object whose color is within a specified color distance from the original point. The perimeter of all included points becomes the object boundary. For reasonably uniform objects that are distinct from the background, this operation provides a fast object outliner. In interactive outlining, as in other parts of our system, we have attempted to build the functions in such a way as to carefully distribute the tasks between the human and machine. For example, the tasks of giving the starting outline for an object or of attaching semantic meaning to objects is left to the user since they are difficult for machines. On the other hand, computation of quantitative image/object features is ideally suited for machines and so is done automatically.  Figure 1: Original image, approximate user outline, and automatically refined outline of sun. IBM0002390 at 2391-92:

Element	'036 Claim Recitation	Exemplary Citations in Reference
		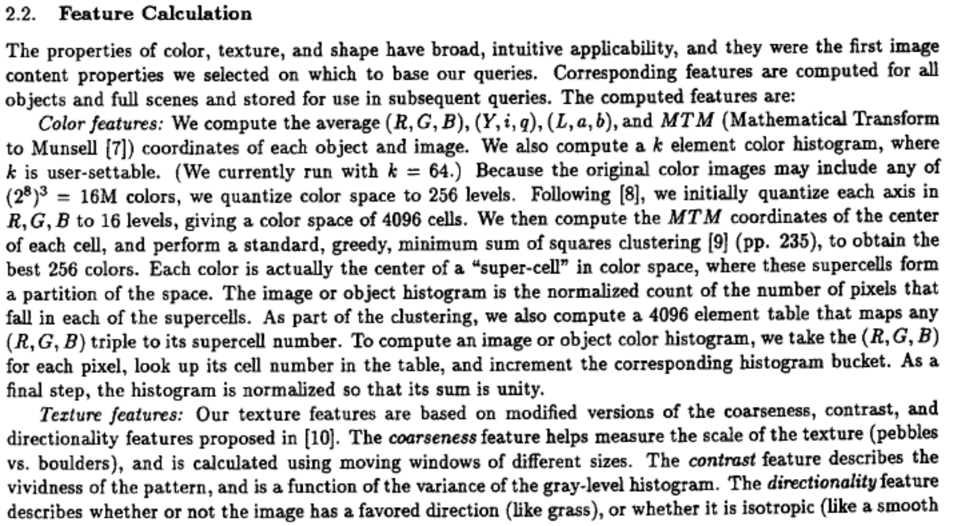 2.2. **Feature Calculation** The properties of color, texture, and shape have broad, intuitive applicability, and they were the first image content properties we selected on which to base our queries. Corresponding features are computed for all objects and full scenes and stored for use in subsequent queries. The computed features are: *Color features:* We compute the average (R, G, B), (Y, i, q), (L, a, b), and MTM (Mathematical Transform to Munsell [7]) coordinates of each object and image. We also compute a k element color histogram, where k is user-settable. (We currently run with $k = 64$.) Because the original color images may include any of $(2^8)^3 = 16$M colors, we quantize color space to 256 levels. Following [8], we initially quantize each axis in R, G, B to 16 levels, giving a color space of 4096 cells. We then compute the MTM coordinates of the center of each cell, and perform a standard, greedy, minimum sum of squares clustering [9] (pp. 235), to obtain the best 256 colors. Each color is actually the center of a "super-cell" in color space, where these supercells form a partition of the space. The image or object histogram is the normalized count of the number of pixels that fall in each of the supercells. As part of the clustering, we also compute a 4096 element table that maps any (R, G, B) triple to its supercell number. To compute an image or object color histogram, we take the (R, G, B) for each pixel, look up its cell number in the table, and increment the corresponding histogram bucket. As a final step, the histogram is normalized so that its sum is unity. *Texture features:* Our texture features are based on modified versions of the coarseness, contrast, and directionality features proposed in [10]. The *coarseness* feature helps measure the scale of the texture (pebbles vs. boulders), and is calculated using moving windows of different sizes. The *contrast* feature describes the vividness of the pattern, and is a function of the variance of the gray-level histogram. The *directionality* feature describes whether or not the image has a favored direction (like grass), or whether it is isotropic (like a smooth

Element	'036 Claim Recitation	Exemplary Citations in Reference
		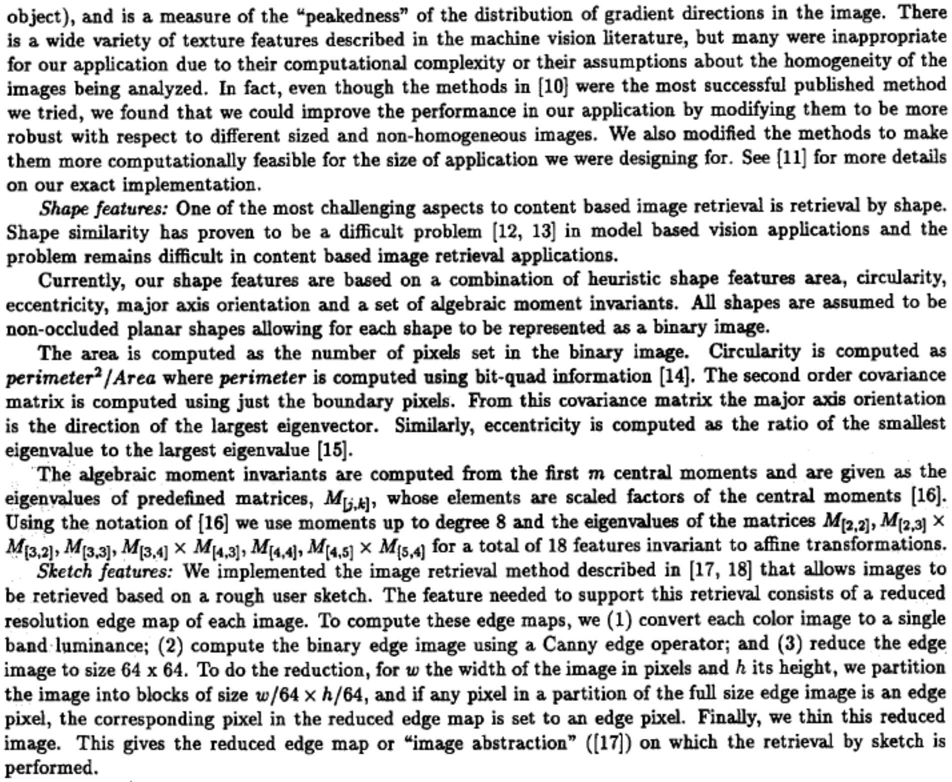 object), and is a measure of the "peakedness" of the distribution of gradient directions in the image. There is a wide variety of texture features described in the machine vision literature, but many were inappropriate for our application due to their computational complexity or their assumptions about the homogeneity of the images being analyzed. In fact, even though the methods in [10] were the most successful published method we tried, we found that we could improve the performance in our application by modifying them to be more robust with respect to different sized and non-homogeneous images. We also modified the methods to make them more computationally feasible for the size of application we were designing for. See [11] for more details on our exact implementation. *Shape features:* One of the most challenging aspects to content based image retrieval is retrieval by shape. Shape similarity has proven to be a difficult problem [12, 13] in model based vision applications and the problem remains difficult in content based image retrieval applications. Currently, our shape features are based on a combination of heuristic shape features area, circularity, eccentricity, major axis orientation and a set of algebraic moment invariants. All shapes are assumed to be non-occluded planar shapes allowing for each shape to be represented as a binary image. The area is computed as the number of pixels set in the binary image. Circularity is computed as $perimeter^2/Area$ where *perimeter* is computed using bit-quad information [14]. The second order covariance matrix is computed using just the boundary pixels. From this covariance matrix the major axis orientation is the direction of the largest eigenvector. Similarly, eccentricity is computed as the ratio of the smallest eigenvalue to the largest eigenvalue [15]. The algebraic moment invariants are computed from the first m central moments and are given as the eigenvalues of predefined matrices, $M_{[j,k]}$, whose elements are scaled factors of the central moments [16]. Using the notation of [16] we use moments up to degree 8 and the eigenvalues of the matrices $M_{[2,2]}, M_{[2,3]} \times M_{[3,2]}, M_{[3,3]}, M_{[3,4]} \times M_{[4,3]}, M_{[4,4]}, M_{[4,5]} \times M_{[5,4]}$ for a total of 18 features invariant to affine transformations. *Sketch features:* We implemented the image retrieval method described in [17, 18] that allows images to be retrieved based on a rough user sketch. The feature needed to support this retrieval consists of a reduced resolution edge map of each image. To compute these edge maps, we (1) convert each color image to a single band luminance; (2) compute the binary edge image using a Canny edge operator; and (3) reduce the edge image to size 64 x 64. To do the reduction, for w the width of the image in pixels and h its height, we partition the image into blocks of size $w/64 \times h/64$, and if any pixel in a partition of the full size edge image is an edge pixel, the corresponding pixel in the reduced edge map is set to an edge pixel. Finally, we thin this reduced image. This gives the reduced edge map or "image abstraction" ([17]) on which the retrieval by sketch is performed. IBM0002390 at 2392-93:

Element	'036 Claim Recitation	Exemplary Citations in Reference
		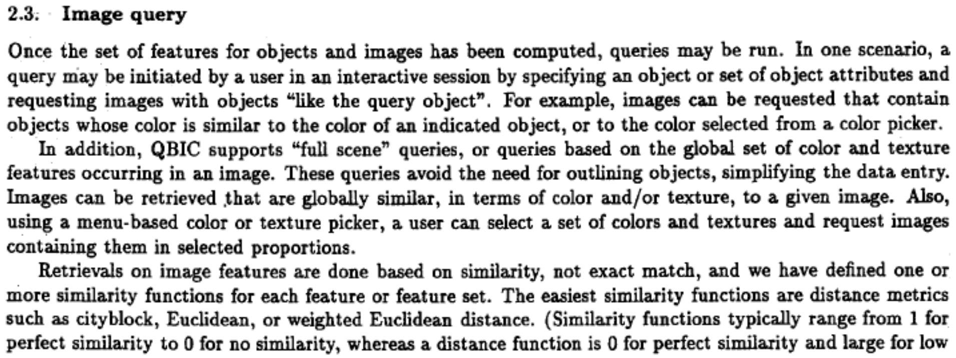 **2.3. Image query** Once the set of features for objects and images has been computed, queries may be run. In one scenario, a query may be initiated by a user in an interactive session by specifying an object or set of object attributes and requesting images with objects "like the query object". For example, images can be requested that contain objects whose color is similar to the color of an indicated object, or to the color selected from a color picker. In addition, QBIC supports "full scene" queries, or queries based on the global set of color and texture features occurring in an image. These queries avoid the need for outlining objects, simplifying the data entry. Images can be retrieved that are globally similar, in terms of color and/or texture, to a given image. Also, using a menu-based color or texture picker, a user can select a set of colors and textures and request images containing them in selected proportions. Retrievals on image features are done based on similarity, not exact match, and we have defined one or more similarity functions for each feature or feature set. The easiest similarity functions are distance metrics such as cityblock, Euclidean, or weighted Euclidean distance. (Similarity functions typically range from 1 for perfect similarity to 0 for no similarity, whereas a distance function is 0 for perfect similarity and large for low

Element	'036 Claim Recitation	Exemplary Citations in Reference
		similarity, so we must convert as necessary to obtain consistent measures.) Different similarity measures may be used, and they may vary from application to application. The matching methods and similarity functions for the different feature types are described below: *Color:* For average color, the distance between a query object and database object is weighted Euclidean distance. The weights are user adjustable so that we may, for example, downweight the L component of an (L, a, b) query. Best results, as judged by our subjective evaluation in on-going experiments, are obtained by representing the average color in the MTM color space, and by inversely weighting each component by its standard deviation over the samples in the database. Improved results are obtained when color matching is done not on average color, but on the distribution of colors that occur in an object or image. In this way, a red and blue object matches other red and blue objects better than to a purple object. To match color histograms, we use a method similar to the one described in [8]. Let X be the query histogram and Y the histogram of an item in the database, both normalized. We compute the element by element difference histogram Z. Then the similarity between X and Y is given by $\lVert Z \rVert = Z^T A Z$, where A is a symmetric color similarity matrix with $a(i,j) \equiv 1 - d(c_i, c_j)/d_{max}$ where c_i and c_j are the ith and jth colors in the color histograms, and $d(c_i, c_j)$ is the MTM color distance, and d_{max} is the maximum distance between any two colors ([8] used Luv distance.) This metric gives the (weighted) length of the difference vector between X and Y, weighted by A, which accounts for the perceptual distance between different pairs of colors. The result is a measure that accounts for both the difference between the amounts in any given color (e.g. a particular shade of red), as well as differences between similar colors (e.g. red and orange). Based on our experience, we also found it useful for the user to be able to request images with $x\%$ of color 1, $y\%$ of color 2, etc. where both the percentages and the colors are explicitly specified by the user. This allows searches such as for a beach scene (25% white (for the sand) and 50% blue (for the sky and water)), and for many scenes for which the user has seen before, and for which he/she can remember approximate color percentages. The color picker that allows this type of query is shown in Figure 2. 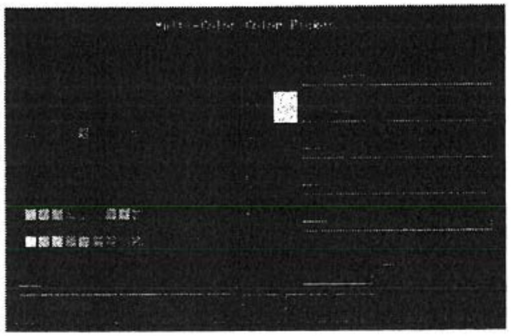 Figure 2: Multi-Color Color Picker with two selected colors. The palette on the left displays the set of selectable colors. Once a color is selected, it appears in a box on the right, and the amount of the color can be specified by the sliders. In this case, 25% light blue and 15% white have been selected. *Texture:* Texture distance is computed as weighted Euclidean distance in the three dimensional texture space. The most common weighting factors are the inverse variances for each component, computed over the

Element	'036 Claim Recitation	Exemplary Citations in Reference
		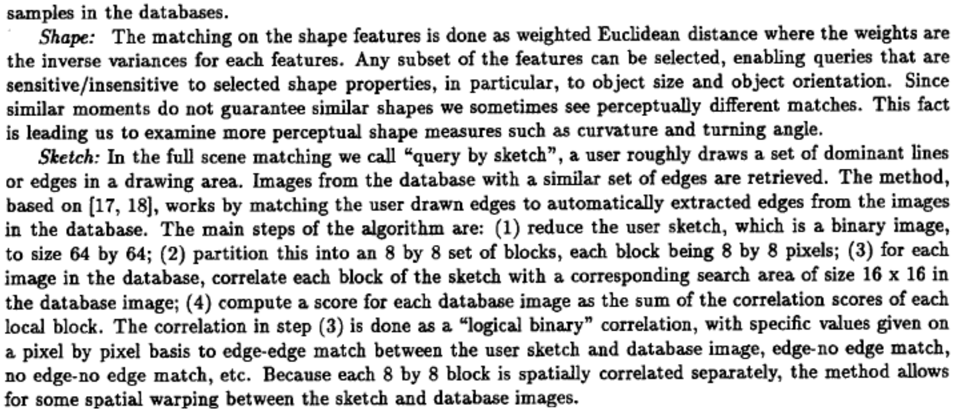 samples in the databases. *Shape:* The matching on the shape features is done as weighted Euclidean distance where the weights are the inverse variances for each features. Any subset of the features can be selected, enabling queries that are sensitive/insensitive to selected shape properties, in particular, to object size and object orientation. Since similar moments do not guarantee similar shapes we sometimes see perceptually different matches. This fact is leading us to examine more perceptual shape measures such as curvature and turning angle. *Sketch:* In the full scene matching we call "query by sketch", a user roughly draws a set of dominant lines or edges in a drawing area. Images from the database with a similar set of edges are retrieved. The method, based on [17, 18], works by matching the user drawn edges to automatically extracted edges from the images in the database. The main steps of the algorithm are: (1) reduce the user sketch, which is a binary image, to size 64 by 64; (2) partition this into an 8 by 8 set of blocks, each block being 8 by 8 pixels; (3) for each image in the database, correlate each block of the sketch with a corresponding search area of size 16 x 16 in the database image; (4) compute a score for each database image as the sum of the correlation scores of each local block. The correlation in step (3) is done as a "logical binary" correlation, with specific values given on a pixel by pixel basis to edge-edge match between the user sketch and database image, edge-no edge match, no edge-no edge match, etc. Because each 8 by 8 block is spatially correlated separately, the method allows for some spatial warping between the sketch and database images. IBM0002390 at 2393-95: QBIC (Query by Image Content) (BOFA00001391–392) at 1:

Element	'036 Claim Recitation	Exemplary Citations in Reference
		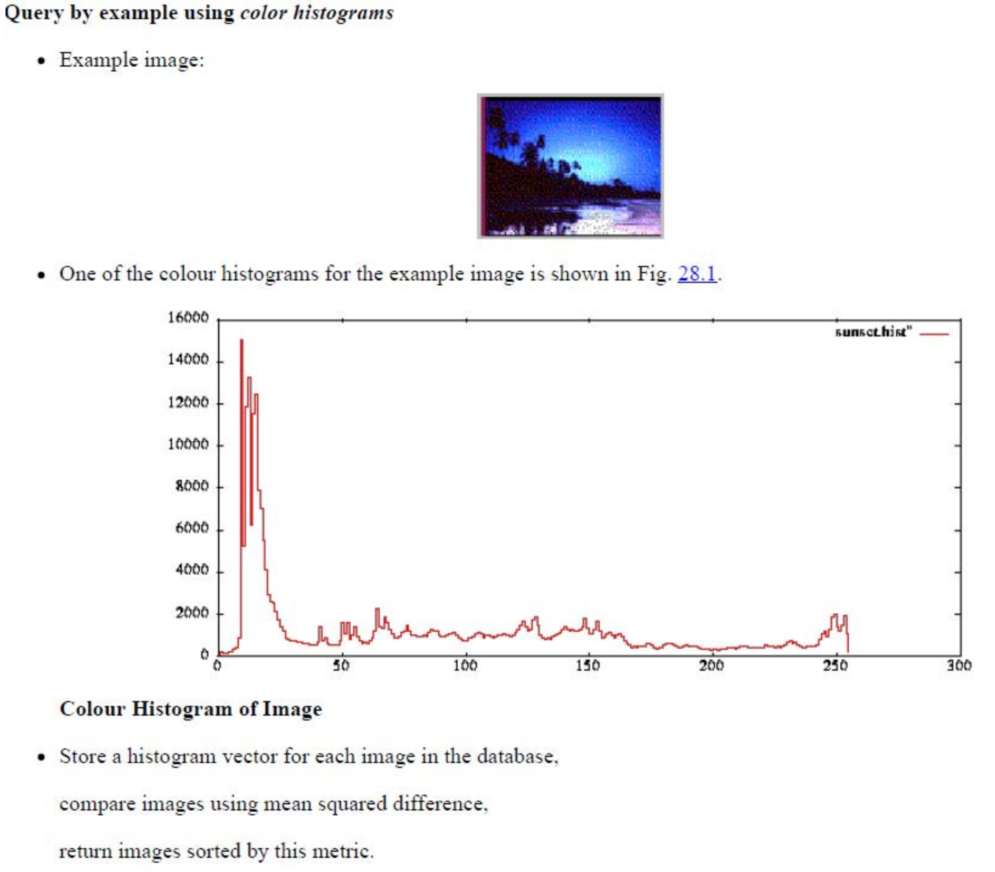 **Query by example using *color histograms*** - Example image: - One of the colour histograms for the example image is shown in Fig. 28.1. **Colour Histogram of Image** - Store a histogram vector for each image in the database, compare images using mean squared difference, return images sorted by this metric. Query by Image and Video Content: The QBIC System (IBM 0000893–902): "To achieve this functionality, QBIC has two main components: database population (the process of creating an image database) and database query. During the population, images and videos are processed to extract features describing their content-colors, textures, shapes, and camera and object motion-and the features are stored in a database. During the query, the user composes a query graphically. Features are generated from the graphical query and then input to a matching

Element	'036 Claim Recitation	Exemplary Citations in Reference
		engine that finds images or videos from the database with similar features. Figure 2 shows the system architecture." *Id.* at p. 3. 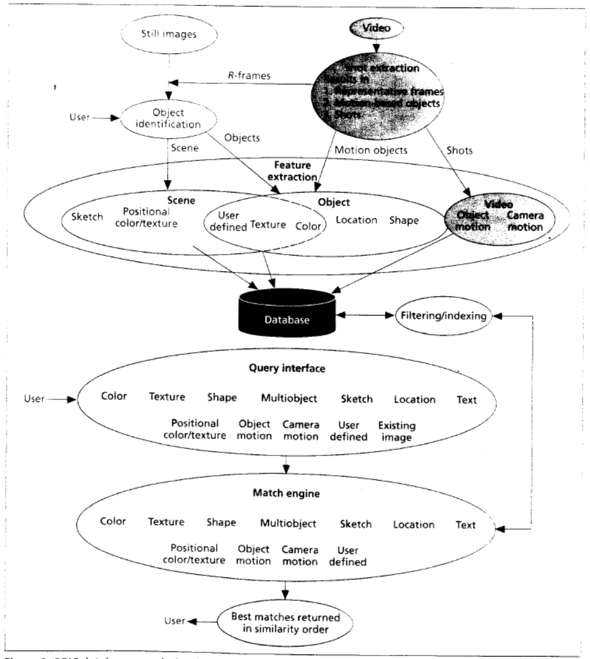 Figure 2. QBIC database population (top) and query (bottom) architecture. *Id*. at Fig. 2. "For both population and query, the QBIC data model has • still images or scenes (full images) that contain objects (subsets of an image), and • video shots that consist of sets of contiguous frames and contain motion objects. For still images, the QBIC data model distinguishes between 'scenes' (or images) and 'objects.' A scene is an image or single representative frame of video. An object is a part of a scene-for example, the fox in Figure 3-or a moving entity in a video. For still image database population, features are extracted from images and objects and stored in a database as shown in the top left part of Figure 2. Videos are broken into clips called shots. Representative frames, or r-frames, are

Element	'036 Claim Recitation	Exemplary Citations in Reference
		generated for each extracted shot. R-frames are treated as still images, and features are extracted and stored in the database. Further processing of shots generates motion objects-for example, a car moving across the screen. Queries are allowed on objects ('Find images with a red, round object'), scenes ('Find images that have approximately 30-percent red and IS-percent blue colors'), shots ('Find all shots panning from left to right'), or any combination ('Find images that have 30 percent red and contain a blue textured object')." *Id.* at p. 3. "**DATABASE POPULATION** In still image database population, the images are reduced to a standard-sized icon called a thumbnail and annotated with any available text information. Object identification is an optional but key part of this step. It lets users manually, semiautomatically, or fully automatically identify interesting regions—which we call objects—in the images. Internally, each object is represented as a binary mask. There may be an arbitrary number of objects per image. Objects can overlap and can consist of multiple disconnected components like the set of dots on a polka-dot dress. Text, like 'baby on beach,' can be associated with an outlined object or with the scene as a whole. **Object-outlining tools** Ideally, object identification would be automatic, but this is generally difficult. The alternative—manual identification—is tedious and can inhibit query-by-content applications. As a result, we have devoted considerable effort to developing tools to aid in this step. In recent work, we have successfully used fully automatic unsupervised segmentation methods along with a foreground/background model to identify objects in a restricted class of images. The images, typical of museums and retail catalogs, have a small number of foreground objects on a generally separable background. Figure 4 shows example results. Even in this domain, robust algorithms are required because of the textured and variegated backgrounds. We also provide semiautomatic tools for identifying objects. One is an enhanced flood-fill technique. Flood-fill methods, found in most photo-editing programs, start from a single object pixel and repeatedly add adjacent pixels whose values

Element	'036 Claim Recitation	Exemplary Citations in Reference
		are within some given threshold of the original pixel. Selecting the threshold, which must change from image to image and object to object, is tedious. We automatically calculate a dynamic threshold by having the user click on background as well as object points. For reasonably uniform objects that are distinct from the background, this operation allows fast object identification without manually adjusting a threshold. The example in Figure 3 shows an object, a fox, identified by using only a few clicks. We designed another outlining tool to help users track object edges. This tool takes a user-drawn curve and automatically aligns it with nearby image edges. Based on the 'snakes' concept developed in recent computer vision research, the tool finds the curve that maximizes the image gradient magnitude along the curve. The spline snake formulation we use allows for smooth solutions to the resulting nonlinear minimization problem. The computation is done at interactive speeds so that, as the user draws a curve, it is 'rubber-banded' to lie along object boundaries. . . . **Video data** For video data, database population has three major components: • shot detection, • representative frame creation for each shot, and • derivation of a layered representation of coherently moving structures/ objects. . . . Once the shot boundaries have been detected, each shot is represented using an r-frame. R-frames are used for several purposes. First, during database population, r-frames are treated as still images in which objects can be identified by using the previously described methods." *Id.* at pp. 3–6. Image, Audio, and Video Extenders Administration and Programming (IBM 000001–617): "Your applications can even search for images by content. Imagine an application that uses visual examples to search for images. With such an application, users could select an example image and have the application find other images that have colors or textures similar to those in the example. With DB2 extenders' Query by Image Content (QBIC) capability, you can create applications that search for images in this visual way." *Id.* at p. 4.

Element	'036 Claim Recitation	Exemplary Citations in Reference
		"A QBIC catalog is a set of files that hold data about the visual features of images. The Image Extender uses this data to search for images by content. You create a QBIC catalog for each column of images in a user table that you want to make available for searching by content. When you create a QBIC catalog you identify the features for which you want the Image Extender to analyze, store, and later query data. You can also add or drop features from a QBIC catalog after the catalog is created. A QBIC catalog can hold data for the following image features: **Average color** The sum of the color values for all pixels in an image divided by the number of pixels in the image. (A pixel is the smallest element of an image that can be assigned color and intensity.) For example, if 50% of an image consists of blue pixels and the other 50% red pixels, the image has an average color value of purple. Average color is used to search for images that have a predominant color. If an image has a predominant color, the average color will be similar to the predominant color. **Histogram color** Measures the distribution of colors in an image against a spectrum of 64 colors. For each of the 64 colors, histogram color identifies the percentage of pixels in an image that have that color. For example, the histogram color of an image might be 40% white pixels, 50% blue, and 10% red; none of the pixels in the image have any of the remaining colors in the histogram spectrum. Histogram color is used to search for images that have a variety of colors. **Positional color** The average color value for the pixels in a specified area in an image. For example, the upper right-hand corner of an image might show a bright yellow sun; the positional color of this area of the image is bright yellow. Positional color is used to search for images that have a predominant color in a particular area. **Texture** Measures the coarseness, contrast, and directionality of an image. Coarseness indicates the size of repeating items in an image (for example, pebbles versus

Element	’036 Claim Recitation	Exemplary Citations in Reference
		boulders). Contrast identifies the brightness variations in an image (light versus dark). Directionality indicates whether a direction predominates in an image (as in the vertical direction of a picket fence) or does not predominate (as in an image of sand). Texture is used to search for images that have a particular pattern. To make an image available for searching by content, you catalog the image. When you catalog an image, the Image Extender analyzes the image, by computing the feature values for the image, and stores the values in a QBIC catalog. When you search for an image by content, your query identifies one or more features for the search (such as average color), a source for each feature (such as an example image), and a target set of cataloged images. The Image Extender computes the feature value of the source and compares it to the cataloged feature values for the target images. It then computes a score that indicates how similar the feature values of the target images are to the source. You can have the Image Extender return the images whose features are most similar to the source. The Image Extender will return the handle of each image and the image score. You can also have the Image Extender return only the score of a single image.” *Id.* at pp. 20–21. QBIC Demo (IBM000747):

Element	'036 Claim Recitation	Exemplary Citations in Reference
		 QBIC Demo (IBM000747): 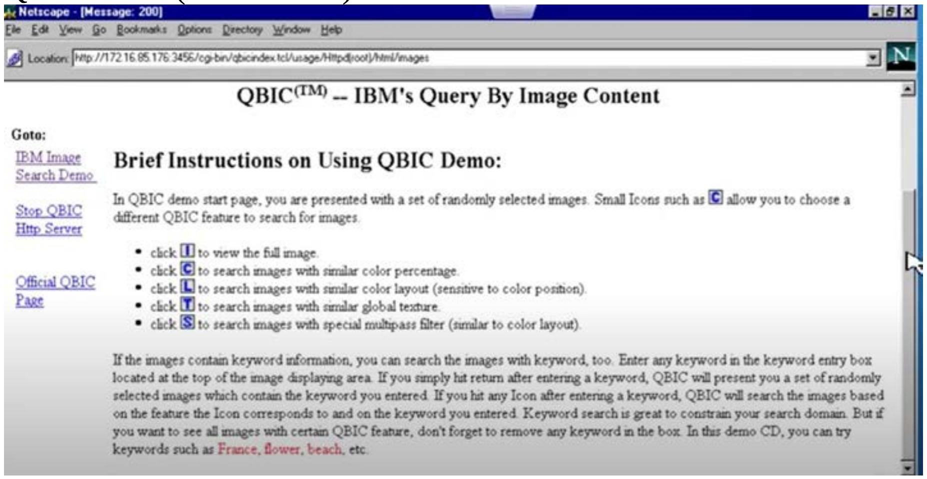 QBIC(TM) -- IBM's Query By Image Content **Brief Instructions on Using QBIC Demo:** In QBIC demo start page, you are presented with a set of randomly selected images. Small Icons such as C allow you to choose a different QBIC feature to search for images. • click I to view the full image. • click C to search images with similar color percentage. • click L to search images with similar color layout (sensitive to color position). • click T to search images with similar global texture. • click S to search images with special multipass filter (similar to color layout). If the images contain keyword information, you can search the images with keyword, too. Enter any keyword in the keyword entry box located at the top of the image displaying area. If you simply hit return after entering a keyword, QBIC will present you a set of randomly selected images which contain the keyword you entered. If you hit any Icon after entering a keyword, QBIC will search the images based on the feature the Icon corresponds to and on the keyword you entered. Keyword search is great to constrain your search domain. But if you want to see all images with certain QBIC feature, don't forget to remove any keyword in the box. In this demo CD, you can try keywords such as France, flower, beach, etc.

Element	'036 Claim Recitation	Exemplary Citations in Reference
		Flickner Depo at 20 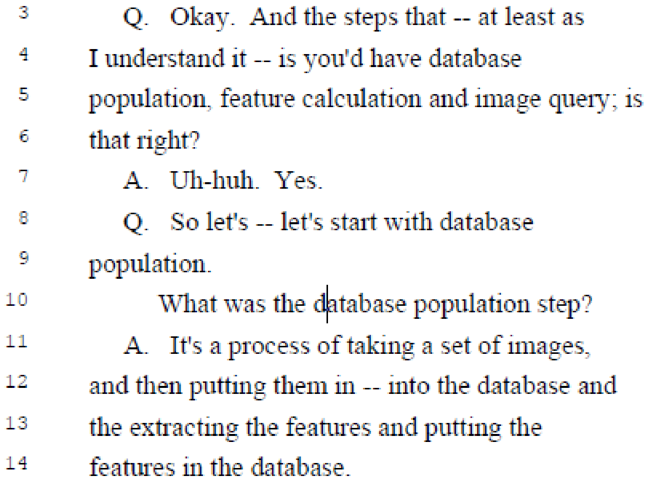 Q. Okay. And the steps that -- at least as I understand it -- is you'd have database population, feature calculation and image query; is that right? A. Uh-huh. Yes. Q. So let's -- let's start with database population. What was the database population step? A. It's a process of taking a set of images, and then putting them in -- into the database and the extracting the features and putting the features in the database. 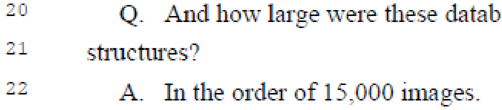 Q. And how large were these database structures? A. In the order of 15,000 images. Flickner Depo at 25-26:

Element	'036 Claim Recitation	Exemplary Citations in Reference
		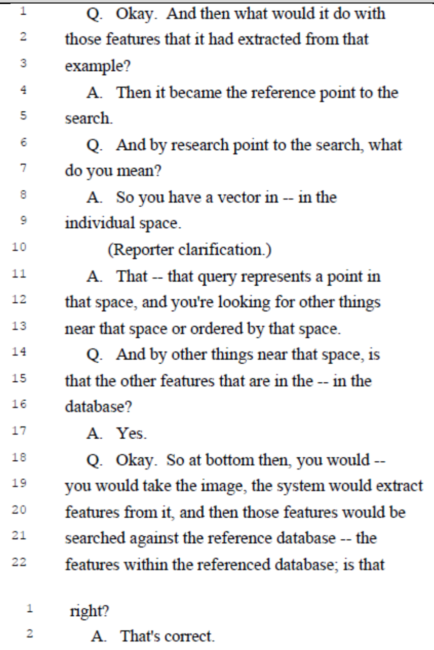 Q. Okay. And then what would it do with those features that it had extracted from that example? A. Then it became the reference point to the search. Q. And by research point to the search, what do you mean? A. So you have a vector in -- in the individual space. (Reporter clarification.) A. That -- that query represents a point in that space, and you're looking for other things near that space or ordered by that space. Q. And by other things near that space, is that the other features that are in the -- in the database? A. Yes. Q. Okay. So at bottom then, you would -- you would take the image, the system would extract features from it, and then those features would be searched against the reference database -- the features within the referenced database; is that right? A. That's correct. Flickner Depo at 36: 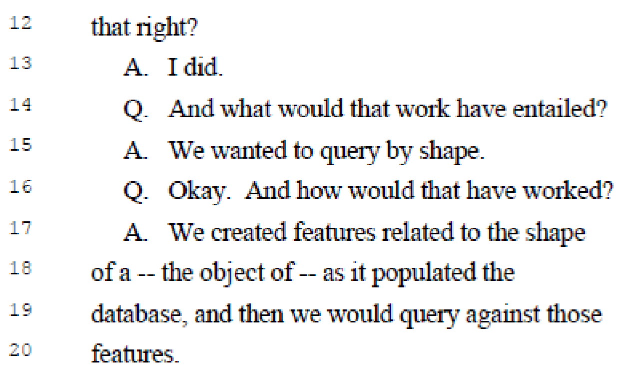 Q. So underlying this paper, you mentioned you did more work on the shape side of things; is that right? A. I did. Q. And what would that work have entailed? A. We wanted to query by shape. Q. Okay. And how would that have worked? A. We created features related to the shape of a -- the object of -- as it populated the database, and then we would query against those features.

Element	'036 Claim Recitation	Exemplary Citations in Reference
		Flickner Depo at 39-40: 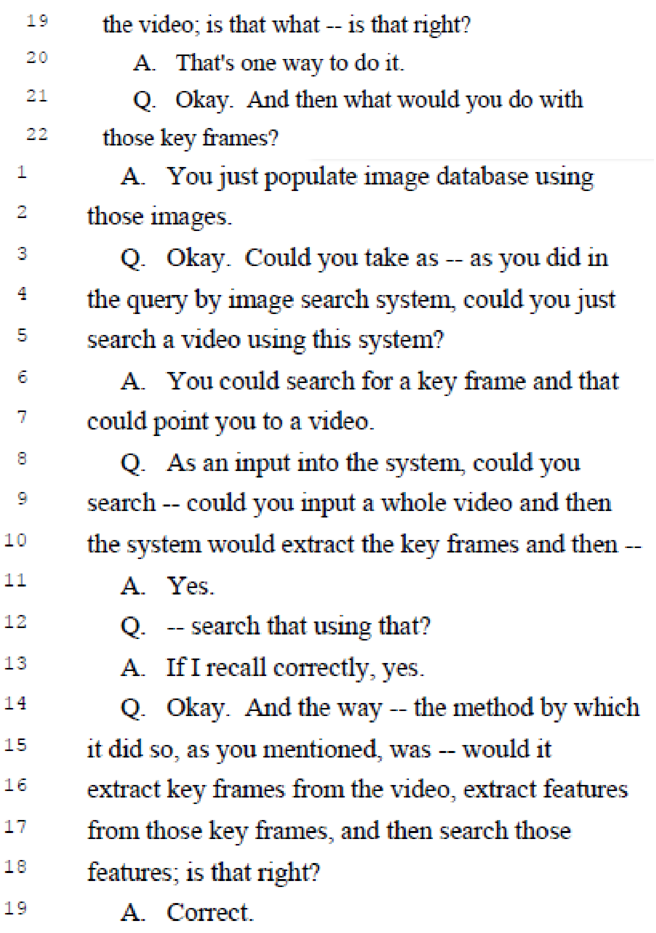 Q. Okay. So you would take -- you would use a histogram methodology to extract key frames from the video; is that what -- is that right? A. That's one way to do it. Q. Okay. And then what would you do with those key frames? A. You just populate image database using those images. Q. Okay. Could you take as -- as you did in the query by image search system, could you just search a video using this system? A. You could search for a key frame and that could point you to a video. Q. As an input into the system, could you search -- could you input a whole video and then the system would extract the key frames and then -- A. Yes. Q. -- search that using that? A. If I recall correctly, yes. Q. Okay. And the way -- the method by which it did so, as you mentioned, was -- would it extract key frames from the video, extract features from those key frames, and then search those features; is that right? A. Correct.

Element	'036 Claim Recitation	Exemplary Citations in Reference
		Flickner Depo at 78-79: 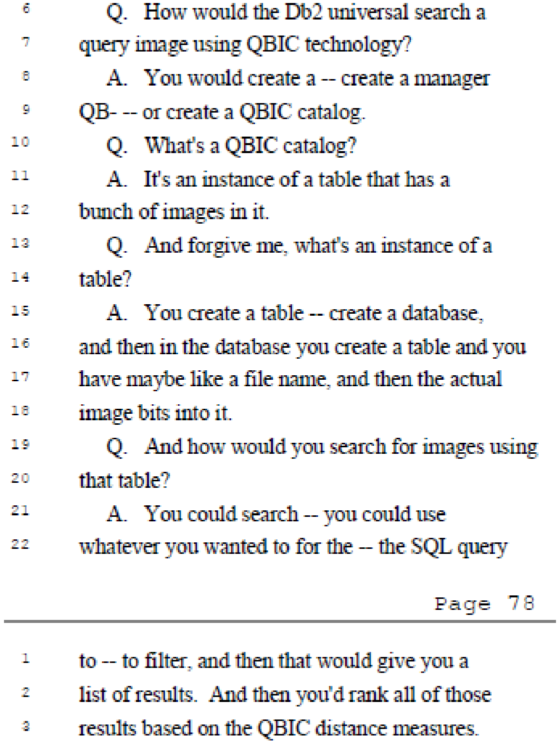 Q. How would the Db2 universal search a query image using QBIC technology? A. You would create a -- create a manager QB- -- or create a QBIC catalog. Q. What's a QBIC catalog? A. It's an instance of a table that has a bunch of images in it. Q. And forgive me, what's an instance of a table? A. You create a table -- create a database, and then in the database you create a table and you have maybe like a file name, and then the actual image bits into it. Q. And how would you search for images using that table? A. You could search -- you could use whatever you wanted to for the -- the SQL query Page 78 to -- to filter, and then that would give you a list of results. And then you'd rank all of those results based on the QBIC distance measures. Flickner Depo at 88-89:

Element	'036 Claim Recitation	Exemplary Citations in Reference
		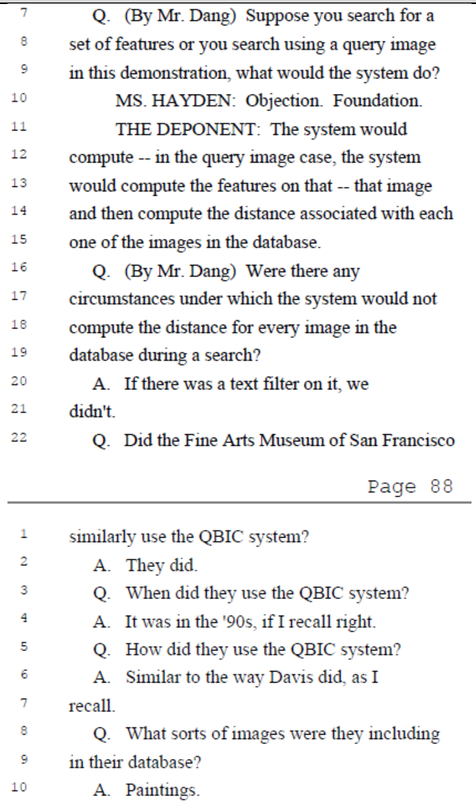 Q. (By Mr. Dang) Suppose you search for a set of features or you search using a query image in this demonstration, what would the system do? MS. HAYDEN: Objection. Foundation. THE DEPONENT: The system would compute -- in the query image case, the system would compute the features on that -- that image and then compute the distance associated with each one of the images in the database. Q. (By Mr. Dang) Were there any circumstances under which the system would not compute the distance for every image in the database during a search? A. If there was a text filter on it, we didn't. Q. Did the Fine Arts Museum of San Francisco Page 88 similarly use the QBIC system? A. They did. Q. When did they use the QBIC system? A. It was in the '90s, if I recall right. Q. How did they use the QBIC system? A. Similar to the way Davis did, as I recall. Q. What sorts of images were they including in their database? A. Paintings. Flickner Depo at 142:

Element	'036 Claim Recitation	Exemplary Citations in Reference
		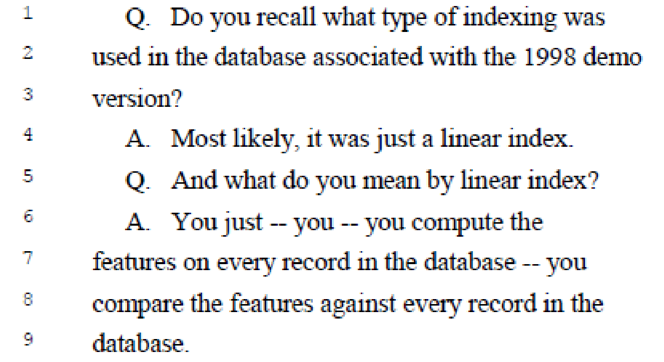 Q. Do you recall what type of indexing was used in the database associated with the 1998 demo version? A. Most likely, it was just a linear index. Q. And what do you mean by linear index? A. You just -- you -- you compute the features on every record in the database -- you compare the features against every record in the database. Flickner Depo at 168-169: 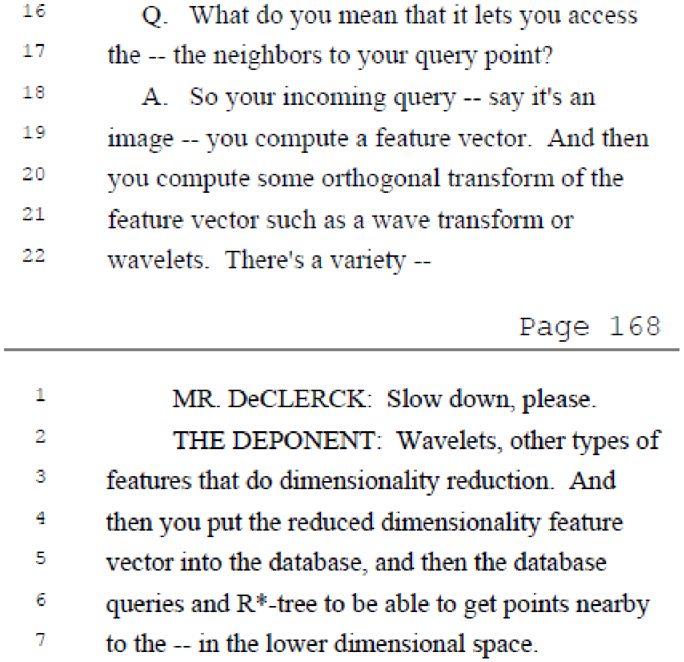 Q. What do you mean that it lets you access the -- the neighbors to your query point? A. So your incoming query -- say it's an image -- you compute a feature vector. And then you compute some orthogonal transform of the feature vector such as a wave transform or wavelets. There's a variety -- Page 168 MR. DeCLERCK: Slow down, please. THE DEPONENT: Wavelets, other types of features that do dimensionality reduction. And then you put the reduced dimensionality feature vector into the database, and then the database queries and R*-tree to be able to get points nearby to the -- in the lower dimensional space.

Element	'036 Claim Recitation	Exemplary Citations in Reference
		Flickner Depo at 181: 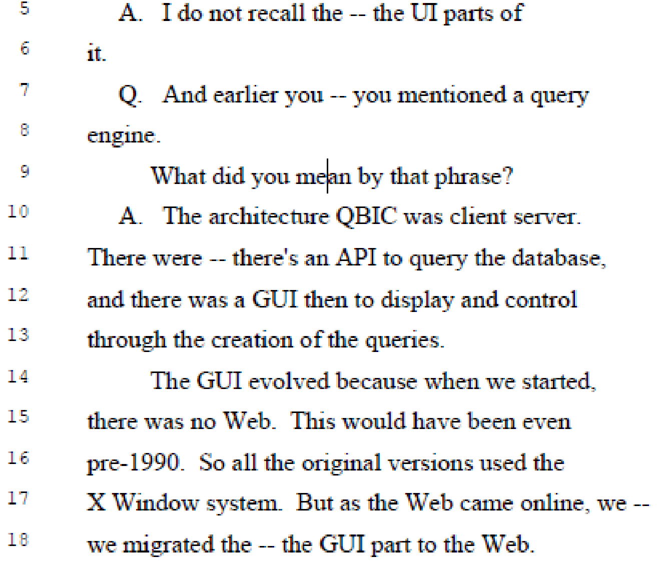 A. I do not recall the -- the UI parts of it. Q. And earlier you -- you mentioned a query engine. What did you mean by that phrase? A. The architecture QBIC was client server. There were -- there's an API to query the database, and there was a GUI then to display and control through the creation of the queries. The GUI evolved because when we started, there was no Web. This would have been even pre-1990. So all the original versions used the X Window system. But as the Web came online, we -- we migrated the -- the GUI part to the Web. Flickner Depo at 201:

Element	'036 Claim Recitation	Exemplary Citations in Reference
		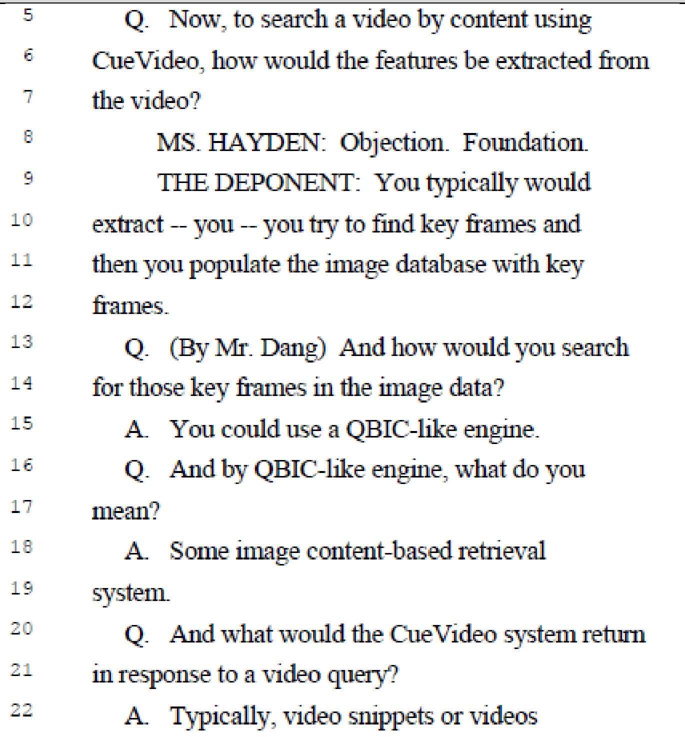 Q. Now, to search a video by content using CueVideo, how would the features be extracted from the video? MS. HAYDEN: Objection. Foundation. THE DEPONENT: You typically would extract -- you -- you try to find key frames and then you populate the image database with key frames. Q. (By Mr. Dang) And how would you search for those key frames in the image data? A. You could use a QBIC-like engine. Q. And by QBIC-like engine, what do you mean? A. Some image content-based retrieval system. Q. And what would the CueVideo system return in response to a video query? A. Typically, video snippets or videos Flickner Depo at 91-92:

Element	'036 Claim Recitation	Exemplary Citations in Reference
		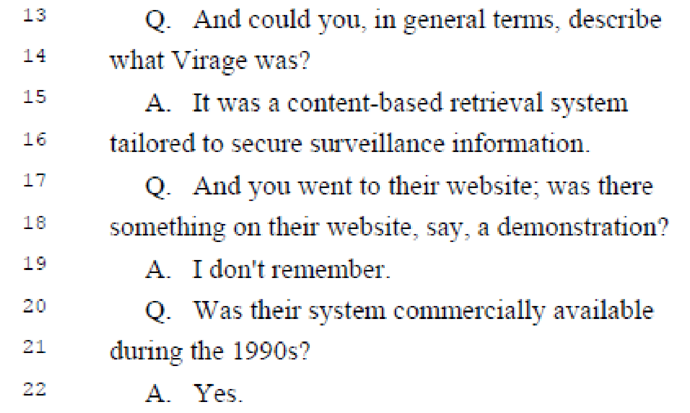 Q. And could you, in general terms, describe what Virage was? A. It was a content-based retrieval system tailored to secure surveillance information. Q. And you went to their website; was there something on their website, say, a demonstration? A. I don't remember. Q. Was their system commercially available during the 1990s? A. Yes. 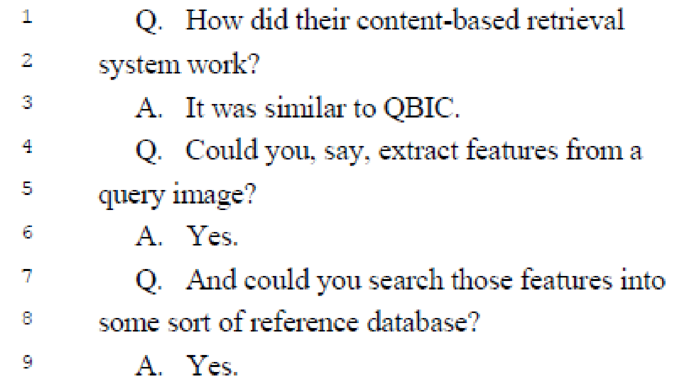 Q. How did their content-based retrieval system work? A. It was similar to QBIC. Q. Could you, say, extract features from a query image? A. Yes. Q. And could you search those features into some sort of reference database? A. Yes. Flickner Depo at 22-25:

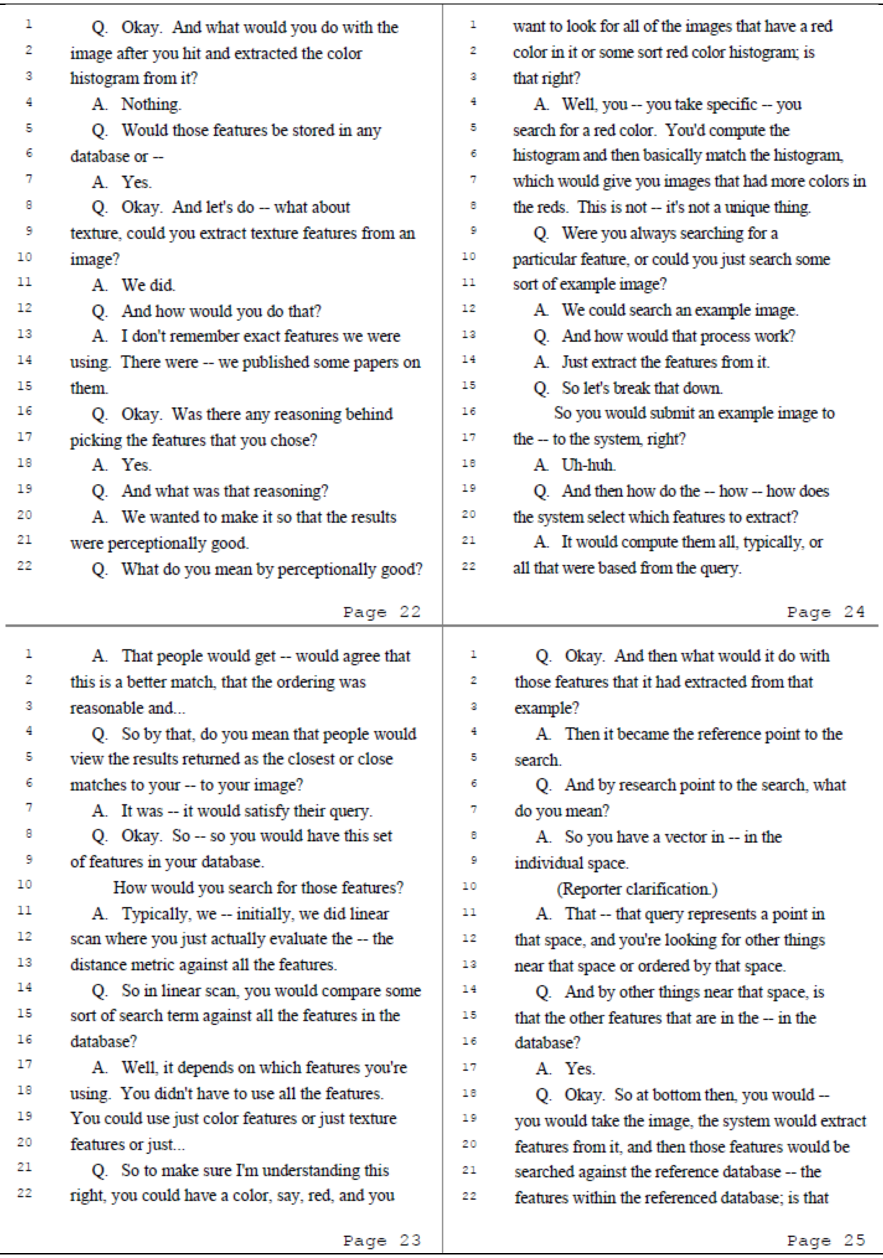

Q. Okay. And what would you do with the image after you hit and extracted the color histogram from it?

A. Nothing.

Q. Would those features be stored in any database or --

A. Yes.

Q. Okay. And let's do -- what about texture, could you extract texture features from an image?

A. We did.

Q. And how would you do that?

A. I don't remember exact features we were using. There were -- we published some papers on them.

Q. Okay. Was there any reasoning behind picking the features that you chose?

A. Yes.

Q. And what was that reasoning?

A. We wanted to make it so that the results were perceptionally good.

Q. What do you mean by perceptionally good?

Page 22

A. That people would get -- would agree that this is a better match, that the ordering was reasonable and...

Q. So by that, do you mean that people would view the results returned as the closest or close matches to your -- to your image?

A. It was -- it would satisfy their query.

Q. Okay. So -- so you would have this set of features in your database.

How would you search for those features?

A. Typically, we -- initially, we did linear scan where you just actually evaluate the -- the distance metric against all the features.

Q. So in linear scan, you would compare some sort of search term against all the features in the database?

A. Well, it depends on which features you're using. You didn't have to use all the features. You could use just color features or just texture features or just...

Q. So to make sure I'm understanding this right, you could have a color, say, red, and you

Page 23

want to look for all of the images that have a red color in it or some sort red color histogram; is that right?

A. Well, you -- you take specific -- you search for a red color. You'd compute the histogram and then basically match the histogram, which would give you images that had more colors in the reds. This is not -- it's not a unique thing.

Q. Were you always searching for a particular feature, or could you just search some sort of example image?

A. We could search an example image.

Q. And how would that process work?

A. Just extract the features from it.

Q. So let's break that down.

So you would submit an example image to the -- to the system, right?

A. Uh-huh.

Q. And then how do the -- how -- how does the system select which features to extract?

A. It would compute them all, typically, or all that were based from the query.

Page 24

Q. Okay. And then what would it do with those features that it had extracted from that example?

A. Then it became the reference point to the search.

Q. And by research point to the search, what do you mean?

A. So you have a vector in -- in the individual space.

(Reporter clarification.)

A. That -- that query represents a point in that space, and you're looking for other things near that space or ordered by that space.

Q. And by other things near that space, is that the other features that are in the -- in the database?

A. Yes.

Q. Okay. So at bottom then, you would -- you would take the image, the system would extract features from it, and then those features would be searched against the reference database -- the features within the referenced database; is that

Page 25

Element	'036 Claim Recitation	Exemplary Citations in Reference
		Flickner Depo at 48-49: 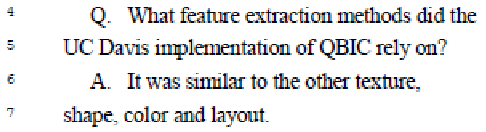 Q. What feature extraction methods did the UC Davis implementation of QBIC rely on? A. It was similar to the other texture, shape, color and layout. Flickner Depo at 159: 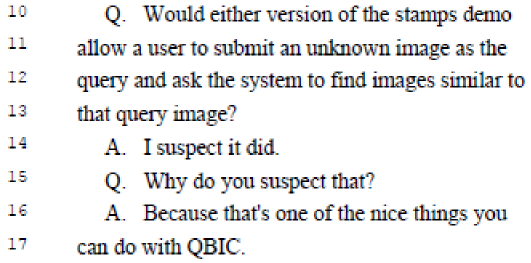 Q. Would either version of the stamps demo allow a user to submit an unknown image as the query and ask the system to find images similar to that query image? A. I suspect it did. Q. Why do you suspect that? A. Because that's one of the nice things you can do with QBIC. Flickner Depo at 209-210: 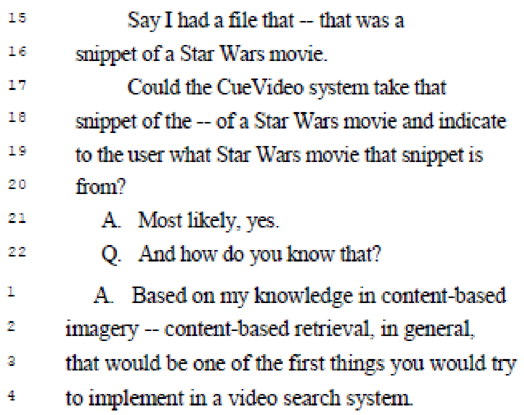 Q. Okay. Let me ask -- let me ask that question using -- using an example. Say I had a file that -- that was a snippet of a Star Wars movie. Could the CueVideo system take that snippet of the -- of a Star Wars movie and indicate to the user what Star Wars movie that snippet is from? A. Most likely, yes. Q. And how do you know that? A. Based on my knowledge in content-based imagery -- content-based retrieval, in general, that would be one of the first things you would try to implement in a video search system.

Element	'036 Claim Recitation	Exemplary Citations in Reference
		See also Flickner Depo at 18, 54-55, 76. Flickner 8/29/23 Depo at 33–37: 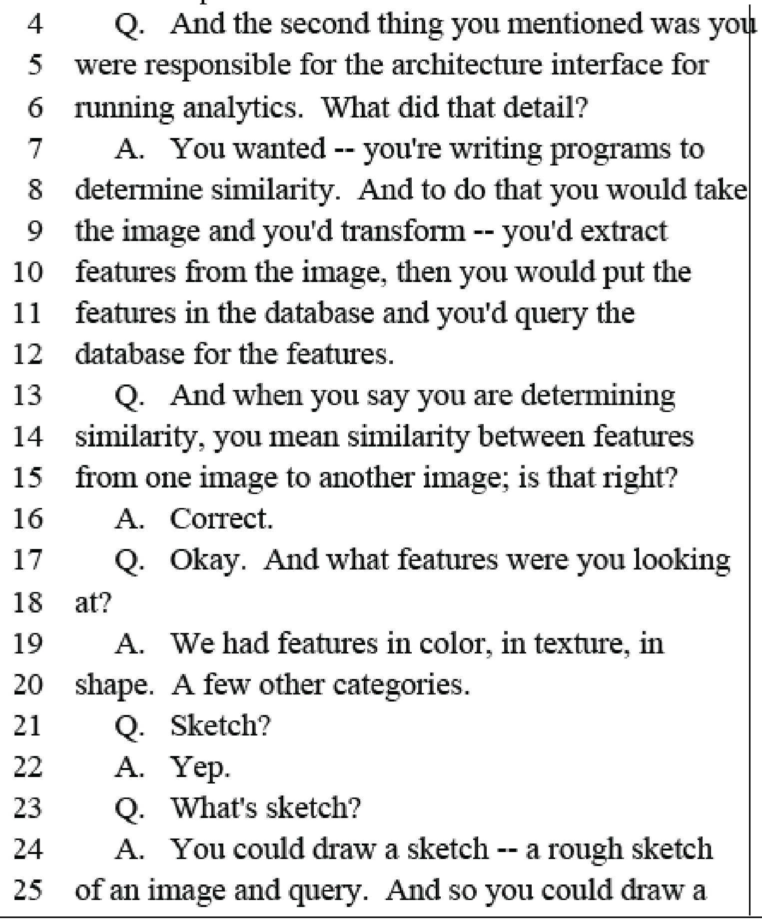 Q. And the second thing you mentioned was you were responsible for the architecture interface for running analytics. What did that detail? A. You wanted -- you're writing programs to determine similarity. And to do that you would take the image and you'd transform -- you'd extract features from the image, then you would put the features in the database and you'd query the database for the features. Q. And when you say you are determining similarity, you mean similarity between features from one image to another image; is that right? A. Correct. Q. Okay. And what features were you looking at? A. We had features in color, in texture, in shape. A few other categories. Q. Sketch? A. Yep. Q. What's sketch? A. You could draw a sketch -- a rough sketch of an image and query. And so you could draw a

Element	'036 Claim Recitation	Exemplary Citations in Reference
		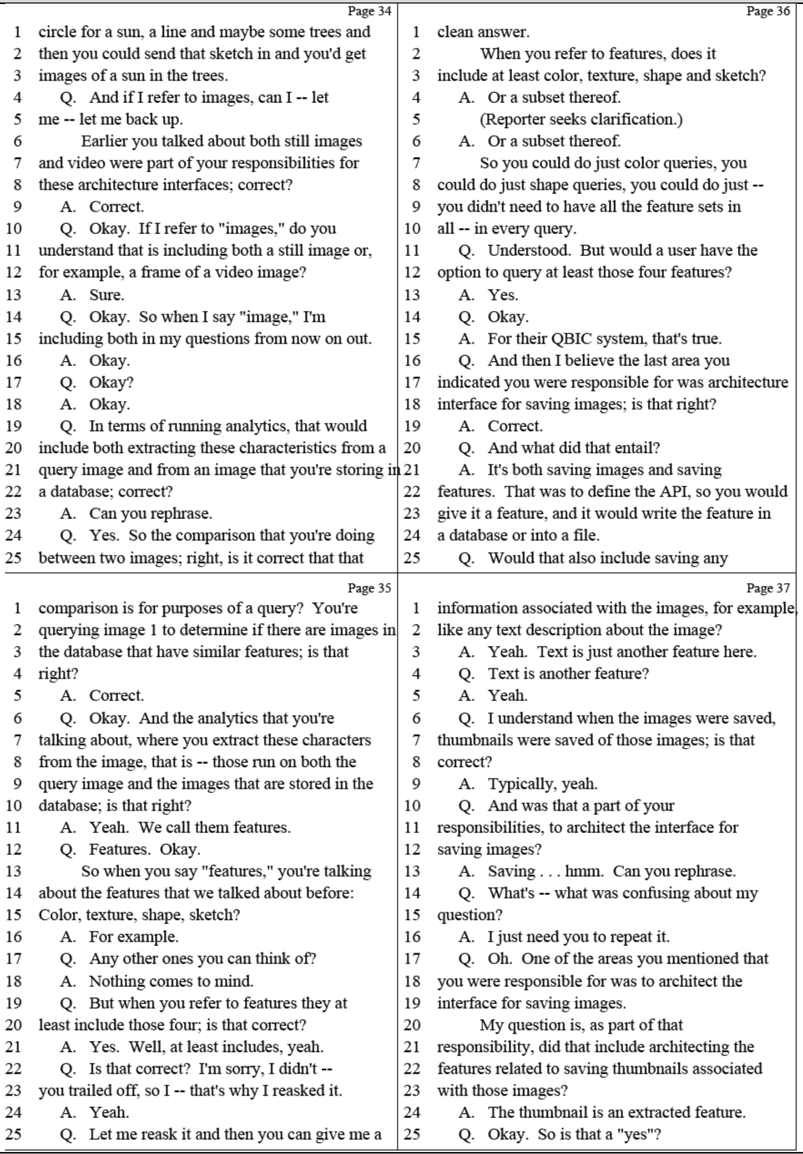 Page 34 circle for a sun, a line and maybe some trees and then you could send that sketch in and you'd get images of a sun in the trees. Q. And if I refer to images, can I -- let me -- let me back up. Earlier you talked about both still images and video were part of your responsibilities for these architecture interfaces; correct? A. Correct. Q. Okay. If I refer to "images," do you understand that is including both a still image or, for example, a frame of a video image? A. Sure. Q. Okay. So when I say "image," I'm including both in my questions from now on out. A. Okay. Q. Okay? A. Okay. Q. In terms of running analytics, that would include both extracting these characteristics from a query image and from an image that you're storing in a database; correct? A. Can you rephrase. Q. Yes. So the comparison that you're doing between two images; right, is it correct that that Page 35 comparison is for purposes of a query? You're querying image 1 to determine if there are images in the database that have similar features; is that right? A. Correct. Q. Okay. And the analytics that you're talking about, where you extract these characters from the image, that is -- those run on both the query image and the images that are stored in the database; is that right? A. Yeah. We call them features. Q. Features. Okay. So when you say "features," you're talking about the features that we talked about before: Color, texture, shape, sketch? A. For example. Q. Any other ones you can think of? A. Nothing comes to mind. Q. But when you refer to features they at least include those four; is that correct? A. Yes. Well, at least includes, yeah. Q. Is that correct? I'm sorry, I didn't -- you trailed off, so I -- that's why I reasked it. A. Yeah. Q. Let me reask it and then you can give me a Page 36 clean answer. When you refer to features, does it include at least color, texture, shape and sketch? A. Or a subset thereof. (Reporter seeks clarification.) A. Or a subset thereof. So you could do just color queries, you could do just shape queries, you could do just -- you didn't need to have all the feature sets in all -- in every query. Q. Understood. But would a user have the option to query at least those four features? A. Yes. Q. Okay. A. For their QBIC system, that's true. Q. And then I believe the last area you indicated you were responsible for was architecture interface for saving images; is that right? A. Correct. Q. And what did that entail? A. It's both saving images and saving features. That was to define the API, so you would give it a feature, and it would write the feature in a database or into a file. Q. Would that also include saving any Page 37 information associated with the images, for example, like any text description about the image? A. Yeah. Text is just another feature here. Q. Text is another feature? A. Yeah. Q. I understand when the images were saved, thumbnails were saved of those images; is that correct? A. Typically, yeah. Q. And was that a part of your responsibilities, to architect the interface for saving images? A. Saving . . . hmm. Can you rephrase. Q. What's -- what was confusing about my question? A. I just need you to repeat it. Q. Oh. One of the areas you mentioned that you were responsible for was to architect the interface for saving images. My question is, as part of that responsibility, did that include architecting the features related to saving thumbnails associated with those images? A. The thumbnail is an extracted feature. Q. Okay. So is that a "yes"?

Element	'036 Claim Recitation	Exemplary Citations in Reference
		Id. at 52–56: 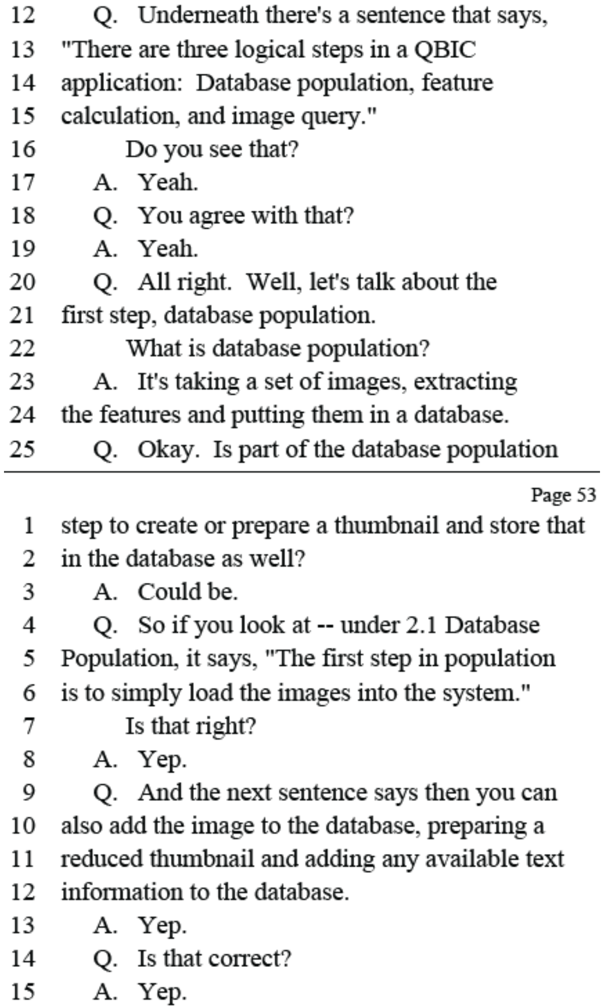 Q. Underneath there's a sentence that says, "There are three logical steps in a QBIC application: Database population, feature calculation, and image query." Do you see that? A. Yeah. Q. You agree with that? A. Yeah. Q. All right. Well, let's talk about the first step, database population. What is database population? A. It's taking a set of images, extracting the features and putting them in a database. Q. Okay. Is part of the database population Page 53 step to create or prepare a thumbnail and store that in the database as well? A. Could be. Q. So if you look at -- under 2.1 Database Population, it says, "The first step in population is to simply load the images into the system." Is that right? A. Yep. Q. And the next sentence says then you can also add the image to the database, preparing a reduced thumbnail and adding any available text information to the database. A. Yep. Q. Is that correct? A. Yep.

Element	'036 Claim Recitation	Exemplary Citations in Reference
		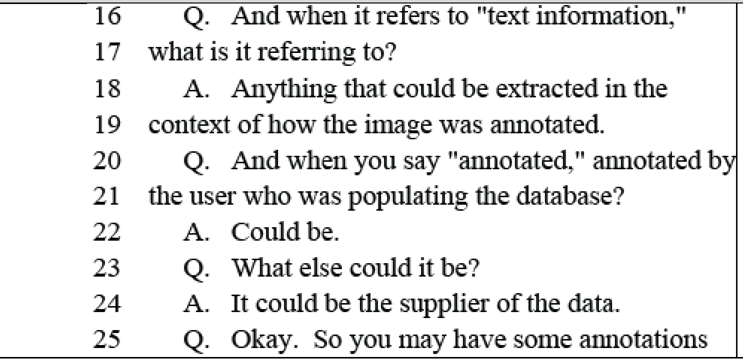 Q. And when it refers to "text information," what is it referring to? A. Anything that could be extracted in the context of how the image was annotated. Q. And when you say "annotated," annotated by the user who was populating the database? A. Could be. Q. What else could it be? A. It could be the supplier of the data. Q. Okay. So you may have some annotations

Element	'036 Claim Recitation	Exemplary Citations in Reference
		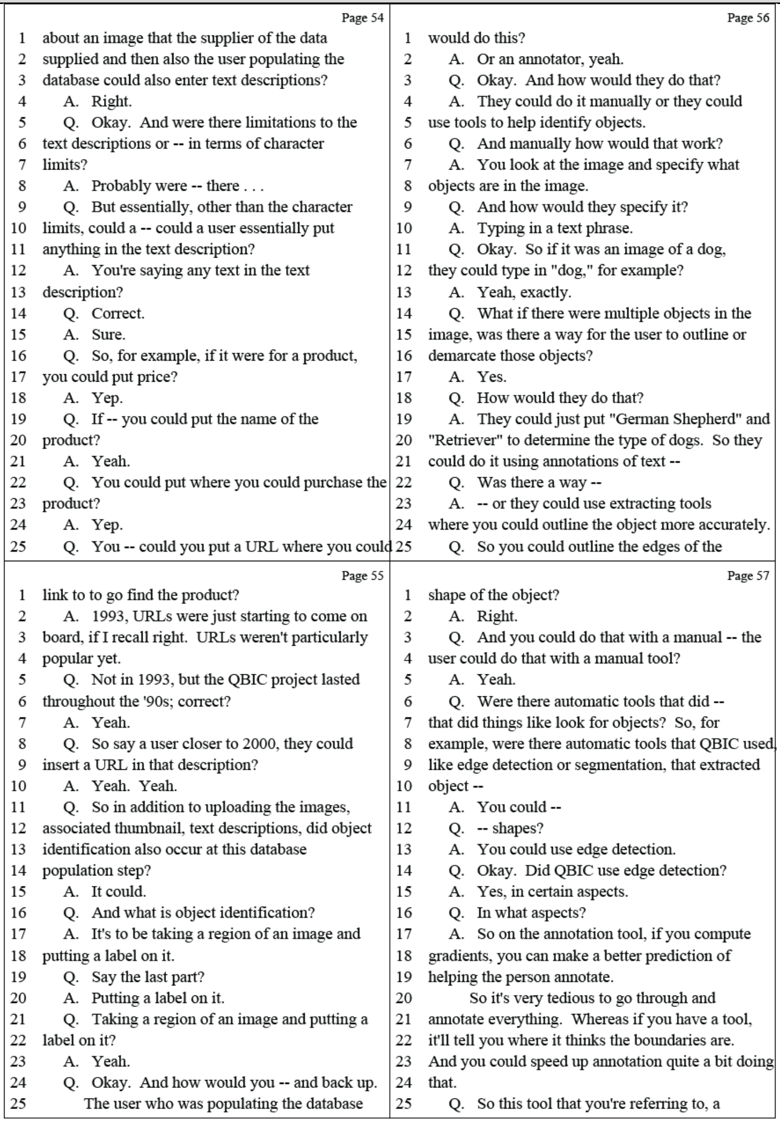

Page 54

about an image that the supplier of the data supplied and then also the user populating the database could also enter text descriptions?
A. Right.
Q. Okay. And were there limitations to the text descriptions or -- in terms of character limits?
A. Probably were -- there . . .
Q. But essentially, other than the character limits, could a -- could a user essentially put anything in the text description?
A. You're saying any text in the text description?
Q. Correct.
A. Sure.
Q. So, for example, if it were for a product, you could put price?
A. Yep.
Q. If -- you could put the name of the product?
A. Yeah.
Q. You could put where you could purchase the product?
A. Yep.
Q. You -- could you put a URL where you could

Page 55

link to to go find the product?
A. 1993, URLs were just starting to come on board, if I recall right. URLs weren't particularly popular yet.
Q. Not in 1993, but the QBIC project lasted throughout the '90s; correct?
A. Yeah.
Q. So say a user closer to 2000, they could insert a URL in that description?
A. Yeah. Yeah.
Q. So in addition to uploading the images, associated thumbnail, text descriptions, did object identification also occur at this database population step?
A. It could.
Q. And what is object identification?
A. It's to be taking a region of an image and putting a label on it.
Q. Say the last part?
A. Putting a label on it.
Q. Taking a region of an image and putting a label on it?
A. Yeah.
Q. Okay. And how would you -- and back up.
The user who was populating the database

Page 56

would do this?
A. Or an annotator, yeah.
Q. Okay. And how would they do that?
A. They could do it manually or they could use tools to help identify objects.
Q. And manually how would that work?
A. You look at the image and specify what objects are in the image.
Q. And how would they specify it?
A. Typing in a text phrase.
Q. Okay. So if it was an image of a dog, they could type in "dog," for example?
A. Yeah, exactly.
Q. What if there were multiple objects in the image, was there a way for the user to outline or demarcate those objects?
A. Yes.
Q. How would they do that?
A. They could just put "German Shepherd" and "Retriever" to determine the type of dogs. So they could do it using annotations of text --
Q. Was there a way --
A. -- or they could use extracting tools where you could outline the object more accurately.
Q. So you could outline the edges of the

Page 57

shape of the object?
A. Right.
Q. And you could do that with a manual -- the user could do that with a manual tool?
A. Yeah.
Q. Were there automatic tools that did -- that did things like look for objects? So, for example, were there automatic tools that QBIC used, like edge detection or segmentation, that extracted object --
A. You could --
Q. -- shapes?
A. You could use edge detection.
Q. Okay. Did QBIC use edge detection?
A. Yes, in certain aspects.
Q. In what aspects?
A. So on the annotation tool, if you compute gradients, you can make a better prediction of helping the person annotate.
So it's very tedious to go through and annotate everything. Whereas if you have a tool, it'll tell you where it thinks the boundaries are. And you could speed up annotation quite a bit doing that.
Q. So this tool that you're referring to, a

Element	'036 Claim Recitation	Exemplary Citations in Reference
		Id. at 60–61: 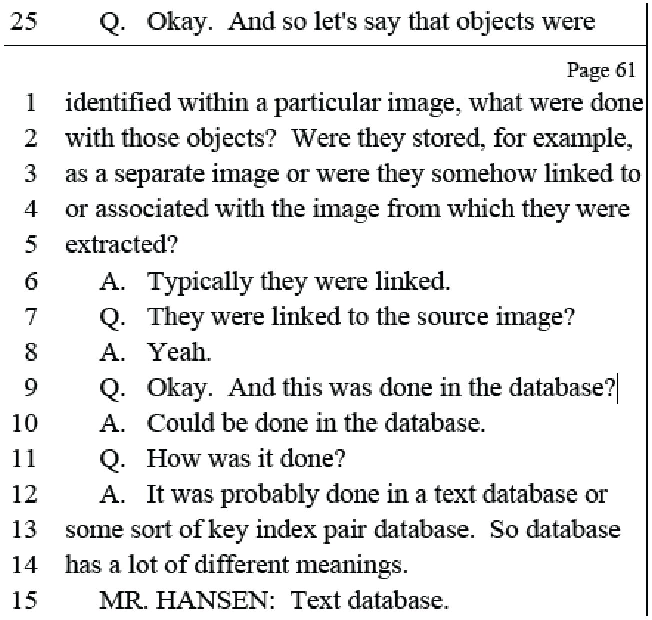Q. Okay. And so let's say that objects were Page 61 identified within a particular image, what were done with those objects? Were they stored, for example, as a separate image or were they somehow linked to or associated with the image from which they were extracted? A. Typically they were linked. Q. They were linked to the source image? A. Yeah. Q. Okay. And this was done in the database? A. Could be done in the database. Q. How was it done? A. It was probably done in a text database or some sort of key index pair database. So database has a lot of different meanings. MR. HANSEN: Text database. *Id.* at 62–63:

Element	'036 Claim Recitation	Exemplary Citations in Reference
		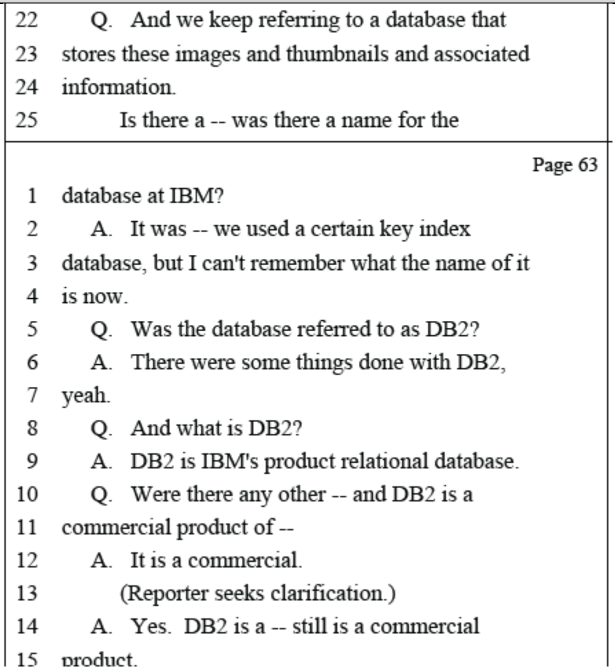 Q. And we keep referring to a database that stores these images and thumbnails and associated information. Is there a -- was there a name for the Page 63 database at IBM? A. It was -- we used a certain key index database, but I can't remember what the name of it is now. Q. Was the database referred to as DB2? A. There were some things done with DB2, yeah. Q. And what is DB2? A. DB2 is IBM's product relational database. Q. Were there any other -- and DB2 is a commercial product of -- A. It is a commercial. (Reporter seeks clarification.) A. Yes. DB2 is a -- still is a commercial product. *Id.* at 77: 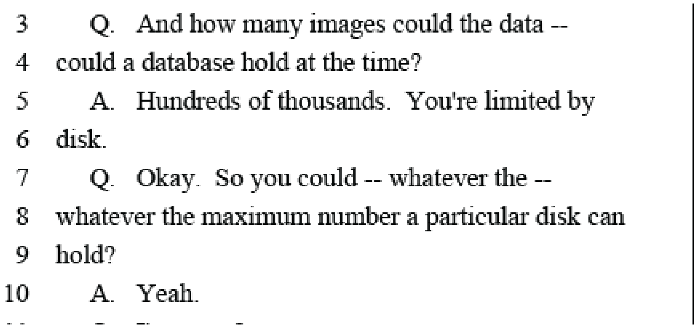 Q. And how many images could the data -- could a database hold at the time? A. Hundreds of thousands. You're limited by disk. Q. Okay. So you could -- whatever the -- whatever the maximum number a particular disk can hold? A. Yeah. *Id.* at 77–78:

Element	'036 Claim Recitation	Exemplary Citations in Reference
		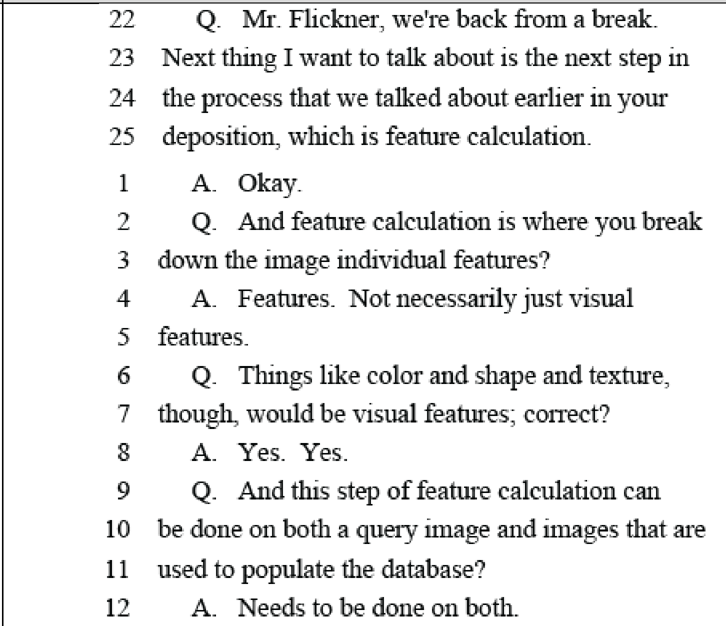 Q. Mr. Flickner, we're back from a break. Next thing I want to talk about is the next step in the process that we talked about earlier in your deposition, which is feature calculation. A. Okay. Q. And feature calculation is where you break down the image individual features? A. Features. Not necessarily just visual features. Q. Things like color and shape and texture, though, would be visual features; correct? A. Yes. Yes. Q. And this step of feature calculation can be done on both a query image and images that are used to populate the database? A. Needs to be done on both. *Id.* at 99: 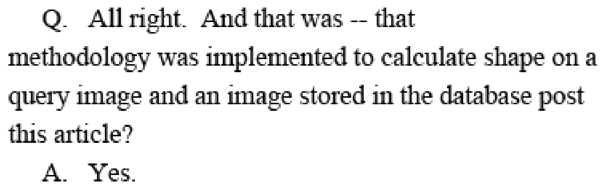 Q. All right. And that was -- that methodology was implemented to calculate shape on a query image and an image stored in the database post this article? A. Yes. *Id.* at 103:

Element	'036 Claim Recitation	Exemplary Citations in Reference
		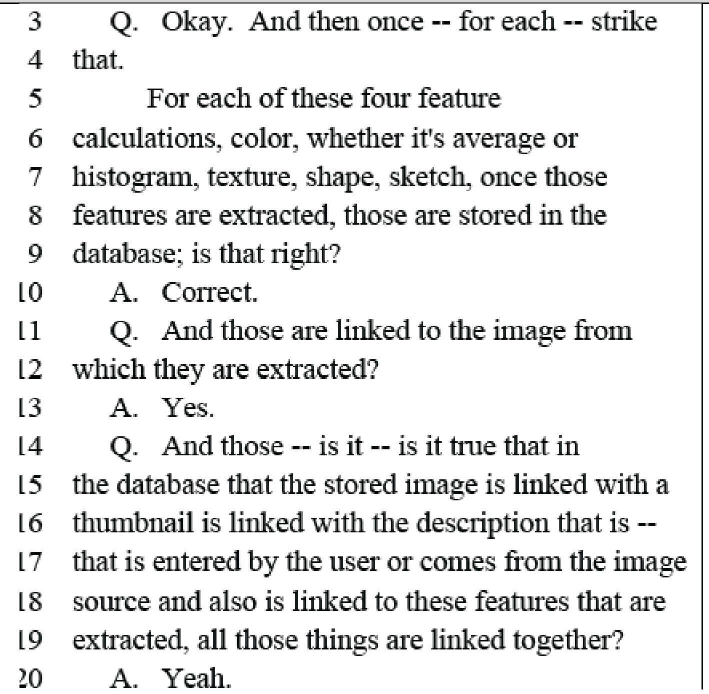 Q. Okay. And then once -- for each -- strike that. For each of these four feature calculations, color, whether it's average or histogram, texture, shape, sketch, once those features are extracted, those are stored in the database; is that right? A. Correct. Q. And those are linked to the image from which they are extracted? A. Yes. Q. And those -- is it -- is it true that in the database that the stored image is linked with a thumbnail is linked with the description that is -- that is entered by the user or comes from the image source and also is linked to these features that are extracted, all those things are linked together? A. Yeah. *Id.* at 161:

Element	'036 Claim Recitation	Exemplary Citations in Reference
		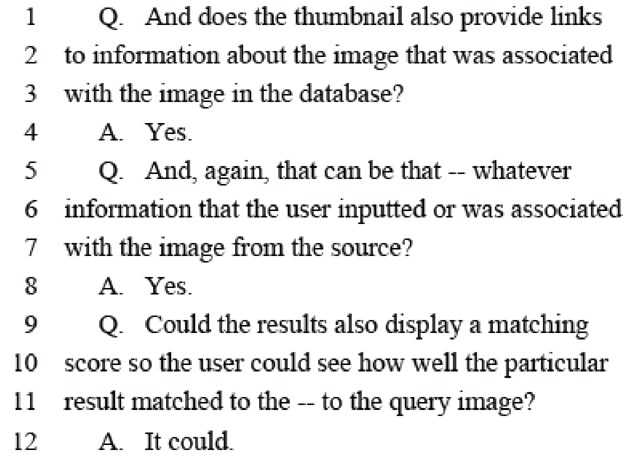 Q. And does the thumbnail also provide links to information about the image that was associated with the image in the database? A. Yes. Q. And, again, that can be that -- whatever information that the user inputted or was associated with the image from the source? A. Yes. Q. Could the results also display a matching score so the user could see how well the particular result matched to the -- to the query image? A. It could. *Id.* at 174–182:

Element	'036 Claim Recitation	Exemplary Citations in Reference
		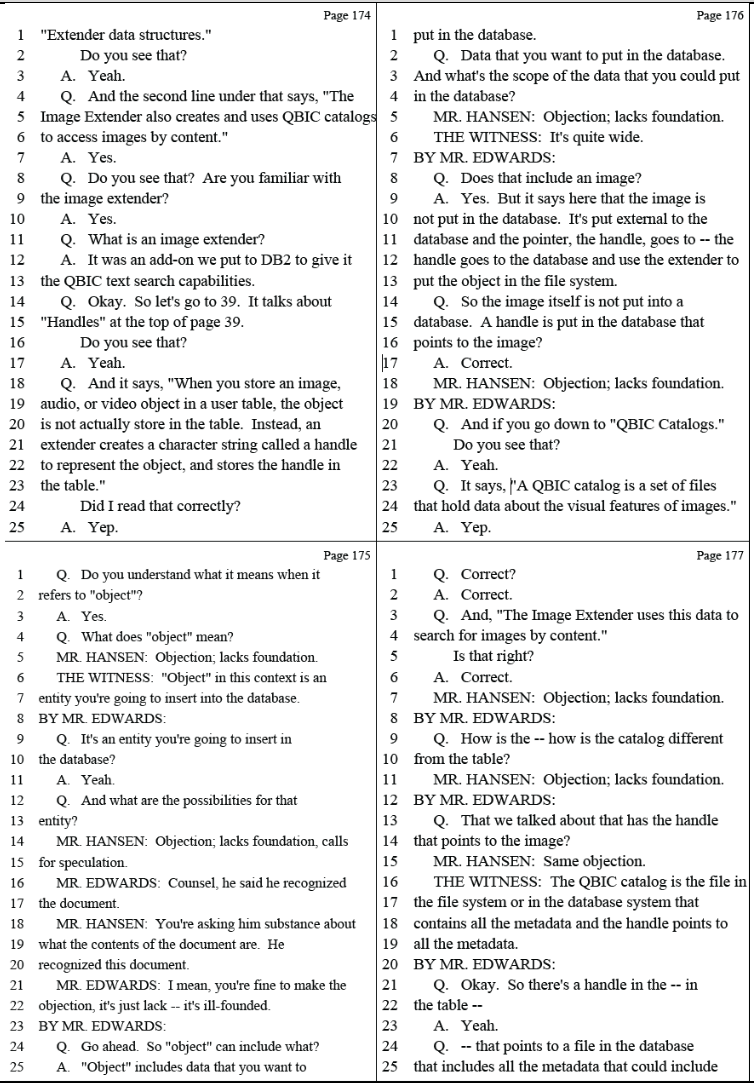

Page 174

"Extender data structures."
Do you see that?
A. Yeah.
Q. And the second line under that says, "The Image Extender also creates and uses QBIC catalogs to access images by content."
A. Yes.
Q. Do you see that? Are you familiar with the image extender?
A. Yes.
Q. What is an image extender?
A. It was an add-on we put to DB2 to give it the QBIC text search capabilities.
Q. Okay. So let's go to 39. It talks about "Handles" at the top of page 39.
Do you see that?
A. Yeah.
Q. And it says, "When you store an image, audio, or video object in a user table, the object is not actually store in the table. Instead, an extender creates a character string called a handle to represent the object, and stores the handle in the table."
Did I read that correctly?
A. Yep.

Page 175

Q. Do you understand what it means when it refers to "object"?
A. Yes.
Q. What does "object" mean?
MR. HANSEN: Objection; lacks foundation.
THE WITNESS: "Object" in this context is an entity you're going to insert into the database.
BY MR. EDWARDS:
Q. It's an entity you're going to insert in the database?
A. Yeah.
Q. And what are the possibilities for that entity?
MR. HANSEN: Objection; lacks foundation, calls for speculation.
MR. EDWARDS: Counsel, he said he recognized the document.
MR. HANSEN: You're asking him substance about what the contents of the document are. He recognized this document.
MR. EDWARDS: I mean, you're fine to make the objection, it's just lack -- it's ill-founded.
BY MR. EDWARDS:
Q. Go ahead. So "object" can include what?
A. "Object" includes data that you want to

Page 176

put in the database.
Q. Data that you want to put in the database. And what's the scope of the data that you could put in the database?
MR. HANSEN: Objection; lacks foundation.
THE WITNESS: It's quite wide.
BY MR. EDWARDS:
Q. Does that include an image?
A. Yes. But it says here that the image is not put in the database. It's put external to the database and the pointer, the handle, goes to -- the handle goes to the database and use the extender to put the object in the file system.
Q. So the image itself is not put into a database. A handle is put in the database that points to the image?
A. Correct.
MR. HANSEN: Objection; lacks foundation.
BY MR. EDWARDS:
Q. And if you go down to "QBIC Catalogs."
Do you see that?
A. Yeah.
Q. It says, "A QBIC catalog is a set of files that hold data about the visual features of images."
A. Yep.

Page 177

Q. Correct?
A. Correct.
Q. And, "The Image Extender uses this data to search for images by content."
Is that right?
A. Correct.
MR. HANSEN: Objection; lacks foundation.
BY MR. EDWARDS:
Q. How is the -- how is the catalog different from the table?
MR. HANSEN: Objection; lacks foundation.
BY MR. EDWARDS:
Q. That we talked about that has the handle that points to the image?
MR. HANSEN: Same objection.
THE WITNESS: The QBIC catalog is the file in the file system or in the database system that contains all the metadata and the handle points to all the metadata.
BY MR. EDWARDS:
Q. Okay. So there's a handle in the -- in the table --
A. Yeah.
Q. -- that points to a file in the database that includes all the metadata that could include

Element	'036 Claim Recitation	Exemplary Citations in Reference
		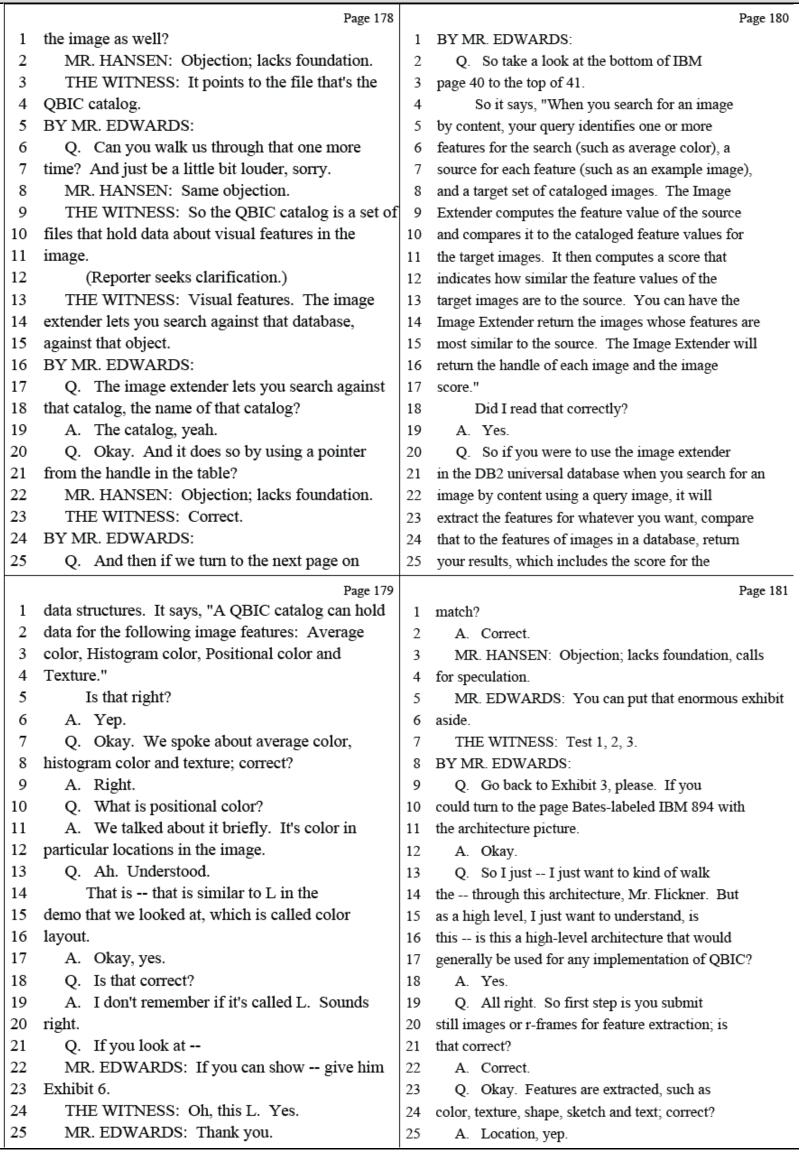 Page 178 the image as well? MR. HANSEN: Objection; lacks foundation. THE WITNESS: It points to the file that's the QBIC catalog. BY MR. EDWARDS: Q. Can you walk us through that one more time? And just be a little bit louder, sorry. MR. HANSEN: Same objection. THE WITNESS: So the QBIC catalog is a set of files that hold data about visual features in the image. (Reporter seeks clarification.) THE WITNESS: Visual features. The image extender lets you search against that database, against that object. BY MR. EDWARDS: Q. The image extender lets you search against that catalog, the name of that catalog? A. The catalog, yeah. Q. Okay. And it does so by using a pointer from the handle in the table? MR. HANSEN: Objection; lacks foundation. THE WITNESS: Correct. BY MR. EDWARDS: Q. And then if we turn to the next page on Page 179 data structures. It says, "A QBIC catalog can hold data for the following image features: Average color, Histogram color, Positional color and Texture." Is that right? A. Yep. Q. Okay. We spoke about average color, histogram color and texture; correct? A. Right. Q. What is positional color? A. We talked about it briefly. It's color in particular locations in the image. Q. Ah. Understood. That is -- that is similar to L in the demo that we looked at, which is called color layout. A. Okay, yes. Q. Is that correct? A. I don't remember if it's called L. Sounds right. Q. If you look at -- MR. EDWARDS: If you can show -- give him Exhibit 6. THE WITNESS: Oh, this L. Yes. MR. EDWARDS: Thank you. Page 180 BY MR. EDWARDS: Q. So take a look at the bottom of IBM page 40 to the top of 41. So it says, "When you search for an image by content, your query identifies one or more features for the search (such as average color), a source for each feature (such as an example image), and a target set of cataloged images. The Image Extender computes the feature value of the source and compares it to the cataloged feature values for the target images. It then computes a score that indicates how similar the feature values of the target images are to the source. You can have the Image Extender return the images whose features are most similar to the source. The Image Extender will return the handle of each image and the image score." Did I read that correctly? A. Yes. Q. So if you were to use the image extender in the DB2 universal database when you search for an image by content using a query image, it will extract the features for whatever you want, compare that to the features of images in a database, return your results, which includes the score for the Page 181 match? A. Correct. MR. HANSEN: Objection; lacks foundation, calls for speculation. MR. EDWARDS: You can put that enormous exhibit aside. THE WITNESS: Test 1, 2, 3. BY MR. EDWARDS: Q. Go back to Exhibit 3, please. If you could turn to the page Bates-labeled IBM 894 with the architecture picture. A. Okay. Q. So I just -- I just want to kind of walk the -- through this architecture, Mr. Flickner. But as a high level, I just want to understand, is this -- is this a high-level architecture that would generally be used for any implementation of QBIC? A. Yes. Q. All right. So first step is you submit still images or r-frames for feature extraction; is that correct? A. Correct. Q. Okay. Features are extracted, such as color, texture, shape, sketch and text; correct? A. Location, yep.

Element	'036 Claim Recitation	Exemplary Citations in Reference
		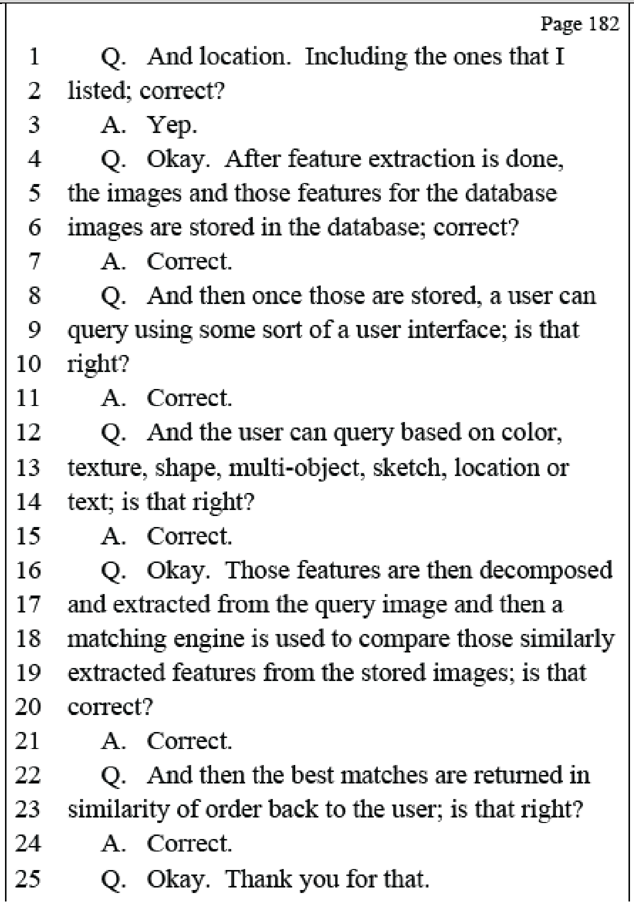 Page 182 Q. And location. Including the ones that I listed; correct? A. Yep. Q. Okay. After feature extraction is done, the images and those features for the database images are stored in the database; correct? A. Correct. Q. And then once those are stored, a user can query using some sort of a user interface; is that right? A. Correct. Q. And the user can query based on color, texture, shape, multi-object, sketch, location or text; is that right? A. Correct. Q. Okay. Those features are then decomposed and extracted from the query image and then a matching engine is used to compare those similarly extracted features from the stored images; is that correct? A. Correct. Q. And then the best matches are returned in similarity of order back to the user; is that right? A. Correct. Q. Okay. Thank you for that. *Id.* at 188–189:

Element	'036 Claim Recitation	Exemplary Citations in Reference
		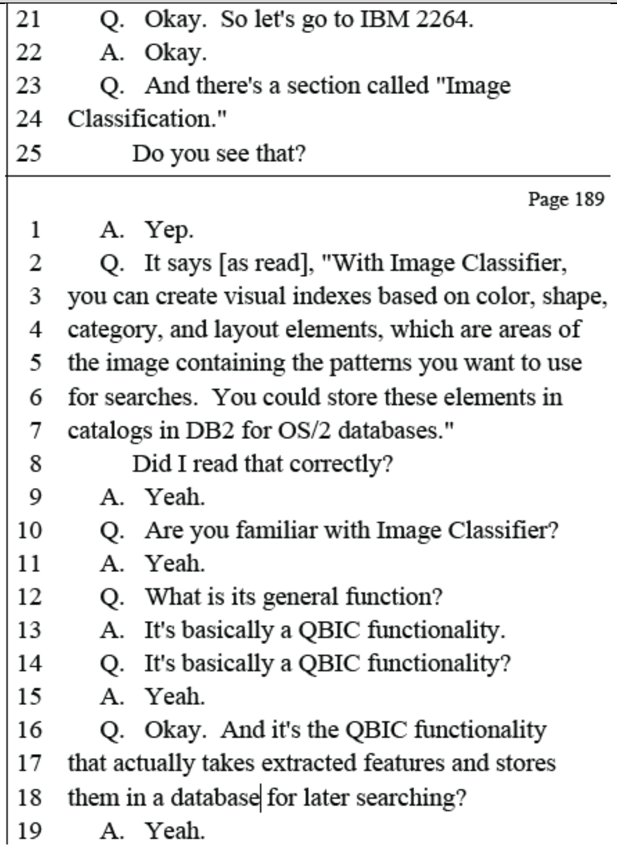 Q. Okay. So let's go to IBM 2264. A. Okay. Q. And there's a section called "Image Classification." Do you see that? Page 189 A. Yep. Q. It says [as read], "With Image Classifier, you can create visual indexes based on color, shape, category, and layout elements, which are areas of the image containing the patterns you want to use for searches. You could store these elements in catalogs in DB2 for OS/2 databases." Did I read that correctly? A. Yeah. Q. Are you familiar with Image Classifier? A. Yeah. Q. What is its general function? A. It's basically a QBIC functionality. Q. It's basically a QBIC functionality? A. Yeah. Q. Okay. And it's the QBIC functionality that actually takes extracted features and stores them in a database for later searching? A. Yeah. *Id.* at 205–206:

Element	'036 Claim Recitation	Exemplary Citations in Reference
		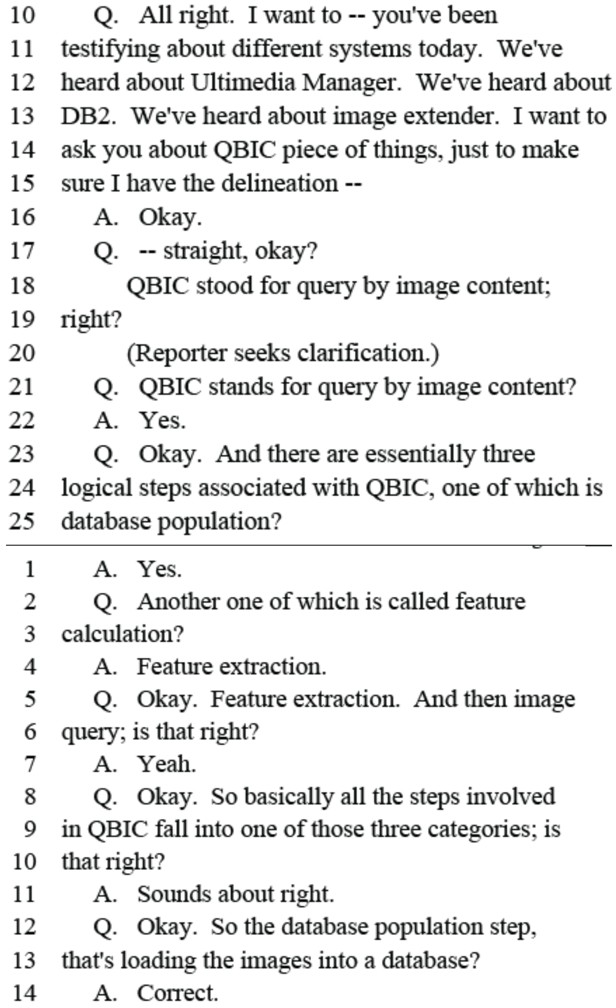 Q. All right. I want to -- you've been testifying about different systems today. We've heard about Ultimedia Manager. We've heard about DB2. We've heard about image extender. I want to ask you about QBIC piece of things, just to make sure I have the delineation -- A. Okay. Q. -- straight, okay? QBIC stood for query by image content; right? (Reporter seeks clarification.) Q. QBIC stands for query by image content? A. Yes. Q. Okay. And there are essentially three logical steps associated with QBIC, one of which is database population? A. Yes. Q. Another one of which is called feature calculation? A. Feature extraction. Q. Okay. Feature extraction. And then image query; is that right? A. Yeah. Q. Okay. So basically all the steps involved in QBIC fall into one of those three categories; is that right? A. Sounds about right. Q. Okay. So the database population step, that's loading the images into a database? A. Correct. *Id*. at 305:

Element	'036 Claim Recitation	Exemplary Citations in Reference
		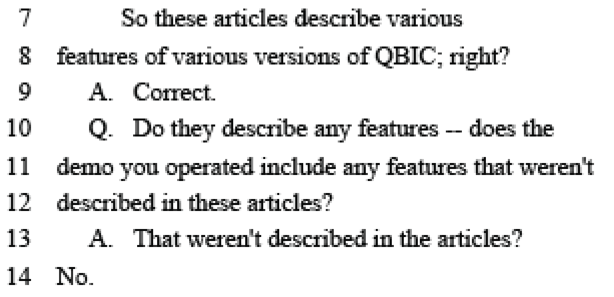 So these articles describe various features of various versions of QBIC; right? A. Correct. Q. Do they describe any features -- does the demo you operated include any features that weren't described in these articles? A. That weren't described in the articles? No. *Id.* at 35: 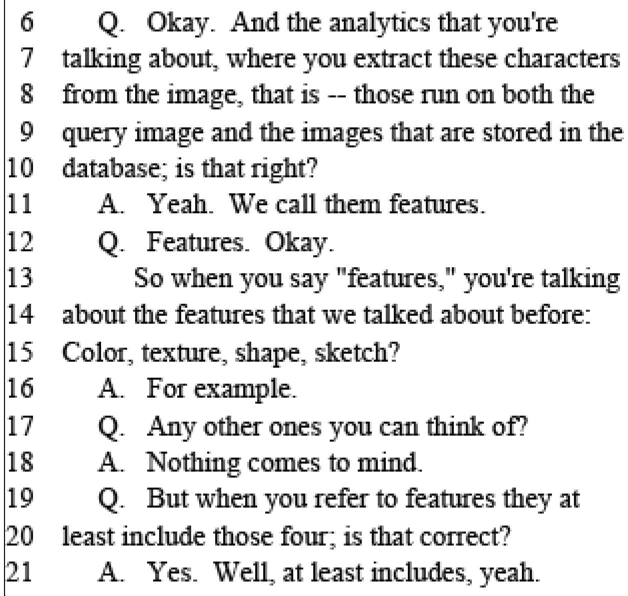 Q. Okay. And the analytics that you're talking about, where you extract these characters from the image, that is -- those run on both the query image and the images that are stored in the database; is that right? A. Yeah. We call them features. Q. Features. Okay. So when you say "features," you're talking about the features that we talked about before: Color, texture, shape, sketch? A. For example. Q. Any other ones you can think of? A. Nothing comes to mind. Q. But when you refer to features they at least include those four; is that correct? A. Yes. Well, at least includes, yeah. *Id.* at 109-110:

Element	'036 Claim Recitation	Exemplary Citations in Reference
		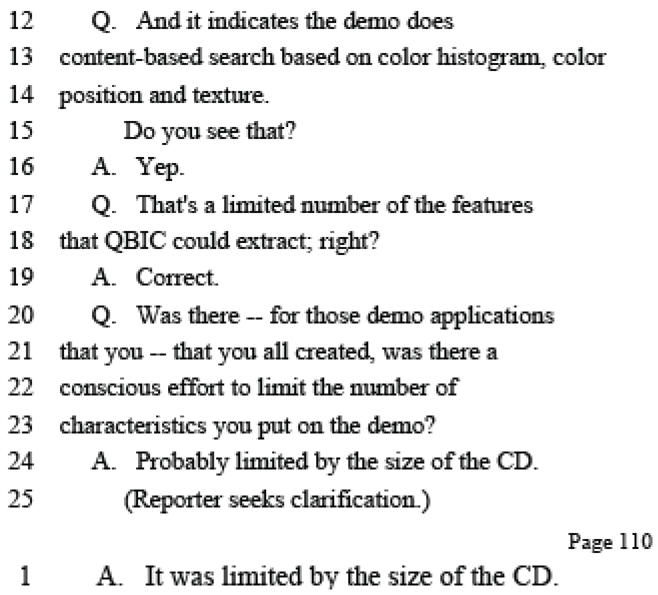 Q. And it indicates the demo does content-based search based on color histogram, color position and texture. Do you see that? A. Yep. Q. That's a limited number of the features that QBIC could extract; right? A. Correct. Q. Was there -- for those demo applications that you -- that you all created, was there a conscious effort to limit the number of characteristics you put on the demo? A. Probably limited by the size of the CD. (Reporter seeks clarification.) Page 110 A. It was limited by the size of the CD. *Id.* at 311: 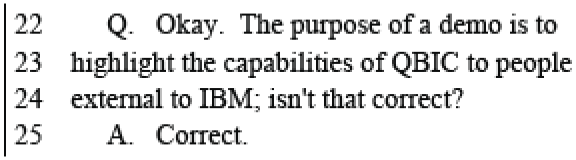 Q. Okay. The purpose of a demo is to highlight the capabilities of QBIC to people external to IBM; isn't that correct? A. Correct. *Id.* at 313-317:

Element	'036 Claim Recitation	Exemplary Citations in Reference
		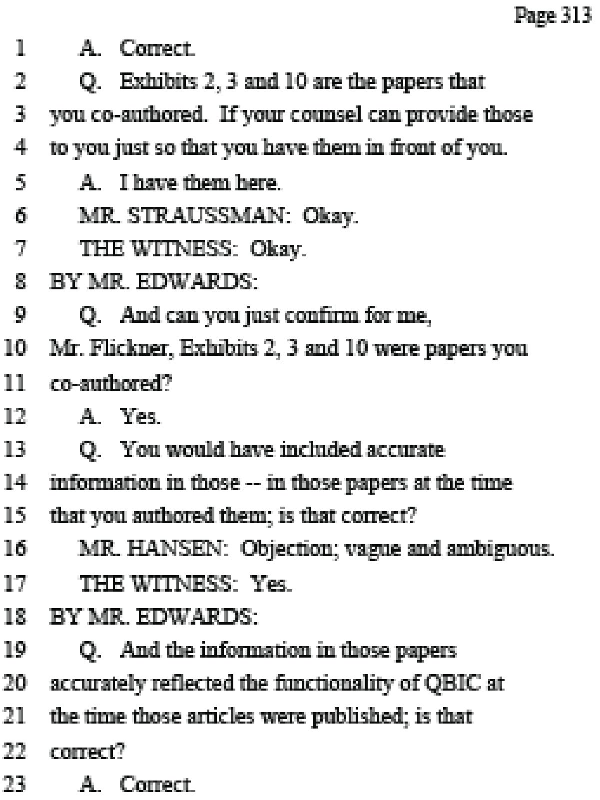 Page 313 A. Correct. Q. Exhibits 2, 3 and 10 are the papers that you co-authored. If your counsel can provide those to you just so that you have them in front of you. A. I have them here. MR. STRAUSSMAN: Okay. THE WITNESS: Okay. BY MR. EDWARDS: Q. And can you just confirm for me, Mr. Flickner, Exhibits 2, 3 and 10 were papers you co-authored? A. Yes. Q. You would have included accurate information in those -- in those papers at the time that you authored them; is that correct? MR. HANSEN: Objection; vague and ambiguous. THE WITNESS: Yes. BY MR. EDWARDS: Q. And the information in those papers accurately reflected the functionality of QBIC at the time those articles were published; is that correct? A. Correct.

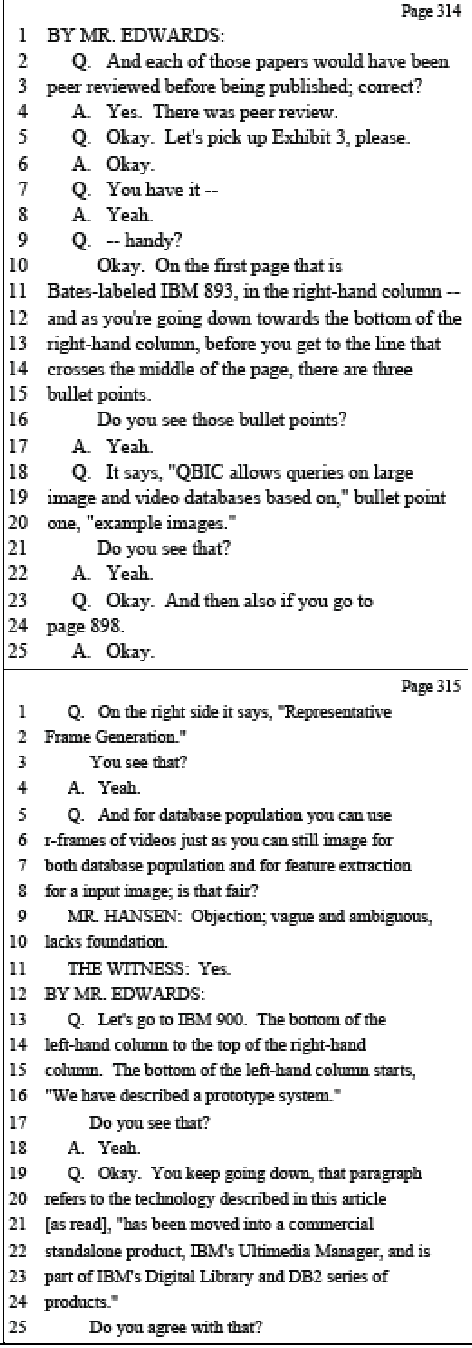

Page 314

BY MR. EDWARDS:
Q. And each of those papers would have been peer reviewed before being published; correct?
A. Yes. There was peer review.
Q. Okay. Let's pick up Exhibit 3, please.
A. Okay.
Q. You have it --
A. Yeah.
Q. -- handy?
Okay. On the first page that is Bates-labeled IBM 893, in the right-hand column -- and as you're going down towards the bottom of the right-hand column, before you get to the line that crosses the middle of the page, there are three bullet points.
Do you see those bullet points?
A. Yeah.
Q. It says, "QBIC allows queries on large image and video databases based on," bullet point one, "example images."
Do you see that?
A. Yeah.
Q. Okay. And then also if you go to page 898.
A. Okay.

Page 315

Q. On the right side it says, "Representative Frame Generation."
You see that?
A. Yeah.
Q. And for database population you can use r-frames of videos just as you can still image for both database population and for feature extraction for a input image; is that fair?
MR. HANSEN: Objection; vague and ambiguous, lacks foundation.
THE WITNESS: Yes.
BY MR. EDWARDS:
Q. Let's go to IBM 900. The bottom of the left-hand column to the top of the right-hand column. The bottom of the left-hand column starts, "We have described a prototype system."
Do you see that?
A. Yeah.
Q. Okay. You keep going down, that paragraph refers to the technology described in this article [as read], "has been moved into a commercial standalone product, IBM's Ultimedia Manager, and is part of IBM's Digital Library and DB2 series of products."
Do you agree with that?

Element	'036 Claim Recitation	Exemplary Citations in Reference
		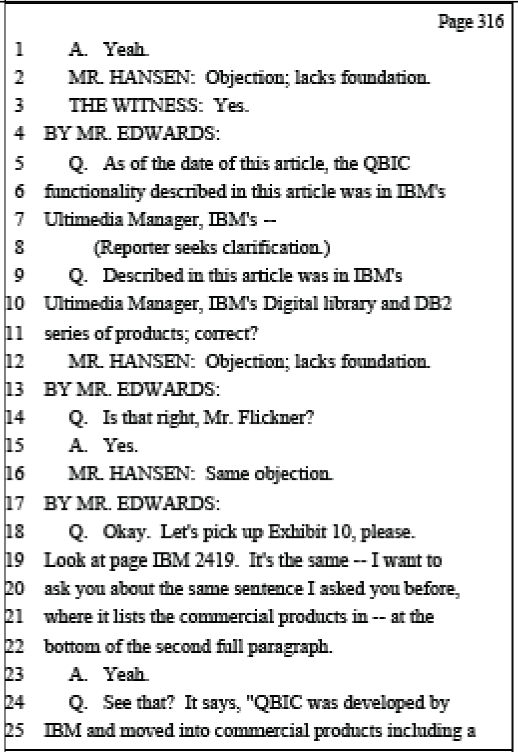 Page 316 A. Yeah. MR. HANSEN: Objection; lacks foundation. THE WITNESS: Yes. BY MR. EDWARDS: Q. As of the date of this article, the QBIC functionality described in this article was in IBM's Ultimedia Manager, IBM's -- (Reporter seeks clarification.) Q. Described in this article was in IBM's Ultimedia Manager, IBM's Digital library and DB2 series of products; correct? MR. HANSEN: Objection; lacks foundation. BY MR. EDWARDS: Q. Is that right, Mr. Flickner? A. Yes. MR. HANSEN: Same objection. BY MR. EDWARDS: Q. Okay. Let's pick up Exhibit 10, please. Look at page IBM 2419. It's the same -- I want to ask you about the same sentence I asked you before, where it lists the commercial products in -- at the bottom of the second full paragraph. A. Yeah. Q. See that? It says, "QBIC was developed by IBM and moved into commercial products including a 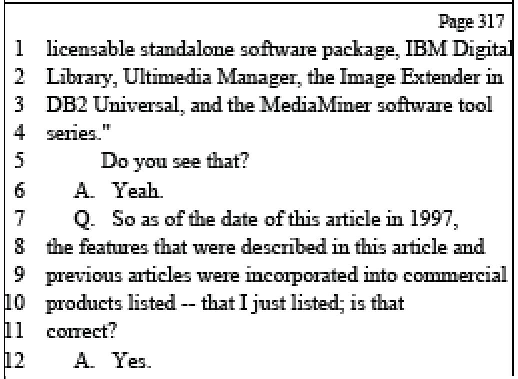 Page 317 licensable standalone software package, IBM Digital Library, Ultimedia Manager, the Image Extender in DB2 Universal, and the MediaMiner software tool series." Do you see that? A. Yeah. Q. So as of the date of this article in 1997, the features that were described in this article and previous articles were incorporated into commercial products listed -- that I just listed; is that correct? A. Yes.

Element	'036 Claim Recitation	Exemplary Citations in Reference
		IBM000039 – 41 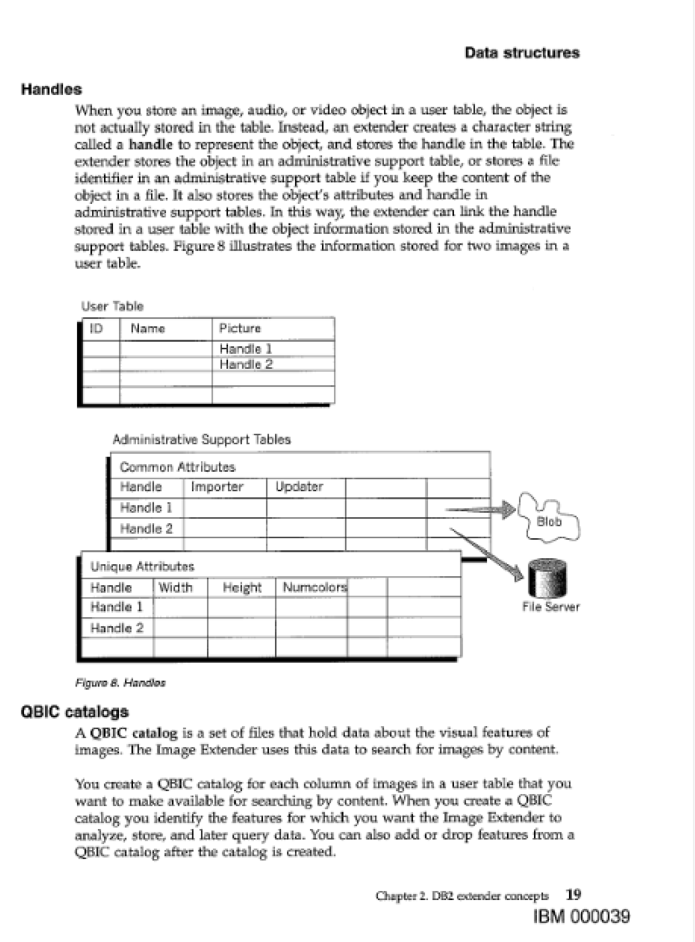

Data structures

Handles

When you store an image, audio, or video object in a user table, the object is not actually stored in the table. Instead, an extender creates a character string called a **handle** to represent the object, and stores the handle in the table. The extender stores the object in an administrative support table, or stores a file identifier in an administrative support table if you keep the content of the object in a file. It also stores the object's attributes and handle in administrative support tables. In this way, the extender can link the handle stored in a user table with the object information stored in the administrative support tables. Figure 8 illustrates the information stored for two images in a user table.

Figure 8. Handles

QBIC catalogs

A **QBIC catalog** is a set of files that hold data about the visual features of images. The Image Extender uses this data to search for images by content.

You create a QBIC catalog for each column of images in a user table that you want to make available for searching by content. When you create a QBIC catalog you identify the features for which you want the Image Extender to analyze, store, and later query data. You can also add or drop features from a QBIC catalog after the catalog is created.

Chapter 2. DB2 extender concepts 19

IBM 000039

Element	'036 Claim Recitation	Exemplary Citations in Reference
		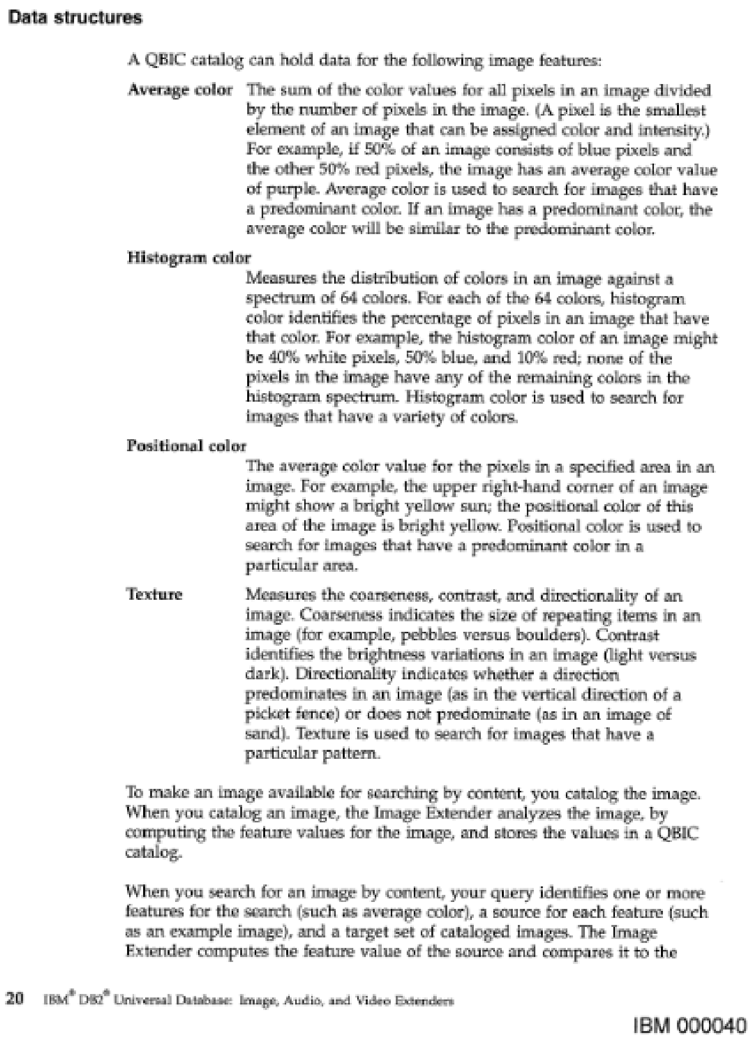

Data structures

A QBIC catalog can hold data for the following image features:

Average color	The sum of the color values for all pixels in an image divided by the number of pixels in the image. (A pixel is the smallest element of an image that can be assigned color and intensity.) For example, if 50% of an image consists of blue pixels and the other 50% red pixels, the image has an average color value of purple. Average color is used to search for images that have a predominant color. If an image has a predominant color, the average color will be similar to the predominant color.
Histogram color	Measures the distribution of colors in an image against a spectrum of 64 colors. For each of the 64 colors, histogram color identifies the percentage of pixels in an image that have that color. For example, the histogram color of an image might be 40% white pixels, 50% blue, and 10% red; none of the pixels in the image have any of the remaining colors in the histogram spectrum. Histogram color is used to search for images that have a variety of colors.
Positional color	The average color value for the pixels in a specified area in an image. For example, the upper right-hand corner of an image might show a bright yellow sun; the positional color of this area of the image is bright yellow. Positional color is used to search for images that have a predominant color in a particular area.
Texture	Measures the coarseness, contrast, and directionality of an image. Coarseness indicates the size of repeating items in an image (for example, pebbles versus boulders). Contrast identifies the brightness variations in an image (light versus dark). Directionality indicates whether a direction predominates in an image (as in the vertical direction of a picket fence) or does not predominate (as in an image of sand). Texture is used to search for images that have a particular pattern.

To make an image available for searching by content, you catalog the image. When you catalog an image, the Image Extender analyzes the image, by computing the feature values for the image, and stores the values in a QBIC catalog.

When you search for an image by content, your query identifies one or more features for the search (such as average color), a source for each feature (such as an example image), and a target set of cataloged images. The Image Extender computes the feature value of the source and compares it to the

20 IBM® DB2® Universal Database: Image, Audio, and Video Extenders

IBM 000040

Element	'036 Claim Recitation	Exemplary Citations in Reference
		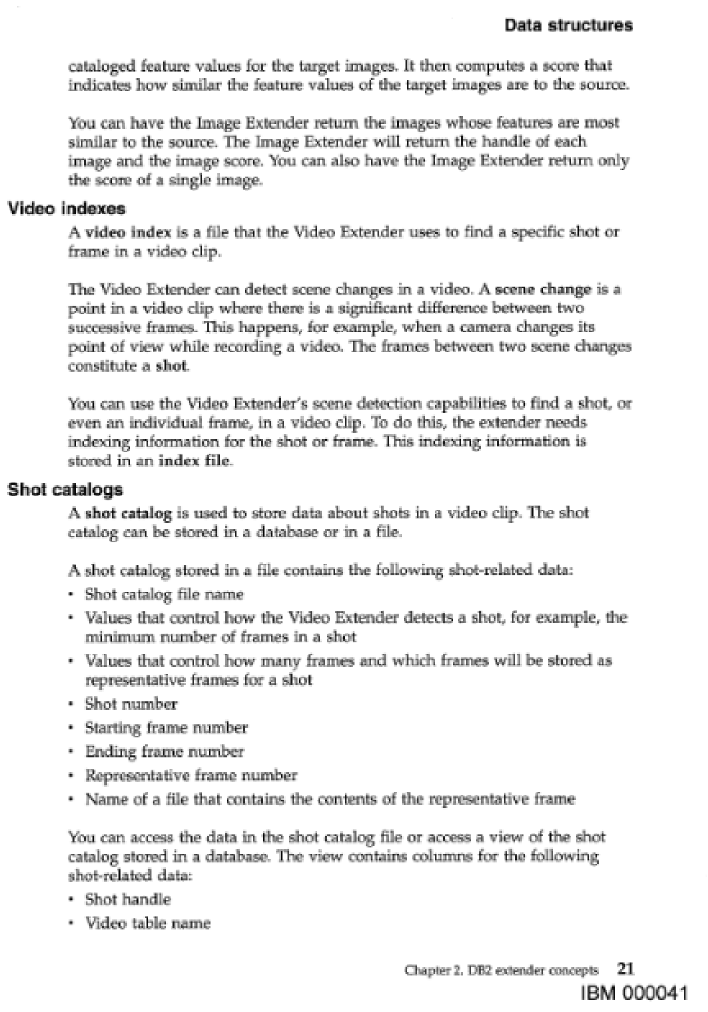 **Data structures** cataloged feature values for the target images. It then computes a score that indicates how similar the feature values of the target images are to the source. You can have the Image Extender return the images whose features are most similar to the source. The Image Extender will return the handle of each image and the image score. You can also have the Image Extender return only the score of a single image. **Video indexes** A **video index** is a file that the Video Extender uses to find a specific shot or frame in a video clip. The Video Extender can detect scene changes in a video. A **scene change** is a point in a video clip where there is a significant difference between two successive frames. This happens, for example, when a camera changes its point of view while recording a video. The frames between two scene changes constitute a **shot**. You can use the Video Extender's scene detection capabilities to find a shot, or even an individual frame, in a video clip. To do this, the extender needs indexing information for the shot or frame. This indexing information is stored in an **index file**. **Shot catalogs** A **shot catalog** is used to store data about shots in a video clip. The shot catalog can be stored in a database or in a file. A shot catalog stored in a file contains the following shot-related data: • Shot catalog file name • Values that control how the Video Extender detects a shot, for example, the minimum number of frames in a shot • Values that control how many frames and which frames will be stored as representative frames for a shot • Shot number • Starting frame number • Ending frame number • Representative frame number • Name of a file that contains the contents of the representative frame You can access the data in the shot catalog file or access a view of the shot catalog stored in a database. The view contains columns for the following shot-related data: • Shot handle • Video table name Chapter 2. DB2 extender concepts 21 IBM 000041 IBM0002264

Element	'036 Claim Recitation	Exemplary Citations in Reference
		 IBM000794

Description

Ultimedia Manager 1.1 opens new dimensions in visual queries, helping you manage your image database more efficiently. With its easy-to-use interface, even entry-level users can perform sophisticated searches. You can search by image content, data description, or a combination of both. Once the images are found, you can view the images in full resolution, use the images in another application, or read the associated textual comments.

You can use Ultimedia Manager 1.1 either as a stand-alone product or in a LAN environment. As a stand-alone product, Ultimedia Manager 1.1 employs full image classification and image query capabilities. Within the LAN environment, the Ultimedia Manager image catalogs can be sorted on a DB2 for OS/2 or DB2 for AIX database server. The image catalogs can be accessed by Windows or OS/2 clients using the Client Search Feature or the Visualizer Ultimedia Query.

Ultimedia Manager runs as an OS/2 application to a DB2 for OS/2 or DB2 for AIX database server in a LAN, or as a stand-alone application with a local DB2 for OS/2 database. The Client Search Feature runs on DOS or Windows, and accesses files in dBASE flat file format. A database manager is not required to access catalogs in dBASE file format.

Image Classification

With the Image Classifier, you can create visual indexes based on color, shape, category, and layout elements, which are areas of the image containing the patterns you want to use for searches. You can store these elements in catalogs in DB2 for OS/2 databases. The catalogs can be used locally or, through the use of utility programs, be transferred to DB2 for OS/2 or AIX servers. In addition, you can convert the catalog data into dBASE flat file format for access by Windows client applications.

The Image Classifier and Analyzer allows automatic or interactive classification of images for later image queries, flexibility in storage of catalog information, and access to multiple users. You can organize large image catalogs so that later searches for images can be performed easily. You can classify automatically using a batch program, or interactively by specifying areas of interest in an image and indicating the visual cue (color, shape, or texture) to be saved for later queries.

Image Query

Once the image catalogs are created, you can use the Image Query component to find image and business information. Business information, such as character and numeric data, are stored in the database tables using normal table population methods. Image Query is provided in Ultimedia Manager and Client Search Feature.

Image Query allows searches by three methods:

- Image content, such as color, shape, texture, or layout
- Data description, such as business data stored in the database (character or numeric) or descriptive text (category)
- Combinations of image content and data description.

Image Content: To search by image content, specify the attributes (color, shape, texture, or layout). You can drag and drop visual samples into an image query, select a color from the color wheel, or sketch a shape.

Data Description: To search by data description, enter character and numeric data values in the database tables using the simple fill-in-the-blanks query panel. You can enter data such as name, number, date, or a range of dates. You can integrate character or numeric data as part of the image query, enabling a more effective search for better results. The simple fill-in-the-blanks method for specifying desired conditions is so easy even non-programmers can use database table information as part of their searches.

Combinations: Image queries using either image content or data description, or combinations of both, allow you to search through thousands of images more quickly and effectively than the traditional text and keyword search methods.

Image Formats

Ultimedia Manager 1.1 understands these image formats:

- Audio Visual Connection (AVC) images
- OS/2 bitmap
- Windows 3.x bitmap
- PCX
- Intel and Motorola TIFF
- TARGA TGA
- TIFF (FAX)
- GIF
- Kodak Photo-CD

Product Positioning

The advanced image search technology in Ultimedia Manager, query by image content, was developed in conjunction with the IBM Almaden Research Center. It provides function not available with conventional, text-only image management systems.

Ultimedia Manager fulfills image search and retrieval needs of business users and creative professionals through use of online image libraries or collections. It is intended for users who have large image libraries and need fast, flexible access to their image data, users who may also want to add character or numeric data to their image queries, and users who feel limited by text-based systems.

Ultimedia Manager 1.1 works with the DB2 Family of products and the Visualizer Ultimedia Query product. Catalogs created by Ultimedia Manager can be stored in DB2 for OS/2 database tables. Image catalogs also can be stored in DB2 for OS/2 or DB2 for AIX database servers in client/server environments.

OS/2 and Windows clients in a LAN environment using the Client Search Feature and the Visualizer Ultimedia Query can access catalogs created with Ultimedia Manager 1.1. Visualizer Ultimedia Query provides users with the capability to:

- Add multimedia columns (such as image, video, audio, and formatted text) to DB2 for OS/2 database tables
- Query and display images
- View or listen to video and audio information
- Display formatted text

Users of Ultimedia Manager 1.1 and Client Search Feature are focused on image management and query applications. Visualizer Ultimedia Query users are

Element	'036 Claim Recitation	Exemplary Citations in Reference
		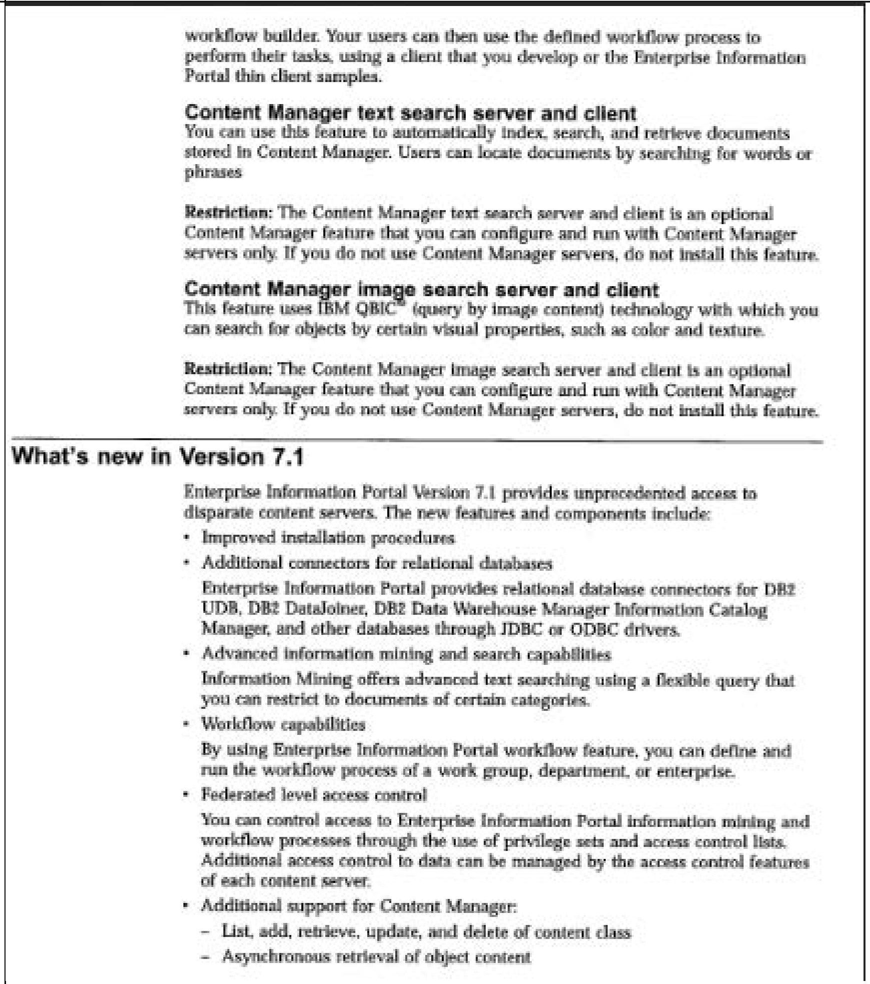 workflow builder. Your users can then use the defined workflow process to perform their tasks, using a client that you develop or the Enterprise Information Portal thin client samples. **Content Manager text search server and client** You can use this feature to automatically index, search, and retrieve documents stored in Content Manager. Users can locate documents by searching for words or phrases **Restriction:** The Content Manager text search server and client is an optional Content Manager feature that you can configure and run with Content Manager servers only. If you do not use Content Manager servers, do not install this feature. **Content Manager image search server and client** This feature uses IBM QBIC® (query by image content) technology with which you can search for objects by certain visual properties, such as color and texture. **Restriction:** The Content Manager image search server and client is an optional Content Manager feature that you can configure and run with Content Manager servers only. If you do not use Content Manager servers, do not install this feature. **What's new in Version 7.1** Enterprise Information Portal Version 7.1 provides unprecedented access to disparate content servers. The new features and components include: • Improved installation procedures • Additional connectors for relational databases Enterprise Information Portal provides relational database connectors for DB2 UDB, DB2 DataJoiner, DB2 Data Warehouse Manager Information Catalog Manager, and other databases through JDBC or ODBC drivers. • Advanced information mining and search capabilities Information Mining offers advanced text searching using a flexible query that you can restrict to documents of certain categories. • Workflow capabilities By using Enterprise Information Portal workflow feature, you can define and run the workflow process of a work group, department, or enterprise. • Federated level access control You can control access to Enterprise Information Portal information mining and workflow processes through the use of privilege sets and access control lists. Additional access control to data can be managed by the access control features of each content server. • Additional support for Content Manager: – List, add, retrieve, update, and delete of content class – Asynchronous retrieval of object content To the extent it is found that QBIC does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that QBIC does not anticipate this claim, QBIC renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other

Element	'036 Claim Recitation	Exemplary Citations in Reference
		prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine QBIC with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1B]	an identification platform coupled with the target database, and that	QBIC discloses an identification platform coupled with the target database. The QBIC Project: Querying Images by Content Using Color, Texture, and Shape (IBM 0002390–404) at 2391: 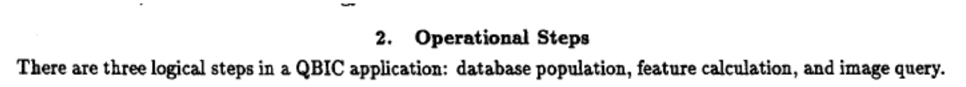 **2. Operational Steps** There are three logical steps in a QBIC application: database population, feature calculation, and image query. IBM0002390 at 2391-92: 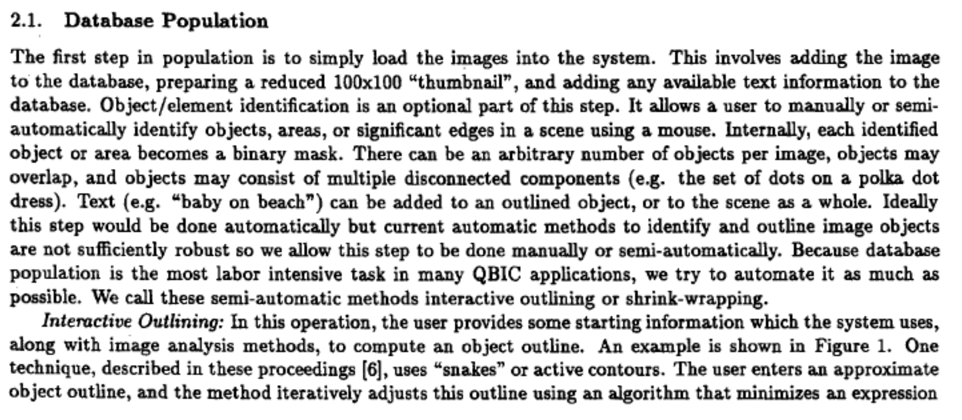 **2.1. Database Population** The first step in population is to simply load the images into the system. This involves adding the image to the database, preparing a reduced 100x100 "thumbnail", and adding any available text information to the database. Object/element identification is an optional part of this step. It allows a user to manually or semi-automatically identify objects, areas, or significant edges in a scene using a mouse. Internally, each identified object or area becomes a binary mask. There can be an arbitrary number of objects per image, objects may overlap, and objects may consist of multiple disconnected components (e.g. the set of dots on a polka dot dress). Text (e.g. "baby on beach") can be added to an outlined object, or to the scene as a whole. Ideally this step would be done automatically but current automatic methods to identify and outline image objects are not sufficiently robust so we allow this step to be done manually or semi-automatically. Because database population is the most labor intensive task in many QBIC applications, we try to automate it as much as possible. We call these semi-automatic methods interactive outlining or shrink-wrapping. *Interactive Outlining:* In this operation, the user provides some starting information which the system uses, along with image analysis methods, to compute an object outline. An example is shown in Figure 1. One technique, described in these proceedings [6], uses "snakes" or active contours. The user enters an approximate object outline, and the method iteratively adjusts this outline using an algorithm that minimizes an expression

Element	'036 Claim Recitation	Exemplary Citations in Reference
		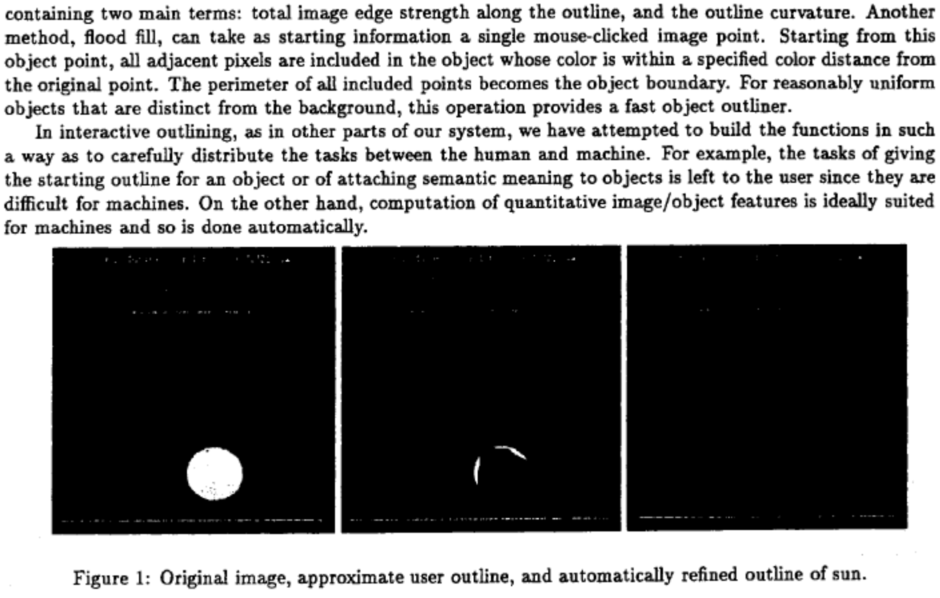 containing two main terms: total image edge strength along the outline, and the outline curvature. Another method, flood fill, can take as starting information a single mouse-clicked image point. Starting from this object point, all adjacent pixels are included in the object whose color is within a specified color distance from the original point. The perimeter of all included points becomes the object boundary. For reasonably uniform objects that are distinct from the background, this operation provides a fast object outliner. In interactive outlining, as in other parts of our system, we have attempted to build the functions in such a way as to carefully distribute the tasks between the human and machine. For example, the tasks of giving the starting outline for an object or of attaching semantic meaning to objects is left to the user since they are difficult for machines. On the other hand, computation of quantitative image/object features is ideally suited for machines and so is done automatically. Figure 1: Original image, approximate user outline, and automatically refined outline of sun. IBM0002390 at 2392-93:

Element	’036 Claim Recitation	Exemplary Citations in Reference
		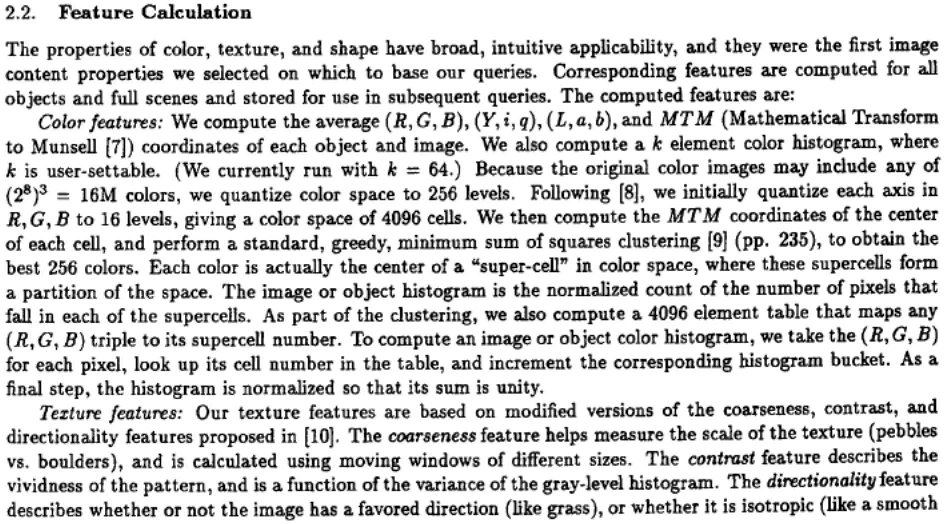 2.2. **Feature Calculation** The properties of color, texture, and shape have broad, intuitive applicability, and they were the first image content properties we selected on which to base our queries. Corresponding features are computed for all objects and full scenes and stored for use in subsequent queries. The computed features are: *Color features:* We compute the average (R, G, B), (Y, i, q), (L, a, b), and MTM (Mathematical Transform to Munsell [7]) coordinates of each object and image. We also compute a k element color histogram, where k is user-settable. (We currently run with $k = 64$.) Because the original color images may include any of $(2^8)^3 = 16$M colors, we quantize color space to 256 levels. Following [8], we initially quantize each axis in R, G, B to 16 levels, giving a color space of 4096 cells. We then compute the MTM coordinates of the center of each cell, and perform a standard, greedy, minimum sum of squares clustering [9] (pp. 235), to obtain the best 256 colors. Each color is actually the center of a "super-cell" in color space, where these supercells form a partition of the space. The image or object histogram is the normalized count of the number of pixels that fall in each of the supercells. As part of the clustering, we also compute a 4096 element table that maps any (R, G, B) triple to its supercell number. To compute an image or object color histogram, we take the (R, G, B) for each pixel, look up its cell number in the table, and increment the corresponding histogram bucket. As a final step, the histogram is normalized so that its sum is unity. *Texture features:* Our texture features are based on modified versions of the coarseness, contrast, and directionality features proposed in [10]. The *coarseness* feature helps measure the scale of the texture (pebbles vs. boulders), and is calculated using moving windows of different sizes. The *contrast* feature describes the vividness of the pattern, and is a function of the variance of the gray-level histogram. The *directionality* feature describes whether or not the image has a favored direction (like grass), or whether it is isotropic (like a smooth

Element	'036 Claim Recitation	Exemplary Citations in Reference
		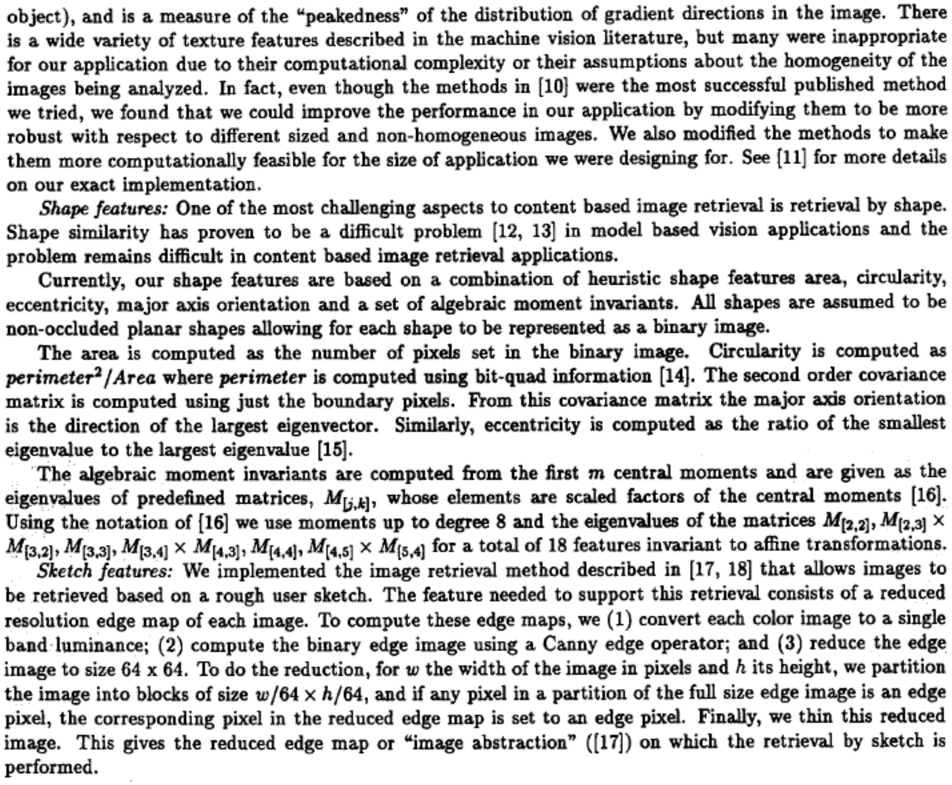 object), and is a measure of the "peakedness" of the distribution of gradient directions in the image. There is a wide variety of texture features described in the machine vision literature, but many were inappropriate for our application due to their computational complexity or their assumptions about the homogeneity of the images being analyzed. In fact, even though the methods in [10] were the most successful published method we tried, we found that we could improve the performance in our application by modifying them to be more robust with respect to different sized and non-homogeneous images. We also modified the methods to make them more computationally feasible for the size of application we were designing for. See [11] for more details on our exact implementation. *Shape features:* One of the most challenging aspects to content based image retrieval is retrieval by shape. Shape similarity has proven to be a difficult problem [12, 13] in model based vision applications and the problem remains difficult in content based image retrieval applications. Currently, our shape features are based on a combination of heuristic shape features area, circularity, eccentricity, major axis orientation and a set of algebraic moment invariants. All shapes are assumed to be non-occluded planar shapes allowing for each shape to be represented as a binary image. The area is computed as the number of pixels set in the binary image. Circularity is computed as $perimeter^2/Area$ where $perimeter$ is computed using bit-quad information [14]. The second order covariance matrix is computed using just the boundary pixels. From this covariance matrix the major axis orientation is the direction of the largest eigenvector. Similarly, eccentricity is computed as the ratio of the smallest eigenvalue to the largest eigenvalue [15]. The algebraic moment invariants are computed from the first m central moments and are given as the eigenvalues of predefined matrices, $M_{[j,k]}$, whose elements are scaled factors of the central moments [16]. Using the notation of [16] we use moments up to degree 8 and the eigenvalues of the matrices $M_{[2,2]}, M_{[2,3]} \times M_{[3,2]}, M_{[3,3]}, M_{[3,4]} \times M_{[4,3]}, M_{[4,4]}, M_{[4,5]} \times M_{[5,4]}$ for a total of 18 features invariant to affine transformations. *Sketch features:* We implemented the image retrieval method described in [17, 18] that allows images to be retrieved based on a rough user sketch. The feature needed to support this retrieval consists of a reduced resolution edge map of each image. To compute these edge maps, we (1) convert each color image to a single band luminance; (2) compute the binary edge image using a Canny edge operator; and (3) reduce the edge image to size 64 x 64. To do the reduction, for w the width of the image in pixels and h its height, we partition the image into blocks of size $w/64 \times h/64$, and if any pixel in a partition of the full size edge image is an edge pixel, the corresponding pixel in the reduced edge map is set to an edge pixel. Finally, we thin this reduced image. This gives the reduced edge map or "image abstraction" ([17]) on which the retrieval by sketch is performed. IBM0002390 at 2393-95:

Element	'036 Claim Recitation	Exemplary Citations in Reference
		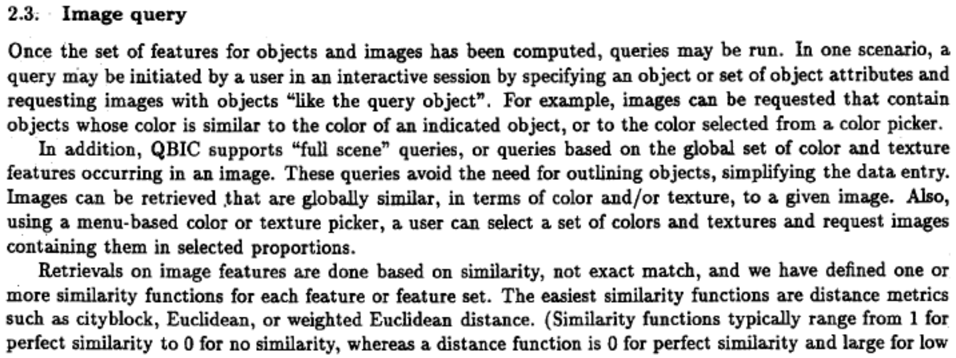 **2.3. Image query** Once the set of features for objects and images has been computed, queries may be run. In one scenario, a query may be initiated by a user in an interactive session by specifying an object or set of object attributes and requesting images with objects "like the query object". For example, images can be requested that contain objects whose color is similar to the color of an indicated object, or to the color selected from a color picker. In addition, QBIC supports "full scene" queries, or queries based on the global set of color and texture features occurring in an image. These queries avoid the need for outlining objects, simplifying the data entry. Images can be retrieved that are globally similar, in terms of color and/or texture, to a given image. Also, using a menu-based color or texture picker, a user can select a set of colors and textures and request images containing them in selected proportions. Retrievals on image features are done based on similarity, not exact match, and we have defined one or more similarity functions for each feature or feature set. The easiest similarity functions are distance metrics such as cityblock, Euclidean, or weighted Euclidean distance. (Similarity functions typically range from 1 for perfect similarity to 0 for no similarity, whereas a distance function is 0 for perfect similarity and large for low

Element	'036 Claim Recitation	Exemplary Citations in Reference
		similarity, so we must convert as necessary to obtain consistent measures.) Different similarity measures may be used, and they may vary from application to application. The matching methods and similarity functions for the different feature types are described below: *Color:* For average color, the distance between a query object and database object is weighted Euclidean distance. The weights are user adjustable so that we may, for example, downweight the L component of an (L, a, b) query. Best results, as judged by our subjective evaluation in on-going experiments, are obtained by representing the average color in the MTM color space, and by inversely weighting each component by its standard deviation over the samples in the database. Improved results are obtained when color matching is done not on average color, but on the distribution of colors that occur in an object or image. In this way, a red and blue object matches other red and blue objects better than to a purple object. To match color histograms, we use a method similar to the one described in [8]. Let X be the query histogram and Y the histogram of an item in the database, both normalized. We compute the element by element difference histogram Z. Then the similarity between X and Y is given by $\|Z\| = Z^T A Z$, where A is a symmetric color similarity matrix with $$a(i,j) \equiv 1 - d(c_i, c_j)/d_{max}$$ where c_i and c_j are the ith and jth colors in the color histograms, and $d(c_i, c_j)$ is the MTM color distance, and d_{max} is the maximum distance between any two colors ([8] used Luv distance.) This metric gives the (weighted) length of the difference vector between X and Y, weighted by A, which accounts for the perceptual distance between different pairs of colors. The result is a measure that accounts for both the difference between the amounts in any given color (e.g. a particular shade of red), as well as differences between similar colors (e.g. red and orange). Based on our experience, we also found it useful for the user to be able to request images with $x\%$ of color 1, $y\%$ of color 2, etc. where both the percentages and the colors are explicitly specified by the user. This allows searches such as for a beach scene (25% white (for the sand) and 50% blue (for the sky and water)), and for many scenes for which the user has seen before, and for which he/she can remember approximate color percentages. The color picker that allows this type of query is shown in Figure 2. 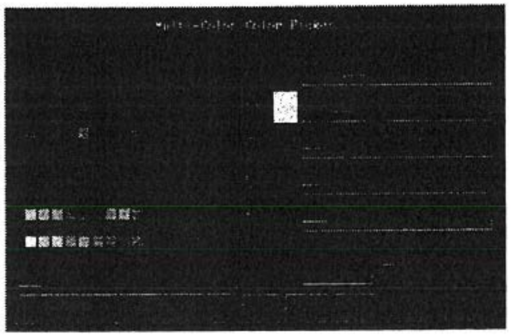 Figure 2: Multi-Color Color Picker with two selected colors. The palette on the left displays the set of selectable colors. Once a color is selected, it appears in a box on the right, and the amount of the color can be specified by the sliders. In this case, 25% light blue and 15% white have been selected. *Texture:* Texture distance is computed as weighted Euclidean distance in the three dimensional texture space. The most common weighting factors are the inverse variances for each component, computed over the

Element	'036 Claim Recitation	Exemplary Citations in Reference
		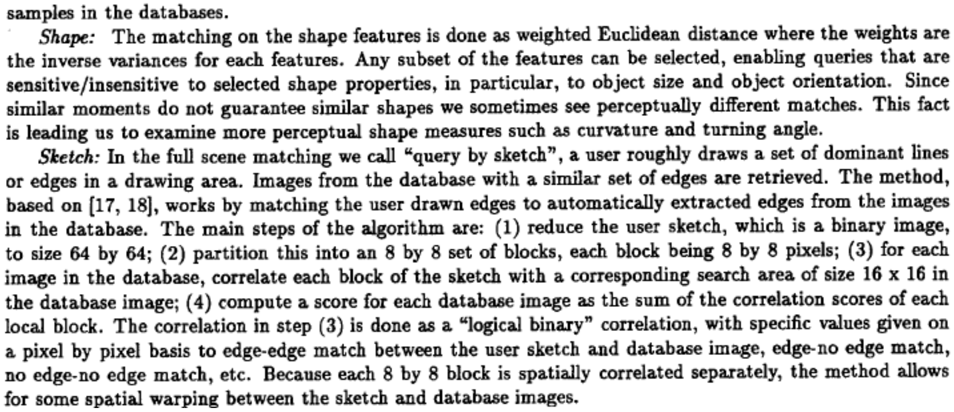 samples in the databases. *Shape:* The matching on the shape features is done as weighted Euclidean distance where the weights are the inverse variances for each features. Any subset of the features can be selected, enabling queries that are sensitive/insensitive to selected shape properties, in particular, to object size and object orientation. Since similar moments do not guarantee similar shapes we sometimes see perceptually different matches. This fact is leading us to examine more perceptual shape measures such as curvature and turning angle. *Sketch:* In the full scene matching we call "query by sketch", a user roughly draws a set of dominant lines or edges in a drawing area. Images from the database with a similar set of edges are retrieved. The method, based on [17, 18], works by matching the user drawn edges to automatically extracted edges from the images in the database. The main steps of the algorithm are: (1) reduce the user sketch, which is a binary image, to size 64 by 64; (2) partition this into an 8 by 8 set of blocks, each block being 8 by 8 pixels; (3) for each image in the database, correlate each block of the sketch with a corresponding search area of size 16 x 16 in the database image; (4) compute a score for each database image as the sum of the correlation scores of each local block. The correlation in step (3) is done as a "logical binary" correlation, with specific values given on a pixel by pixel basis to edge-edge match between the user sketch and database image, edge-no edge match, no edge-no edge match, etc. Because each 8 by 8 block is spatially correlated separately, the method allows for some spatial warping between the sketch and database images. Query by Image and Video Content: The QBIC System (IBM 0000893–902): "To achieve this functionality, QBIC has two main components: database population (the process of creating an image database) and database query. During the population, images and videos are processed to extract features describing their content-colors, textures, shapes, and camera and object motion-and the features are stored in a database. During the query, the user composes a query graphically. Features are generated from the graphical query and then input to a matching engine that finds images or videos from the database with similar features. Figure 2 shows the system architecture." *Id.* at p. 3.

Element	'036 Claim Recitation	Exemplary Citations in Reference
		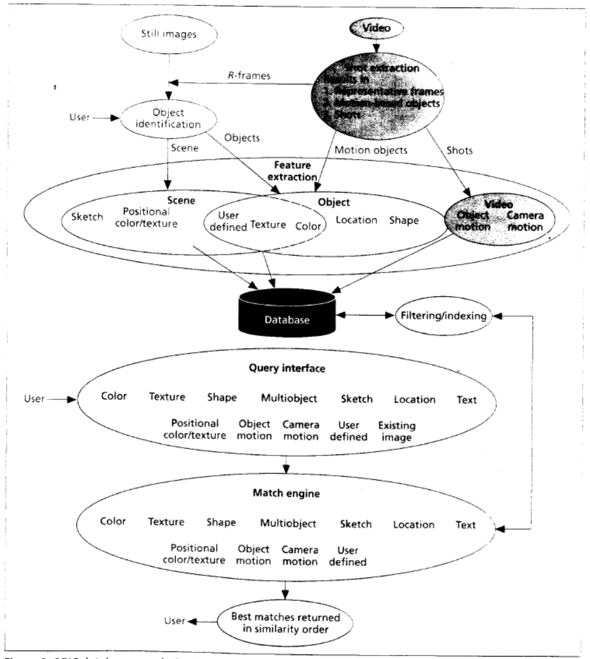 Figure 2. QBIC database population (top) and query (bottom) architecture. *Id*. at Fig. 2. "For both population and query, the QBIC data model has • still images or scenes (full images) that contain objects (subsets of an image), and • video shots that consist of sets of contiguous frames and contain motion objects. For still images, the QBIC data model distinguishes between 'scenes' (or images) and 'objects.' A scene is an image or single representative frame of video. An object is a part of a scene-for example, the fox in Figure 3-or a moving entity in a video. For still image database population, features are extracted from images and objects and stored in a database as shown in the top left part of Figure 2. Videos are broken into clips called shots. Representative frames, or r-frames, are generated for each extracted shot. R-frames are treated as still images, and features are extracted and stored in the database. Further processing of shots generates motion objects-for example, a car moving across the screen. Queries are allowed

Element	'036 Claim Recitation	Exemplary Citations in Reference
		on objects ('Find images with a red, round object'), scenes ('Find images that have approximately 30-percent red and IS-percent blue colors'), shots ('Find all shots panning from left to right'), or any combination ('Find images that have 30 percent red and contain a blue textured object')." *Id.* at p. 3. "In still image database population, the images are reduced to a standard-sized icon called a thumbnail and annotated with any available text information. Object identification is an optional but key part of this step. It lets users manually, semiautomatically, or fully automatically identify interesting regions-which we call objects-in the images. . . . For video data, database population has three major components: • shot detection, • representative frame creation for each shot, and • derivation of a layered representation of coherently moving structures/ objects. . . . Once the shot boundaries have been detected, each shot is represented using an r-frame. R-frames are used for several purposes. First, during database population, r-frames are treated as still images in which objects can be identified by using the previously described methods." *Id.* at pp. 3–6. "**SAMPLE QUERIES** For each full-scene image, identified image object, r-frame, and identified video object resulting from the above processing, a set of features is computed to allow content-based queries. The features are computed and stored during database population. We present a brief description of the features and the associated queries. Mathematical details on the features and matching methods can be found in Ashley et al. and Niblack et al. Average color queries let users find images or objects that are similar to a selected color, say from a color wheel, or to the color of an object. The feature used in the query is a 3D vector of Munsell color coordinates. Histogram color queries return items with matching color distributions—say, a fabric pattern with approximately 40 percent red and 20 percent blue. For this case, the underlying feature is a 256-element histogram computed over a quantized version of the color space. Figure 7 shows a histogram query on still images and a color query on video r-

Element	'036 Claim Recitation	Exemplary Citations in Reference
		frames. Note that in the query specification for the histogram query of Figure 7, the user has selected percentages of two colors (blue and white) by adjusting sliders. Using such a query, an advertising agent could, for example, search for a picture of a beach scene, one predominantly blue (for sky and water) and white (for sand and clouds); or find images with similar color spreads for a uniform ad campaign. The average color query demonstrates a query against a video shot database where the user is searching for red r-frames. Again, the query specification is on the left and the best hits are on the right. Figure 8 shows an example texture query. In this case, the query is specified by selecting from a sampler—a set of prestored example images. The underlying texture features are mathematical representations of coarseness, contrast, and directionality features. Coarseness measures the scale of a texture (pebbles versus boulders), contrast describes its vividness, and directionality describes whether it has a favored direction (like grass) or not (like a smooth object). An object shape query is shown in Figure 8. In this case, the query specification is the drawn shape on the left. Area, circularity, eccentricity, major-axis direction, features derived from the object moments, and a set of tangent angles around the object perimeter are the features used to characterize and match shapes. Figure 9 illustrates query by sketch. In this case, the query specification is a freehand drawing of the dominant lines and edges in the image. The sketch feature is an automatically extracted reduced-resolution 'edge map.' Matching is done by using a template-matching technique. A multiobject query asking for images that contain both a red round object and a green textured object is shown in the bottom of Figure 9. The features are standard color and texture. The matching is done by combining the color and texture distances. Combining distances is applied to arbitrary sets of objects and features to implement logical And semantics. . . . . **ANNOTATION AND DATABASE POPULATION TOOLS.** Automatic methods (such as our Positional Color query) that don't rely on object identification, methods that identify objects automatically as in the museum image example, fast and easy-to-use semiautomatic outlining tools, and motion-based

Element	'036 Claim Recitation	Exemplary Citations in Reference
		segmentation algorithms will enable additional application areas. **FEATURE EXTRACTION AND MATCHING METHODS**. New mathematical representations of video, image, and object attributes that capture 'interesting' features for retrieval are needed. Features that describe new image properties such as alternate texture measures or that are based on fractals or wavelet representations, for example, may offer advantages of representation, indexability, and ease of similarity matching." *Id.* at pp. 7–8. Image, Audio, and Video Extenders Administration and Programming (IBM 000001–617): "Your applications can even search for images by content. Imagine an application that uses visual examples to search for images. With such an application, users could select an example image and have the application find other images that have colors or textures similar to those in the example. With DB2 extenders' Query by Image Content (QBIC) capability, you can create applications that search for images in this visual way." *Id.* at p. 4. "To make an image available for searching by content, you catalog the image. When you catalog an image, the Image Extender analyzes the image, by computing the feature values for the image, and stores the values in a QBIC catalog. When you search for an image by content, your query identifies one or more features for the search (such as average color), a source for each feature (such as an example image), and a target set of cataloged images. The Image Extender computes the feature value of the source and compares it to the cataloged feature values for the target images. It then computes a score that indicates how similar the feature values of the target images are to the source. You can have the Image Extender return the images whose features are most similar to the source. The Image Extender will return the handle of each image and the image score. You can also have the Image Extender return only the score of a single image." *Id.* at pp. 20–21.

Element	'036 Claim Recitation	Exemplary Citations in Reference
		QBIC Demo (IBM000747):  Flickner 8/29/23 Depo at Ex. 6 at 1:

Element	'036 Claim Recitation	Exemplary Citations in Reference
		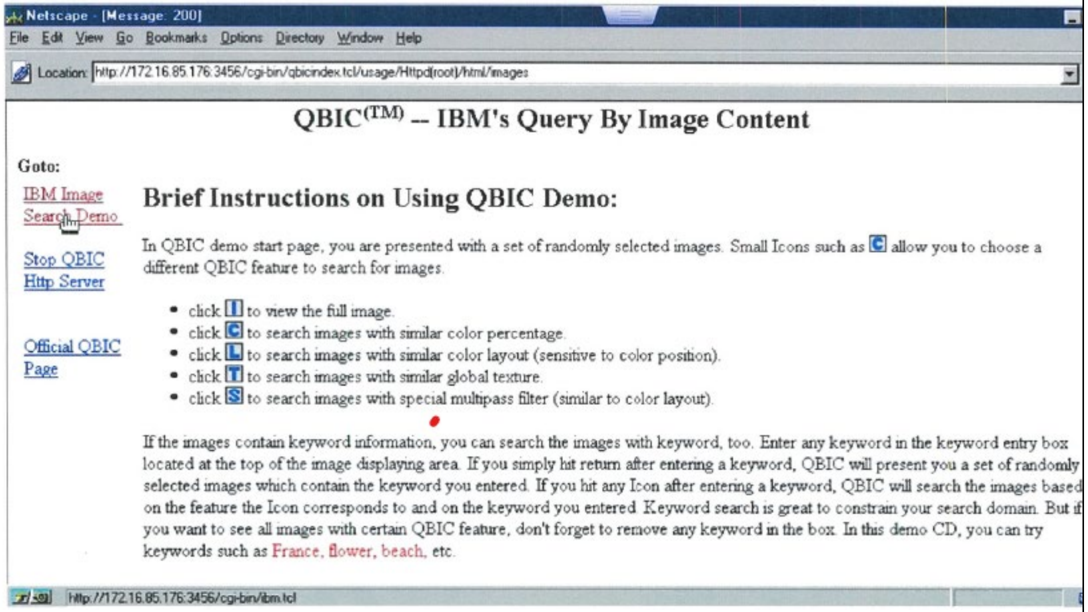 QBIC$^{(TM)}$ -- IBM's Query By Image Content Goto: IBM Image Search Demo Stop QBIC Http Server Official QBIC Page **Brief Instructions on Using QBIC Demo:** In QBIC demo start page, you are presented with a set of randomly selected images. Small Icons such as C allow you to choose a different QBIC feature to search for images. • click I to view the full image. • click C to search images with similar color percentage. • click L to search images with similar color layout (sensitive to color position). • click T to search images with similar global texture. • click S to search images with special multipass filter (similar to color layout). If the images contain keyword information, you can search the images with keyword, too. Enter any keyword in the keyword entry box located at the top of the image displaying area. If you simply hit return after entering a keyword, QBIC will present you a set of randomly selected images which contain the keyword you entered. If you hit any Icon after entering a keyword, QBIC will search the images based on the feature the Icon corresponds to and on the keyword you entered. Keyword search is great to constrain your search domain. But if you want to see all images with certain QBIC feature, don't forget to remove any keyword in the box. In this demo CD, you can try keywords such as France, flower, beach, etc. *Id.* at 2:

Element	'036 Claim Recitation	Exemplary Citations in Reference
		 *Id.* at 305: 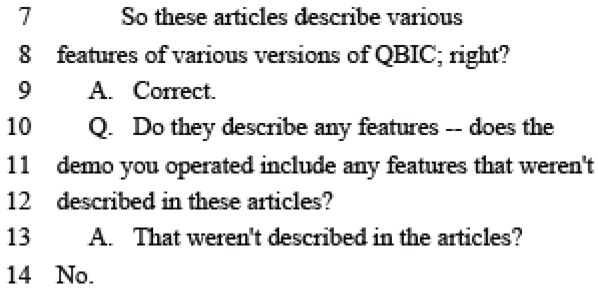So these articles describe various features of various versions of QBIC; right? A. Correct. Q. Do they describe any features -- does the demo you operated include any features that weren't described in these articles? A. That weren't described in the articles? No. *Id.* at 35:

Element	'036 Claim Recitation	Exemplary Citations in Reference
		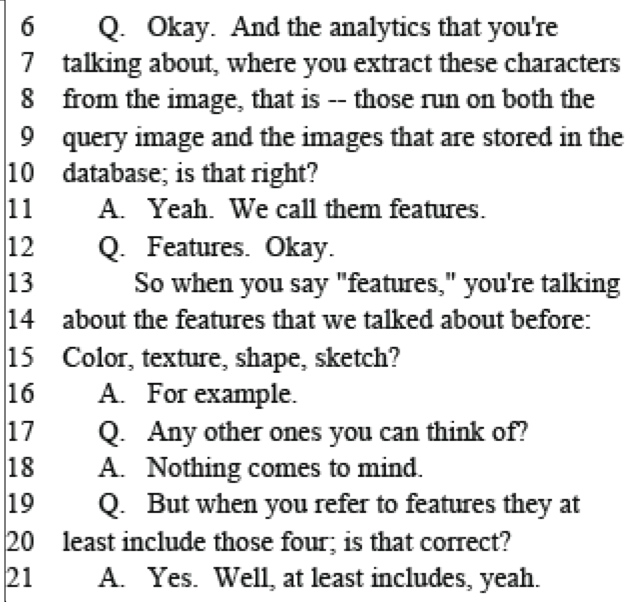 Q. Okay. And the analytics that you're talking about, where you extract these characters from the image, that is -- those run on both the query image and the images that are stored in the database; is that right? A. Yeah. We call them features. Q. Features. Okay. So when you say "features," you're talking about the features that we talked about before: Color, texture, shape, sketch? A. For example. Q. Any other ones you can think of? A. Nothing comes to mind. Q. But when you refer to features they at least include those four; is that correct? A. Yes. Well, at least includes, yeah. *Id.* at 109-110:

Element	'036 Claim Recitation	Exemplary Citations in Reference
		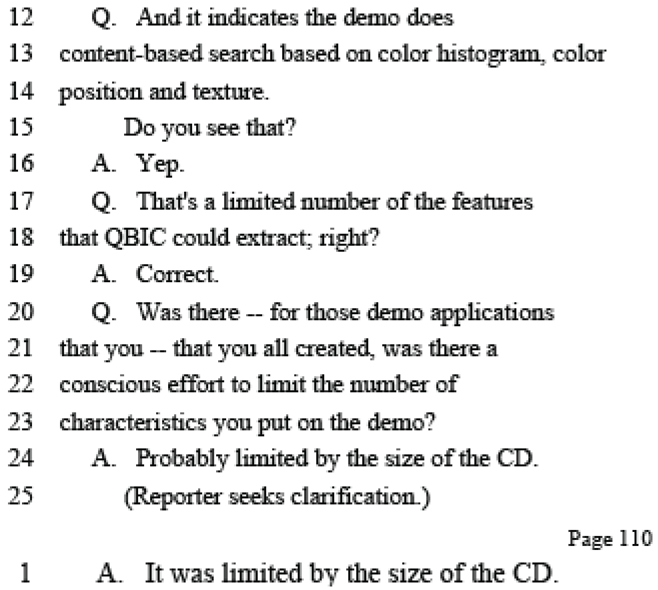 Q. And it indicates the demo does content-based search based on color histogram, color position and texture. Do you see that? A. Yep. Q. That's a limited number of the features that QBIC could extract; right? A. Correct. Q. Was there -- for those demo applications that you -- that you all created, was there a conscious effort to limit the number of characteristics you put on the demo? A. Probably limited by the size of the CD. (Reporter seeks clarification.) Page 110 A. It was limited by the size of the CD. *Id.* at 311: 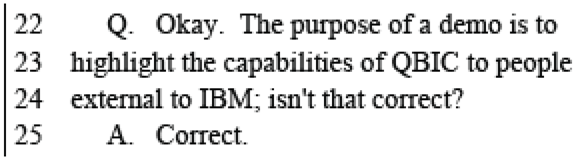 Q. Okay. The purpose of a demo is to highlight the capabilities of QBIC to people external to IBM; isn't that correct? A. Correct. *Id.* at 313-317:

Element	'036 Claim Recitation	Exemplary Citations in Reference
		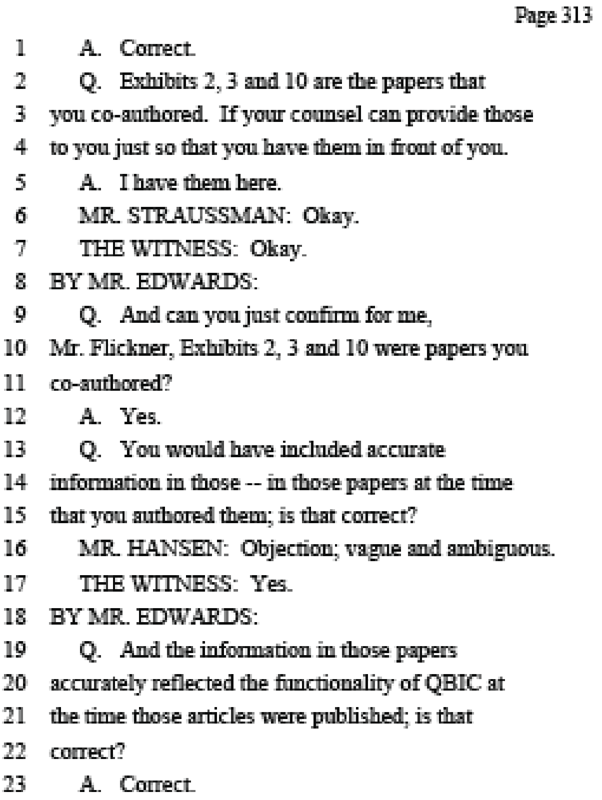 Page 313 A. Correct. Q. Exhibits 2, 3 and 10 are the papers that you co-authored. If your counsel can provide those to you just so that you have them in front of you. A. I have them here. MR. STRAUSSMAN: Okay. THE WITNESS: Okay. BY MR. EDWARDS: Q. And can you just confirm for me, Mr. Flickner, Exhibits 2, 3 and 10 were papers you co-authored? A. Yes. Q. You would have included accurate information in those -- in those papers at the time that you authored them; is that correct? MR. HANSEN: Objection; vague and ambiguous. THE WITNESS: Yes. BY MR. EDWARDS: Q. And the information in those papers accurately reflected the functionality of QBIC at the time those articles were published; is that correct? A. Correct.

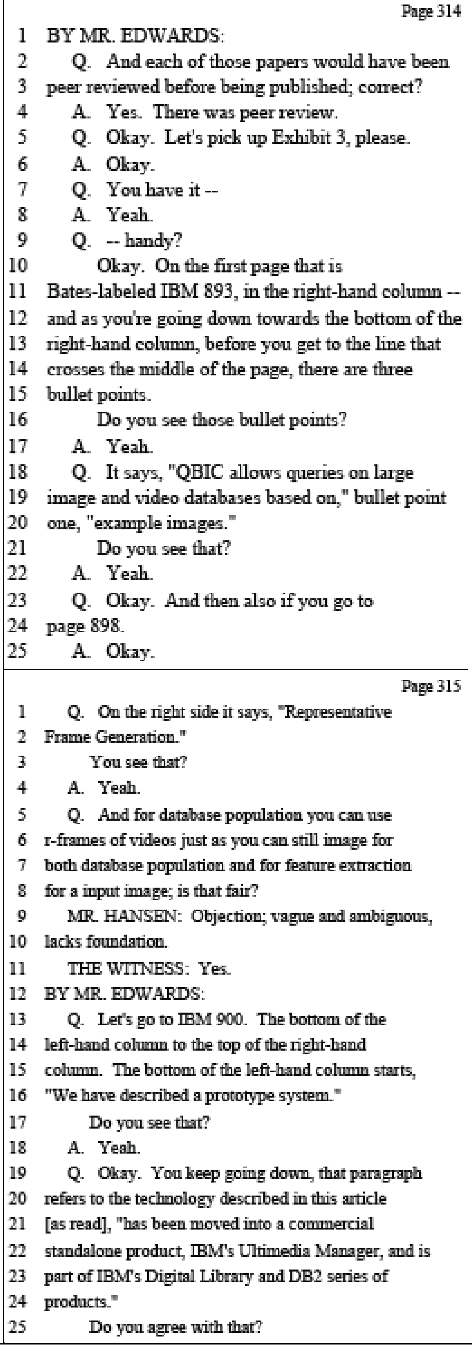

Page 314

BY MR. EDWARDS:
Q. And each of those papers would have been peer reviewed before being published; correct?
A. Yes. There was peer review.
Q. Okay. Let's pick up Exhibit 3, please.
A. Okay.
Q. You have it --
A. Yeah.
Q. -- handy?
Okay. On the first page that is Bates-labeled IBM 893, in the right-hand column -- and as you're going down towards the bottom of the right-hand column, before you get to the line that crosses the middle of the page, there are three bullet points.
Do you see those bullet points?
A. Yeah.
Q. It says, "QBIC allows queries on large image and video databases based on," bullet point one, "example images."
Do you see that?
A. Yeah.
Q. Okay. And then also if you go to page 898.
A. Okay.

Page 315

Q. On the right side it says, "Representative Frame Generation."
You see that?
A. Yeah.
Q. And for database population you can use r-frames of videos just as you can still image for both database population and for feature extraction for a input image; is that fair?
MR. HANSEN: Objection; vague and ambiguous, lacks foundation.
THE WITNESS: Yes.
BY MR. EDWARDS:
Q. Let's go to IBM 900. The bottom of the left-hand column to the top of the right-hand column. The bottom of the left-hand column starts, "We have described a prototype system."
Do you see that?
A. Yeah.
Q. Okay. You keep going down, that paragraph refers to the technology described in this article [as read], "has been moved into a commercial standalone product, IBM's Ultimedia Manager, and is part of IBM's Digital Library and DB2 series of products."
Do you agree with that?

Element	'036 Claim Recitation	Exemplary Citations in Reference
		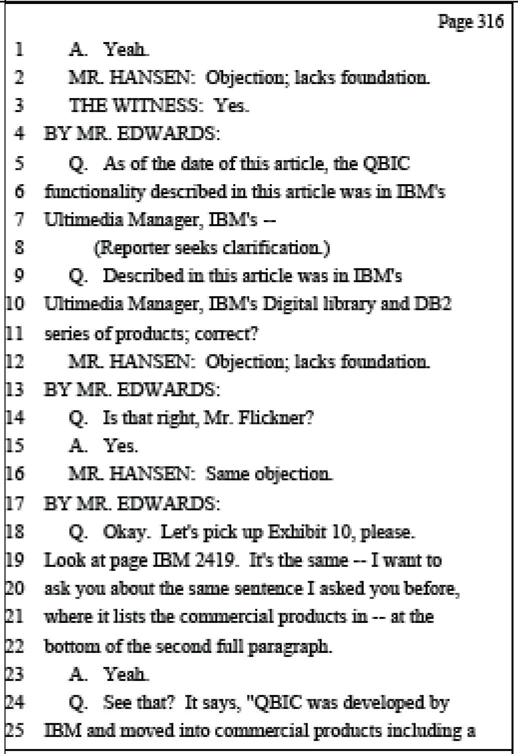 Page 316 A. Yeah. MR. HANSEN: Objection; lacks foundation. THE WITNESS: Yes. BY MR. EDWARDS: Q. As of the date of this article, the QBIC functionality described in this article was in IBM's Ultimedia Manager, IBM's -- (Reporter seeks clarification.) Q. Described in this article was in IBM's Ultimedia Manager, IBM's Digital library and DB2 series of products; correct? MR. HANSEN: Objection; lacks foundation. BY MR. EDWARDS: Q. Is that right, Mr. Flickner? A. Yes. MR. HANSEN: Same objection. BY MR. EDWARDS: Q. Okay. Let's pick up Exhibit 10, please. Look at page IBM 2419. It's the same -- I want to ask you about the same sentence I asked you before, where it lists the commercial products in -- at the bottom of the second full paragraph. A. Yeah. Q. See that? It says, "QBIC was developed by IBM and moved into commercial products including a 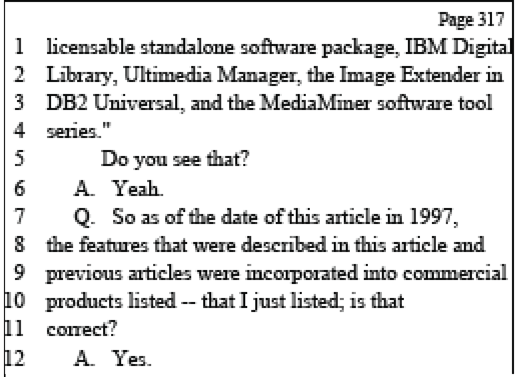 Page 317 licensable standalone software package, IBM Digital Library, Ultimedia Manager, the Image Extender in DB2 Universal, and the MediaMiner software tool series." Do you see that? A. Yeah. Q. So as of the date of this article in 1997, the features that were described in this article and previous articles were incorporated into commercial products listed -- that I just listed; is that correct? A. Yes.

Element	'036 Claim Recitation	Exemplary Citations in Reference
		See also Flickner Depo at 39-40, 54-55, 76, 88-89, 159; Flickner 8/29/23 Depo at 62–63, 74, 166–174, 180–181, 189, 193, 202, 316–317 To the extent it is found that QBIC does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that QBIC does not anticipate this claim, QBIC renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine QBIC with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1B.i]	communicates with a mobile device capable of acquiring a digital representation of a scene containing at least a portion of a target;	QBIC discloses an identification platform that communicates with a mobile device capable of acquiring a digital representation of a scene containing at least a portion of a target. 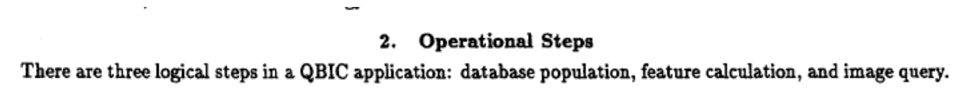 **2. Operational Steps** There are three logical steps in a QBIC application: database population, feature calculation, and image query. IBM0002390 at 2391:

Element	'036 Claim Recitation	Exemplary Citations in Reference
		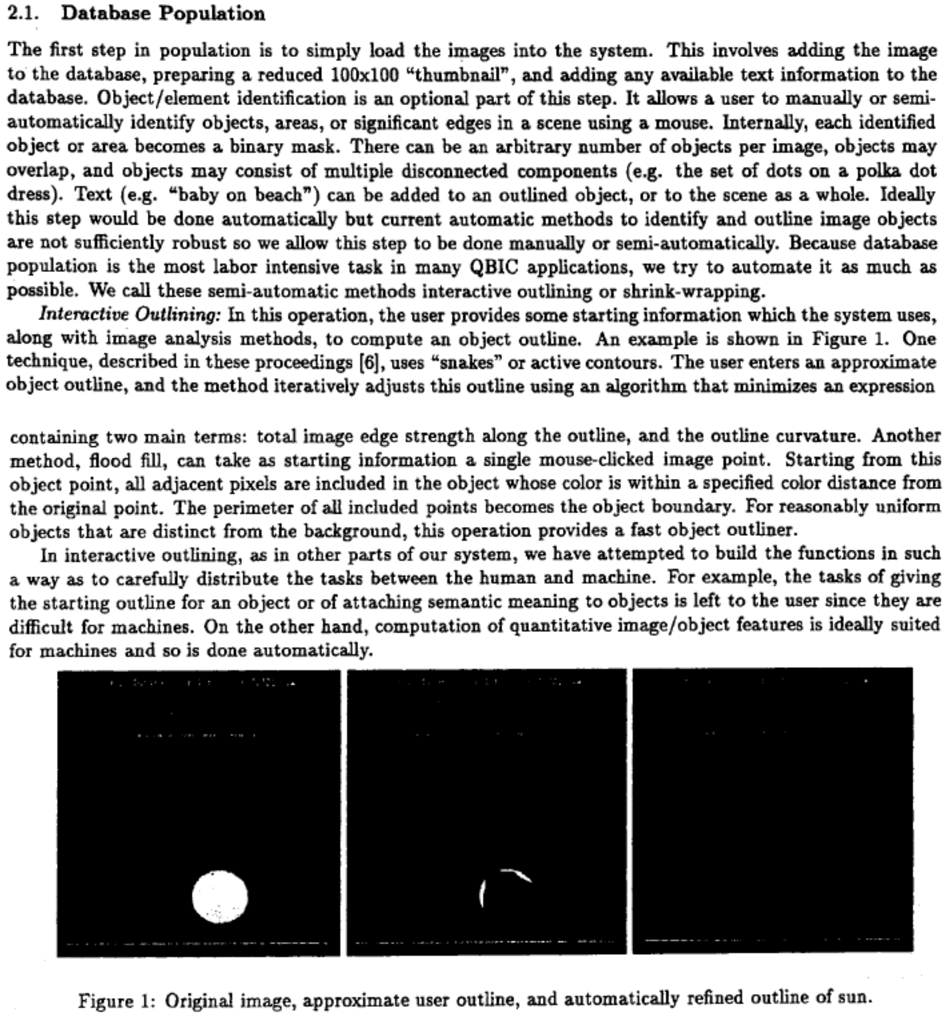 2.1. **Database Population** The first step in population is to simply load the images into the system. This involves adding the image to the database, preparing a reduced 100x100 "thumbnail", and adding any available text information to the database. Object/element identification is an optional part of this step. It allows a user to manually or semi-automatically identify objects, areas, or significant edges in a scene using a mouse. Internally, each identified object or area becomes a binary mask. There can be an arbitrary number of objects per image, objects may overlap, and objects may consist of multiple disconnected components (e.g. the set of dots on a polka dot dress). Text (e.g. "baby on beach") can be added to an outlined object, or to the scene as a whole. Ideally this step would be done automatically but current automatic methods to identify and outline image objects are not sufficiently robust so we allow this step to be done manually or semi-automatically. Because database population is the most labor intensive task in many QBIC applications, we try to automate it as much as possible. We call these semi-automatic methods interactive outlining or shrink-wrapping. *Interactive Outlining:* In this operation, the user provides some starting information which the system uses, along with image analysis methods, to compute an object outline. An example is shown in Figure 1. One technique, described in these proceedings [6], uses "snakes" or active contours. The user enters an approximate object outline, and the method iteratively adjusts this outline using an algorithm that minimizes an expression containing two main terms: total image edge strength along the outline, and the outline curvature. Another method, flood fill, can take as starting information a single mouse-clicked image point. Starting from this object point, all adjacent pixels are included in the object whose color is within a specified color distance from the original point. The perimeter of all included points becomes the object boundary. For reasonably uniform objects that are distinct from the background, this operation provides a fast object outliner. In interactive outlining, as in other parts of our system, we have attempted to build the functions in such a way as to carefully distribute the tasks between the human and machine. For example, the tasks of giving the starting outline for an object or of attaching semantic meaning to objects is left to the user since they are difficult for machines. On the other hand, computation of quantitative image/object features is ideally suited for machines and so is done automatically. Figure 1: Original image, approximate user outline, and automatically refined outline of sun. IBM0002390 at 2391-92:

Element	’036 Claim Recitation	Exemplary Citations in Reference
		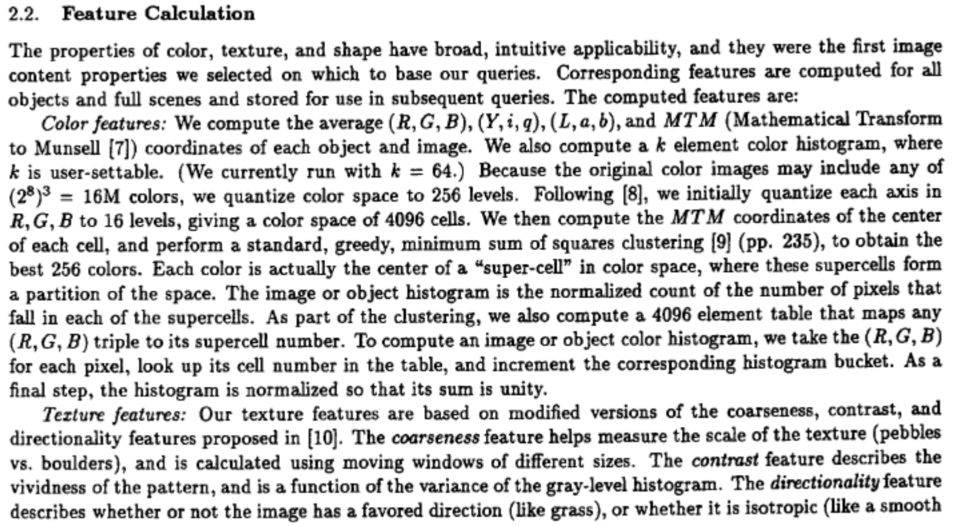 2.2. **Feature Calculation** The properties of color, texture, and shape have broad, intuitive applicability, and they were the first image content properties we selected on which to base our queries. Corresponding features are computed for all objects and full scenes and stored for use in subsequent queries. The computed features are: *Color features:* We compute the average (R, G, B), (Y, i, q), (L, a, b), and MTM (Mathematical Transform to Munsell [7]) coordinates of each object and image. We also compute a k element color histogram, where k is user-settable. (We currently run with $k = 64$.) Because the original color images may include any of $(2^8)^3 = 16$M colors, we quantize color space to 256 levels. Following [8], we initially quantize each axis in R, G, B to 16 levels, giving a color space of 4096 cells. We then compute the MTM coordinates of the center of each cell, and perform a standard, greedy, minimum sum of squares clustering [9] (pp. 235), to obtain the best 256 colors. Each color is actually the center of a “super-cell” in color space, where these supercells form a partition of the space. The image or object histogram is the normalized count of the number of pixels that fall in each of the supercells. As part of the clustering, we also compute a 4096 element table that maps any (R, G, B) triple to its supercell number. To compute an image or object color histogram, we take the (R, G, B) for each pixel, look up its cell number in the table, and increment the corresponding histogram bucket. As a final step, the histogram is normalized so that its sum is unity. *Texture features:* Our texture features are based on modified versions of the coarseness, contrast, and directionality features proposed in [10]. The *coarseness* feature helps measure the scale of the texture (pebbles vs. boulders), and is calculated using moving windows of different sizes. The *contrast* feature describes the vividness of the pattern, and is a function of the variance of the gray-level histogram. The *directionality* feature describes whether or not the image has a favored direction (like grass), or whether it is isotropic (like a smooth

Element	'036 Claim Recitation	Exemplary Citations in Reference
		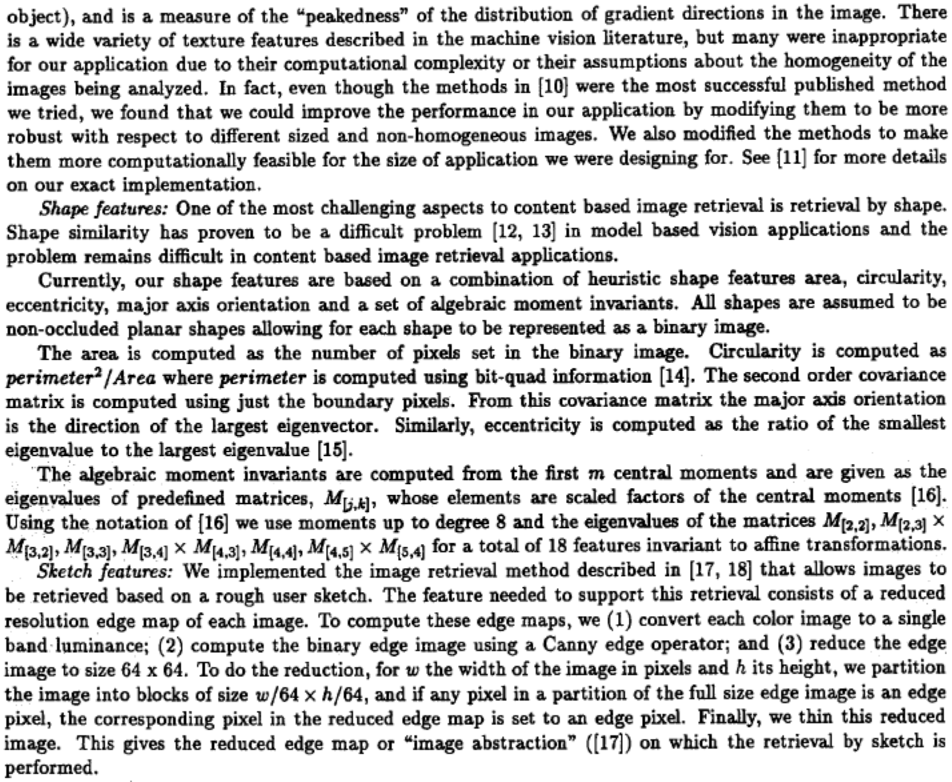 object), and is a measure of the "peakedness" of the distribution of gradient directions in the image. There is a wide variety of texture features described in the machine vision literature, but many were inappropriate for our application due to their computational complexity or their assumptions about the homogeneity of the images being analyzed. In fact, even though the methods in [10] were the most successful published method we tried, we found that we could improve the performance in our application by modifying them to be more robust with respect to different sized and non-homogeneous images. We also modified the methods to make them more computationally feasible for the size of application we were designing for. See [11] for more details on our exact implementation. *Shape features:* One of the most challenging aspects to content based image retrieval is retrieval by shape. Shape similarity has proven to be a difficult problem [12, 13] in model based vision applications and the problem remains difficult in content based image retrieval applications. Currently, our shape features are based on a combination of heuristic shape features area, circularity, eccentricity, major axis orientation and a set of algebraic moment invariants. All shapes are assumed to be non-occluded planar shapes allowing for each shape to be represented as a binary image. The area is computed as the number of pixels set in the binary image. Circularity is computed as $perimeter^2/Area$ where $perimeter$ is computed using bit-quad information [14]. The second order covariance matrix is computed using just the boundary pixels. From this covariance matrix the major axis orientation is the direction of the largest eigenvector. Similarly, eccentricity is computed as the ratio of the smallest eigenvalue to the largest eigenvalue [15]. The algebraic moment invariants are computed from the first m central moments and are given as the eigenvalues of predefined matrices, $M_{[j,k]}$, whose elements are scaled factors of the central moments [16]. Using the notation of [16] we use moments up to degree 8 and the eigenvalues of the matrices $M_{[2,2]}, M_{[2,3]} \times M_{[3,2]}, M_{[3,3]}, M_{[3,4]} \times M_{[4,3]}, M_{[4,4]}, M_{[4,5]} \times M_{[5,4]}$ for a total of 18 features invariant to affine transformations. *Sketch features:* We implemented the image retrieval method described in [17, 18] that allows images to be retrieved based on a rough user sketch. The feature needed to support this retrieval consists of a reduced resolution edge map of each image. To compute these edge maps, we (1) convert each color image to a single band luminance; (2) compute the binary edge image using a Canny edge operator; and (3) reduce the edge image to size 64 x 64. To do the reduction, for w the width of the image in pixels and h its height, we partition the image into blocks of size $w/64 \times h/64$, and if any pixel in a partition of the full size edge image is an edge pixel, the corresponding pixel in the reduced edge map is set to an edge pixel. Finally, we thin this reduced image. This gives the reduced edge map or "image abstraction" ([17]) on which the retrieval by sketch is performed. IBM0002390 at 2392-93:

Element	'036 Claim Recitation	Exemplary Citations in Reference
		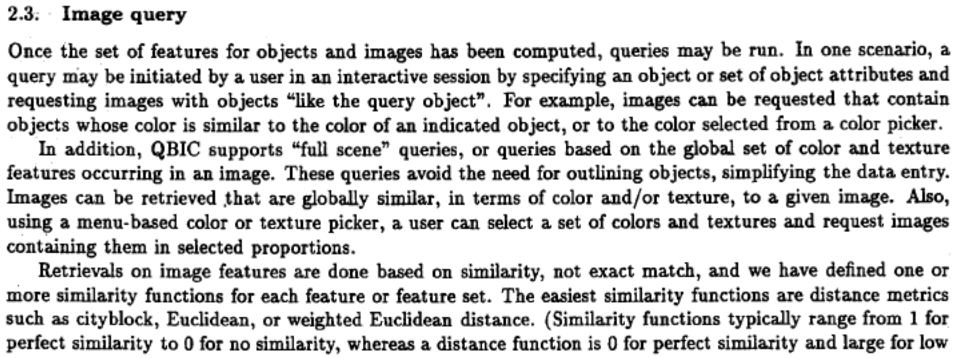 **2.3. Image query** Once the set of features for objects and images has been computed, queries may be run. In one scenario, a query may be initiated by a user in an interactive session by specifying an object or set of object attributes and requesting images with objects "like the query object". For example, images can be requested that contain objects whose color is similar to the color of an indicated object, or to the color selected from a color picker. In addition, QBIC supports "full scene" queries, or queries based on the global set of color and texture features occurring in an image. These queries avoid the need for outlining objects, simplifying the data entry. Images can be retrieved that are globally similar, in terms of color and/or texture, to a given image. Also, using a menu-based color or texture picker, a user can select a set of colors and textures and request images containing them in selected proportions. Retrievals on image features are done based on similarity, not exact match, and we have defined one or more similarity functions for each feature or feature set. The easiest similarity functions are distance metrics such as cityblock, Euclidean, or weighted Euclidean distance. (Similarity functions typically range from 1 for perfect similarity to 0 for no similarity, whereas a distance function is 0 for perfect similarity and large for low

Element	'036 Claim Recitation	Exemplary Citations in Reference
		similarity, so we must convert as necessary to obtain consistent measures.) Different similarity measures may be used, and they may vary from application to application. The matching methods and similarity functions for the different feature types are described below: *Color:* For average color, the distance between a query object and database object is weighted Euclidean distance. The weights are user adjustable so that we may, for example, downweight the L component of an (L, a, b) query. Best results, as judged by our subjective evaluation in on-going experiments, are obtained by representing the average color in the MTM color space, and by inversely weighting each component by its standard deviation over the samples in the database. Improved results are obtained when color matching is done not on average color, but on the distribution of colors that occur in an object or image. In this way, a red and blue object matches other red and blue objects better than to a purple object. To match color histograms, we use a method similar to the one described in [8]. Let X be the query histogram and Y the histogram of an item in the database, both normalized. We compute the element by element difference histogram Z. Then the similarity between X and Y is given by $\|Z\| = Z^T A Z$, where A is a symmetric color similarity matrix with $$a(i,j) \equiv 1 - d(c_i, c_j)/d_{max}$$ where c_i and c_j are the ith and jth colors in the color histograms, and $d(c_i, c_j)$ is the MTM color distance, and d_{max} is the maximum distance between any two colors ([8] used Luv distance.) This metric gives the (weighted) length of the difference vector between X and Y, weighted by A, which accounts for the perceptual distance between different pairs of colors. The result is a measure that accounts for both the difference between the amounts in any given color (e.g. a particular shade of red), as well as differences between similar colors (e.g. red and orange). Based on our experience, we also found it useful for the user to be able to request images with $x\%$ of color 1, $y\%$ of color 2, etc. where both the percentages and the colors are explicitly specified by the user. This allows searches such as for a beach scene (25% white (for the sand) and 50% blue (for the sky and water)), and for many scenes for which the user has seen before, and for which he/she can remember approximate color percentages. The color picker that allows this type of query is shown in Figure 2. 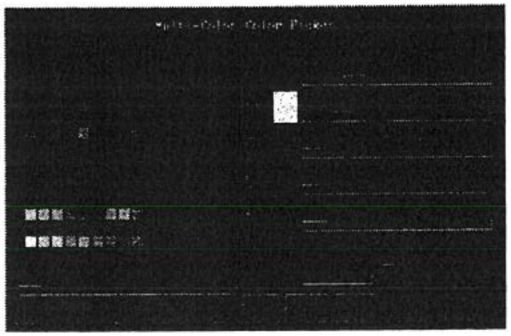 Figure 2: Multi-Color Color Picker with two selected colors. The palette on the left displays the set of selectable colors. Once a color is selected, it appears in a box on the right, and the amount of the color can be specified by the sliders. In this case, 25% light blue and 15% white have been selected. *Texture:* Texture distance is computed as weighted Euclidean distance in the three dimensional texture space. The most common weighting factors are the inverse variances for each component, computed over the

Element	'036 Claim Recitation	Exemplary Citations in Reference
		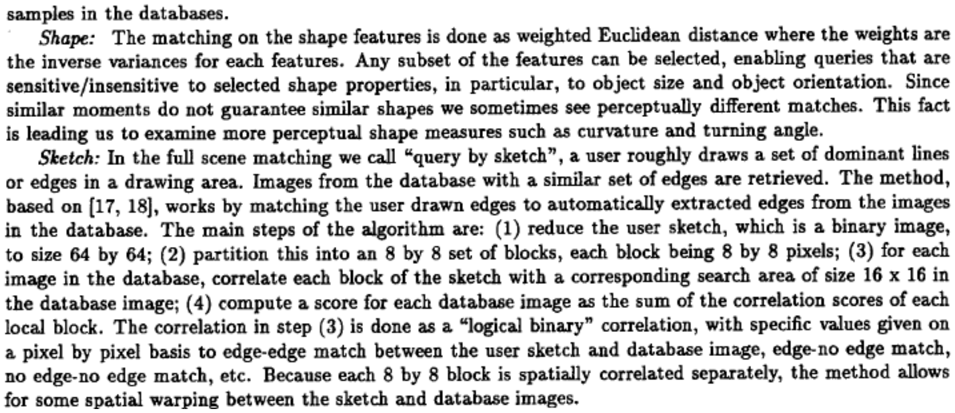 samples in the databases. *Shape:* The matching on the shape features is done as weighted Euclidean distance where the weights are the inverse variances for each features. Any subset of the features can be selected, enabling queries that are sensitive/insensitive to selected shape properties, in particular, to object size and object orientation. Since similar moments do not guarantee similar shapes we sometimes see perceptually different matches. This fact is leading us to examine more perceptual shape measures such as curvature and turning angle. *Sketch:* In the full scene matching we call "query by sketch", a user roughly draws a set of dominant lines or edges in a drawing area. Images from the database with a similar set of edges are retrieved. The method, based on [17, 18], works by matching the user drawn edges to automatically extracted edges from the images in the database. The main steps of the algorithm are: (1) reduce the user sketch, which is a binary image, to size 64 by 64; (2) partition this into an 8 by 8 set of blocks, each block being 8 by 8 pixels; (3) for each image in the database, correlate each block of the sketch with a corresponding search area of size 16 x 16 in the database image; (4) compute a score for each database image as the sum of the correlation scores of each local block. The correlation in step (3) is done as a "logical binary" correlation, with specific values given on a pixel by pixel basis to edge-edge match between the user sketch and database image, edge-no edge match, no edge-no edge match, etc. Because each 8 by 8 block is spatially correlated separately, the method allows for some spatial warping between the sketch and database images. IBM0002390 at 2393-95: Query by Image and Video Content: The QBIC System (IBM 0000893–902): "To achieve this functionality, QBIC has two main components: database population (the process of creating an image database) and database query. During the population, images and videos are processed to extract features describing their content-colors, textures, shapes, and camera and object motion-and the features are stored in a database. . . . For both population and query, the QBIC data model has • still images or scenes (full images) that contain objects (subsets of an image), and • video shots that consist of sets of contiguous frames and contain motion objects." *Id*. at p. 3. Image, Audio, and Video Extenders Administration and Programming (IBM 000001–617): "To make an image available for searching by content, you catalog the image. When you catalog an image, the Image Extender analyzes the image, by computing the feature values for the image, and stores the values in a QBIC catalog. When you search for an image by content, your query identifies one or

Element	'036 Claim Recitation	Exemplary Citations in Reference
		more features for the search (such as average color), a source for each feature (such as an example image), and a target set of cataloged images." *Id*. at pp. 20–21. QBIC could integrate with a camera, such as a portable camera; a user could take an image using a portable camera and submit that image to the QBIC system for query. Flickner Depo at, e.g., 24, 41-42, 62-63 (user could take a picture or video and submit it to the QBIC system).

Element	'036 Claim Recitation	Exemplary Citations in Reference
		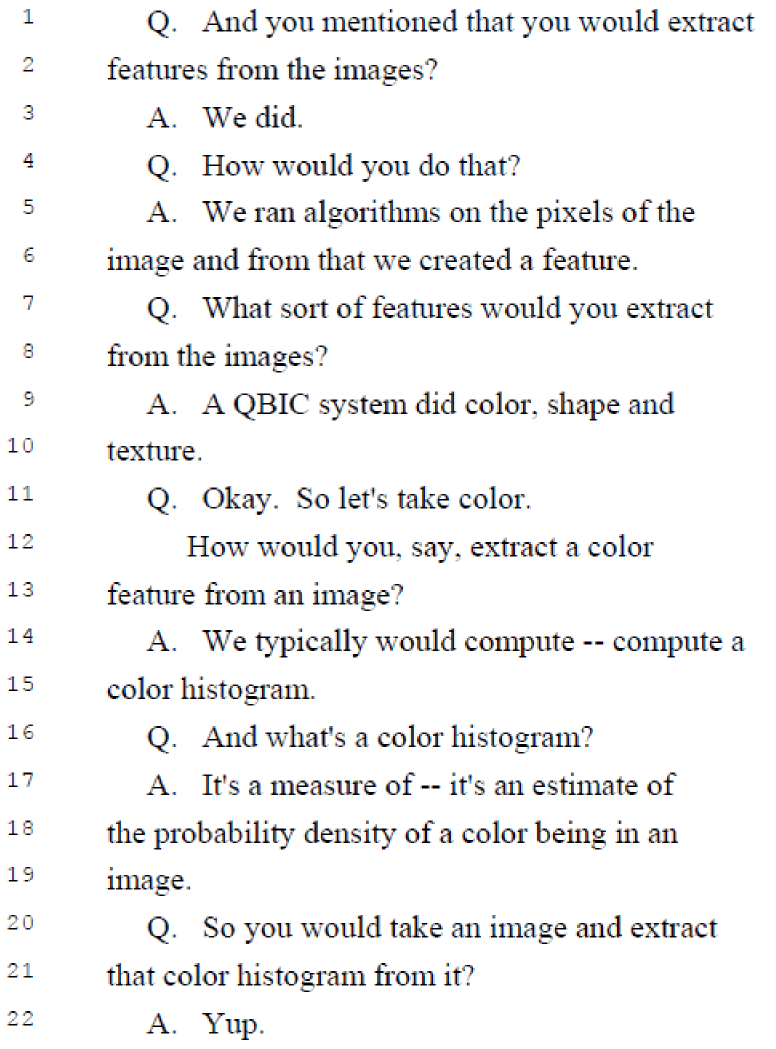 Q. And you mentioned that you would extract features from the images? A. We did. Q. How would you do that? A. We ran algorithms on the pixels of the image and from that we created a feature. Q. What sort of features would you extract from the images? A. A QBIC system did color, shape and texture. Q. Okay. So let's take color. How would you, say, extract a color feature from an image? A. We typically would compute -- compute a color histogram. Q. And what's a color histogram? A. It's a measure of -- it's an estimate of the probability density of a color being in an image. Q. So you would take an image and extract that color histogram from it? A. Yup.

		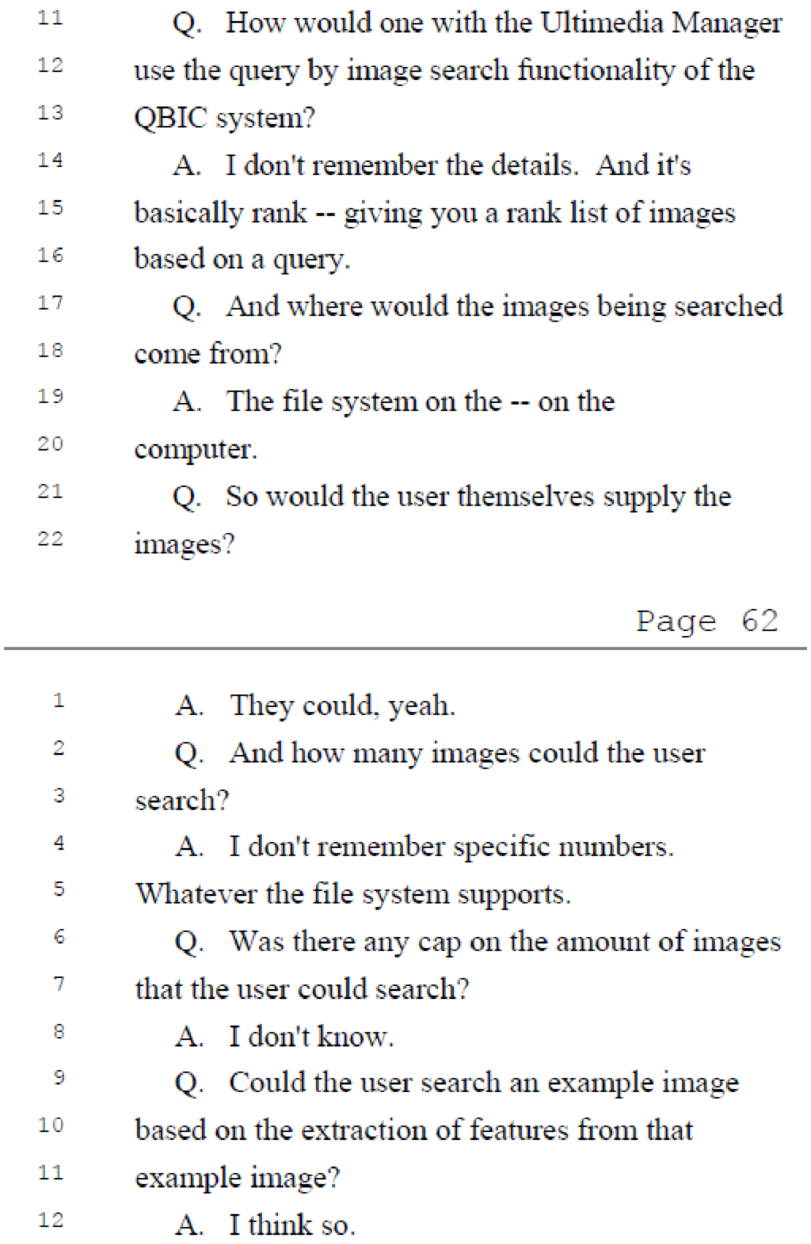 Q. How would one with the Ultimedia Manager use the query by image search functionality of the QBIC system? A. I don't remember the details. And it's basically rank -- giving you a rank list of images based on a query. Q. And where would the images being searched come from? A. The file system on the -- on the computer. Q. So would the user themselves supply the images? Page 62 A. They could, yeah. Q. And how many images could the user search? A. I don't remember specific numbers. Whatever the file system supports. Q. Was there any cap on the amount of images that the user could search? A. I don't know. Q. Could the user search an example image based on the extraction of features from that example image? A. I think so.

Element	'036 Claim Recitation	Exemplary Citations in Reference
		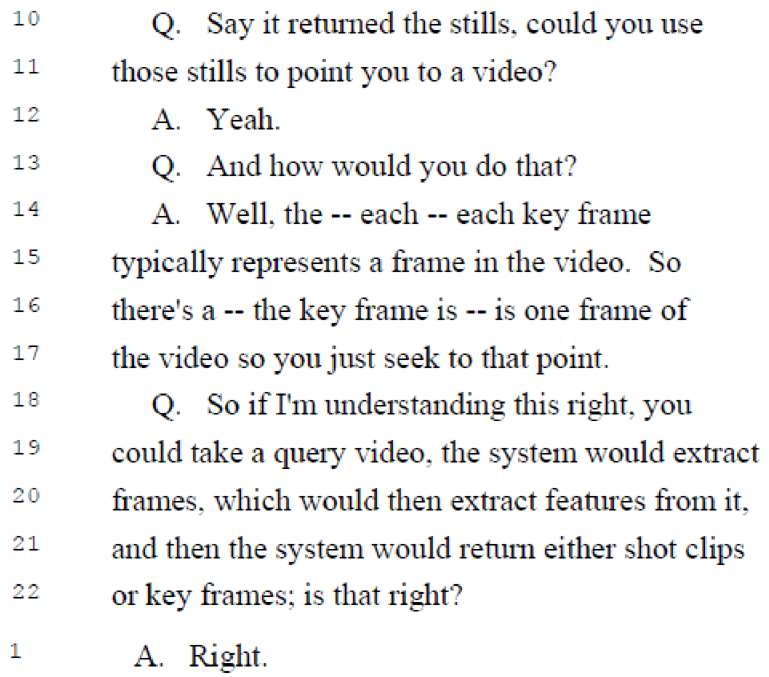 Q. Say it returned the stills, could you use those stills to point you to a video? A. Yeah. Q. And how would you do that? A. Well, the -- each -- each key frame typically represents a frame in the video. So there's a -- the key frame is -- is one frame of the video so you just seek to that point. Q. So if I'm understanding this right, you could take a query video, the system would extract frames, which would then extract features from it, and then the system would return either shot clips or key frames; is that right? A. Right. Flickner 8/29/23 Depo at 53:

Element	'036 Claim Recitation	Exemplary Citations in Reference
		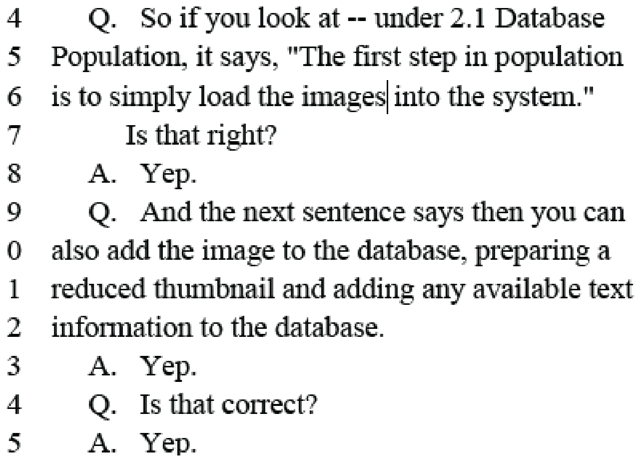 Q. So if you look at -- under 2.1 Database Population, it says, "The first step in population is to simply load the images into the system." Is that right? A. Yep. Q. And the next sentence says then you can also add the image to the database, preparing a reduced thumbnail and adding any available text information to the database. A. Yep. Q. Is that correct? A. Yep. *Id.* at 60–61:

Element	'036 Claim Recitation	Exemplary Citations in Reference
		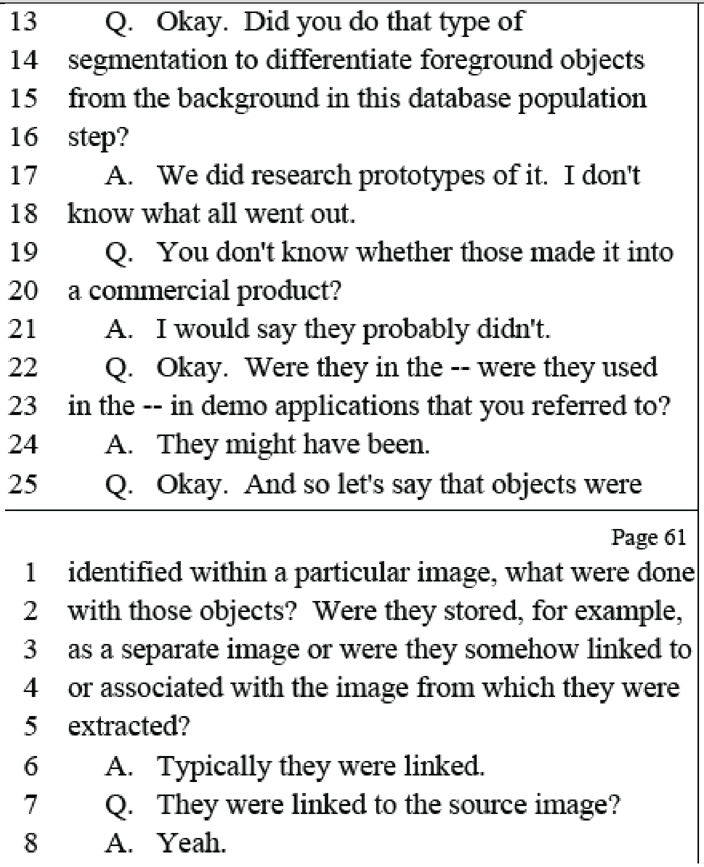 Q. Okay. Did you do that type of segmentation to differentiate foreground objects from the background in this database population step? A. We did research prototypes of it. I don't know what all went out. Q. You don't know whether those made it into a commercial product? A. I would say they probably didn't. Q. Okay. Were they in the -- were they used in the -- in demo applications that you referred to? A. They might have been. Q. Okay. And so let's say that objects were Page 61 identified within a particular image, what were done with those objects? Were they stored, for example, as a separate image or were they somehow linked to or associated with the image from which they were extracted? A. Typically they were linked. Q. They were linked to the source image? A. Yeah. *Id.* at 32–33:

Element	'036 Claim Recitation	Exemplary Citations in Reference
		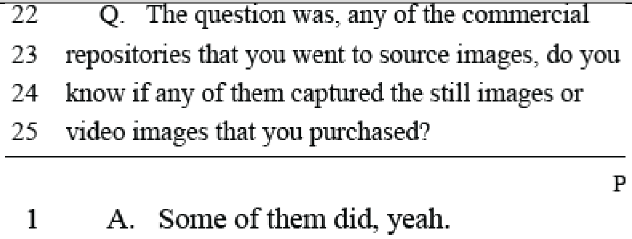 Q. The question was, any of the commercial repositories that you went to source images, do you know if any of them captured the still images or video images that you purchased? P A. Some of them did, yeah. *Id.* at 74–75:

			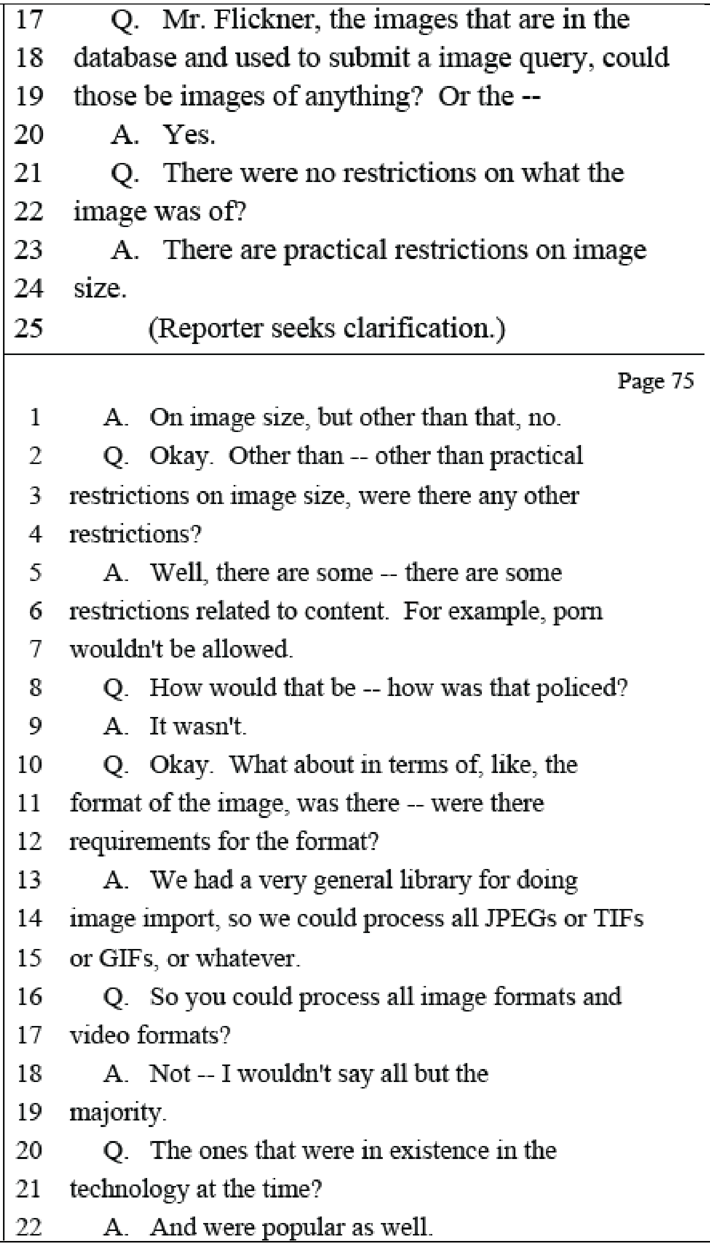 Q. Mr. Flickner, the images that are in the database and used to submit a image query, could those be images of anything? Or the -- A. Yes. Q. There were no restrictions on what the image was of? A. There are practical restrictions on image size. (Reporter seeks clarification.) Page 75 A. On image size, but other than that, no. Q. Okay. Other than -- other than practical restrictions on image size, were there any other restrictions? A. Well, there are some -- there are some restrictions related to content. For example, porn wouldn't be allowed. Q. How would that be -- how was that policed? A. It wasn't. Q. Okay. What about in terms of, like, the format of the image, was there -- were there requirements for the format? A. We had a very general library for doing image import, so we could process all JPEGs or TIFs or GIFs, or whatever. Q. So you could process all image formats and video formats? A. Not -- I wouldn't say all but the majority. Q. The ones that were in existence in the technology at the time? A. And were popular as well.

Element	'036 Claim Recitation	Exemplary Citations in Reference
		Id. at 114 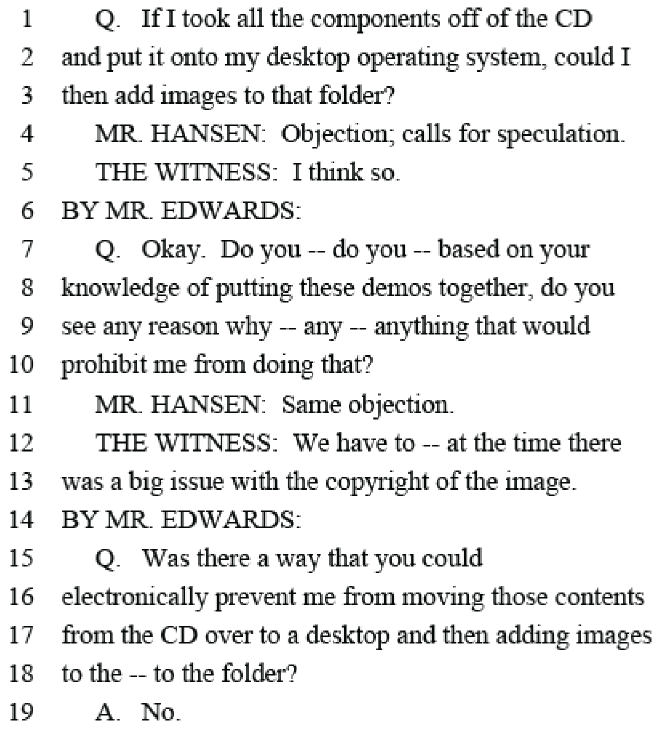 Q. If I took all the components off of the CD and put it onto my desktop operating system, could I then add images to that folder? MR. HANSEN: Objection; calls for speculation. THE WITNESS: I think so. BY MR. EDWARDS: Q. Okay. Do you -- do you -- based on your knowledge of putting these demos together, do you see any reason why -- any -- anything that would prohibit me from doing that? MR. HANSEN: Same objection. THE WITNESS: We have to -- at the time there was a big issue with the copyright of the image. BY MR. EDWARDS: Q. Was there a way that you could electronically prevent me from moving those contents from the CD over to a desktop and then adding images to the -- to the folder? A. No. *Id.* at 202:

Element	'036 Claim Recitation	Exemplary Citations in Reference
		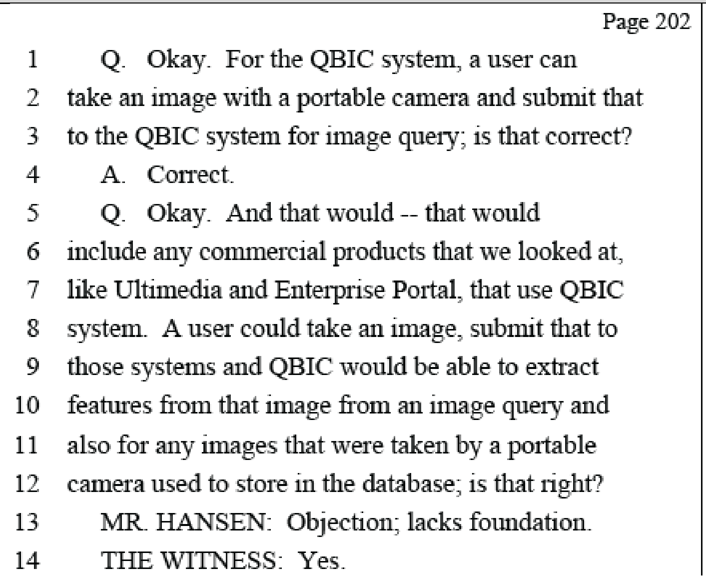 Page 202 Q. Okay. For the QBIC system, a user can take an image with a portable camera and submit that to the QBIC system for image query; is that correct? A. Correct. Q. Okay. And that would -- that would include any commercial products that we looked at, like Ultimedia and Enterprise Portal, that use QBIC system. A user could take an image, submit that to those systems and QBIC would be able to extract features from that image from an image query and also for any images that were taken by a portable camera used to store in the database; is that right? MR. HANSEN: Objection; lacks foundation. THE WITNESS: Yes. *Id.* at 300: 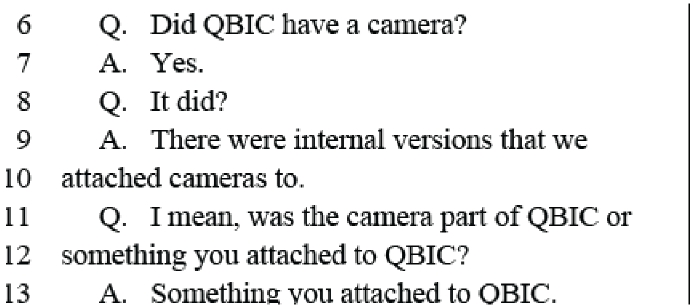 Q. Did QBIC have a camera? A. Yes. Q. It did? A. There were internal versions that we attached cameras to. Q. I mean, was the camera part of QBIC or something you attached to QBIC? A. Something you attached to QBIC. *Id.* at 312:

Element	'036 Claim Recitation	Exemplary Citations in Reference
		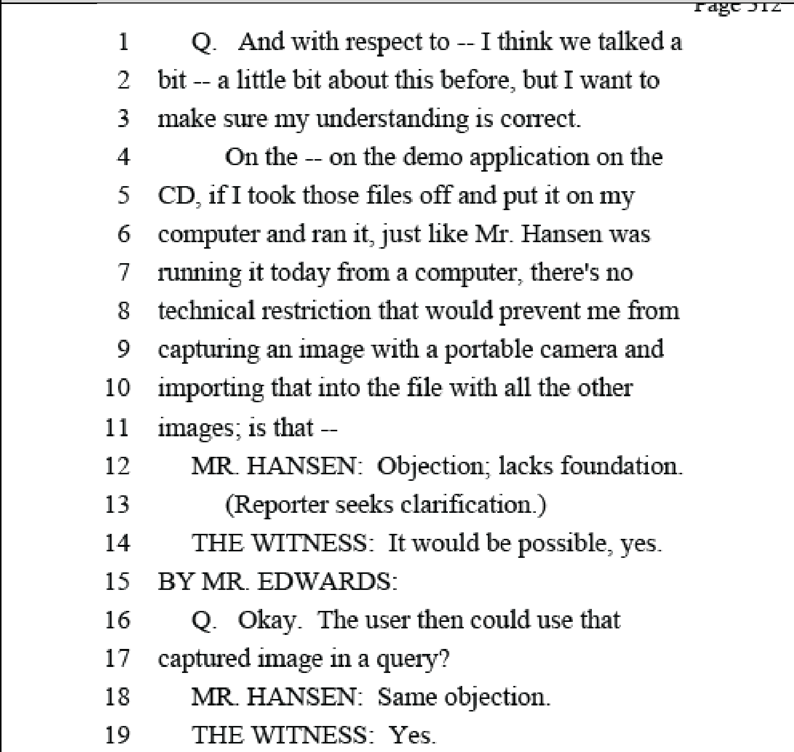 Page 312 Q. And with respect to -- I think we talked a bit -- a little bit about this before, but I want to make sure my understanding is correct. On the -- on the demo application on the CD, if I took those files off and put it on my computer and ran it, just like Mr. Hansen was running it today from a computer, there's no technical restriction that would prevent me from capturing an image with a portable camera and importing that into the file with all the other images; is that -- MR. HANSEN: Objection; lacks foundation. (Reporter seeks clarification.) THE WITNESS: It would be possible, yes. BY MR. EDWARDS: Q. Okay. The user then could use that captured image in a query? MR. HANSEN: Same objection. THE WITNESS: Yes. *Id.* at 305: 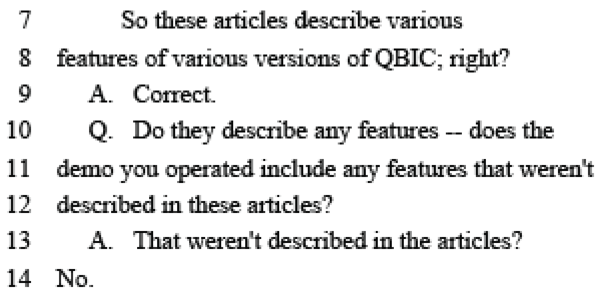 So these articles describe various features of various versions of QBIC; right? A. Correct. Q. Do they describe any features -- does the demo you operated include any features that weren't described in these articles? A. That weren't described in the articles? No.

Element	'036 Claim Recitation	Exemplary Citations in Reference
		Id. at 35: 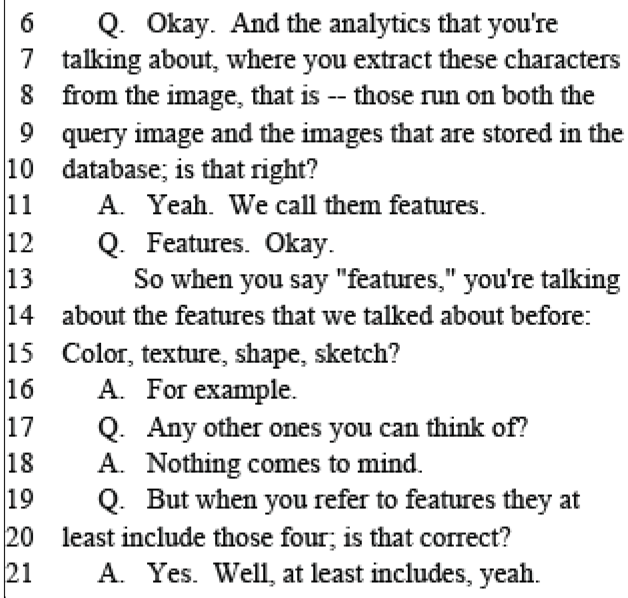 Q. Okay. And the analytics that you're talking about, where you extract these characters from the image, that is -- those run on both the query image and the images that are stored in the database; is that right? A. Yeah. We call them features. Q. Features. Okay. So when you say "features," you're talking about the features that we talked about before: Color, texture, shape, sketch? A. For example. Q. Any other ones you can think of? A. Nothing comes to mind. Q. But when you refer to features they at least include those four; is that correct? A. Yes. Well, at least includes, yeah. *Id.* at 109-110:

Element	'036 Claim Recitation	Exemplary Citations in Reference
		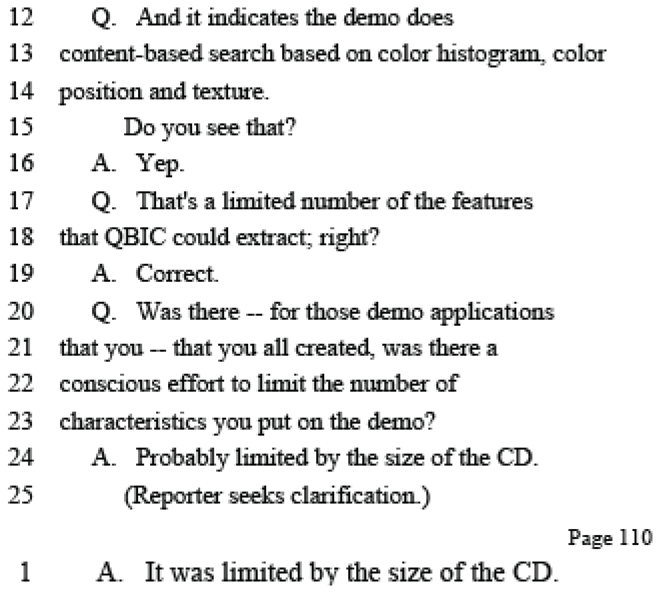 Q. And it indicates the demo does content-based search based on color histogram, color position and texture. Do you see that? A. Yep. Q. That's a limited number of the features that QBIC could extract; right? A. Correct. Q. Was there -- for those demo applications that you -- that you all created, was there a conscious effort to limit the number of characteristics you put on the demo? A. Probably limited by the size of the CD. (Reporter seeks clarification.) Page 110 A. It was limited by the size of the CD. *Id.* at 311: 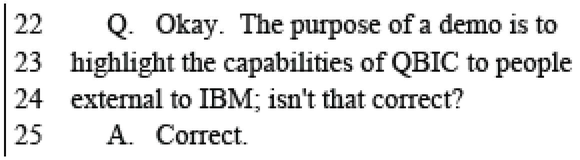 Q. Okay. The purpose of a demo is to highlight the capabilities of QBIC to people external to IBM; isn't that correct? A. Correct. *Id.* at 313-317:

Element	'036 Claim Recitation	Exemplary Citations in Reference
		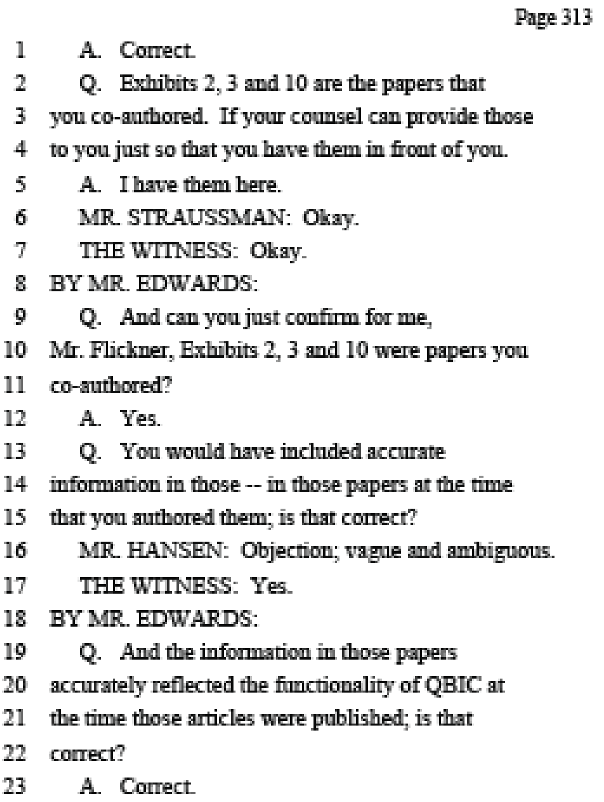 Page 313 A. Correct. Q. Exhibits 2, 3 and 10 are the papers that you co-authored. If your counsel can provide those to you just so that you have them in front of you. A. I have them here. MR. STRAUSSMAN: Okay. THE WITNESS: Okay. BY MR. EDWARDS: Q. And can you just confirm for me, Mr. Flickner, Exhibits 2, 3 and 10 were papers you co-authored? A. Yes. Q. You would have included accurate information in those -- in those papers at the time that you authored them; is that correct? MR. HANSEN: Objection; vague and ambiguous. THE WITNESS: Yes. BY MR. EDWARDS: Q. And the information in those papers accurately reflected the functionality of QBIC at the time those articles were published; is that correct? A. Correct.

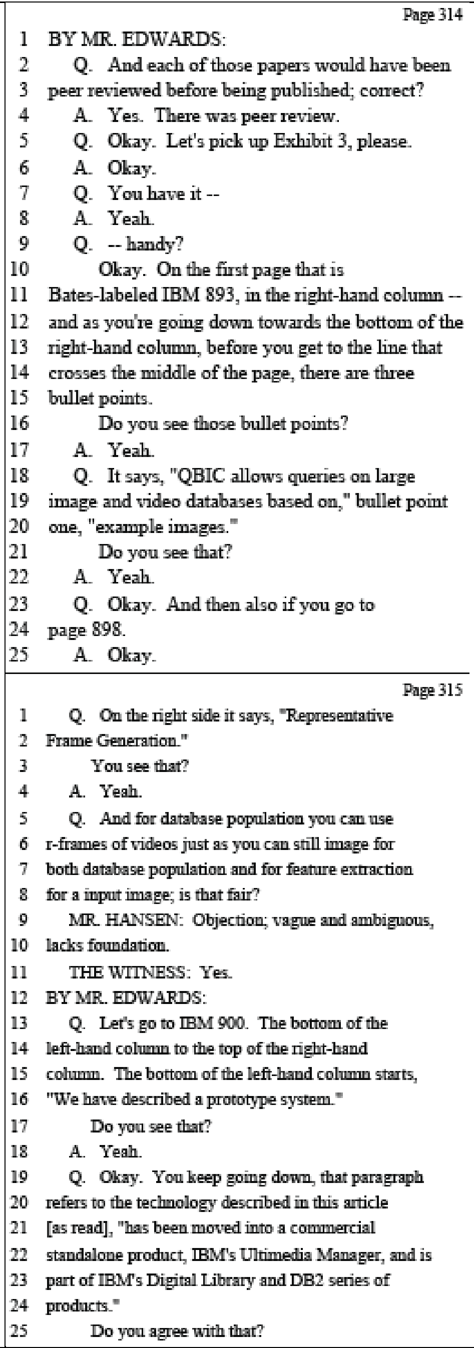

Page 314

BY MR. EDWARDS:

Q. And each of those papers would have been peer reviewed before being published; correct?

A. Yes. There was peer review.

Q. Okay. Let's pick up Exhibit 3, please.

A. Okay.

Q. You have it --

A. Yeah.

Q. -- handy?

Okay. On the first page that is Bates-labeled IBM 893, in the right-hand column -- and as you're going down towards the bottom of the right-hand column, before you get to the line that crosses the middle of the page, there are three bullet points.

Do you see those bullet points?

A. Yeah.

Q. It says, "QBIC allows queries on large image and video databases based on," bullet point one, "example images."

Do you see that?

A. Yeah.

Q. Okay. And then also if you go to page 898.

A. Okay.

Page 315

Q. On the right side it says, "Representative Frame Generation."

You see that?

A. Yeah.

Q. And for database population you can use r-frames of videos just as you can still image for both database population and for feature extraction for a input image; is that fair?

MR. HANSEN: Objection; vague and ambiguous, lacks foundation.

THE WITNESS: Yes.

BY MR. EDWARDS:

Q. Let's go to IBM 900. The bottom of the left-hand column to the top of the right-hand column. The bottom of the left-hand column starts, "We have described a prototype system."

Do you see that?

A. Yeah.

Q. Okay. You keep going down, that paragraph refers to the technology described in this article [as read], "has been moved into a commercial standalone product, IBM's Ultimedia Manager, and is part of IBM's Digital Library and DB2 series of products."

Do you agree with that?

Element	'036 Claim Recitation	Exemplary Citations in Reference
		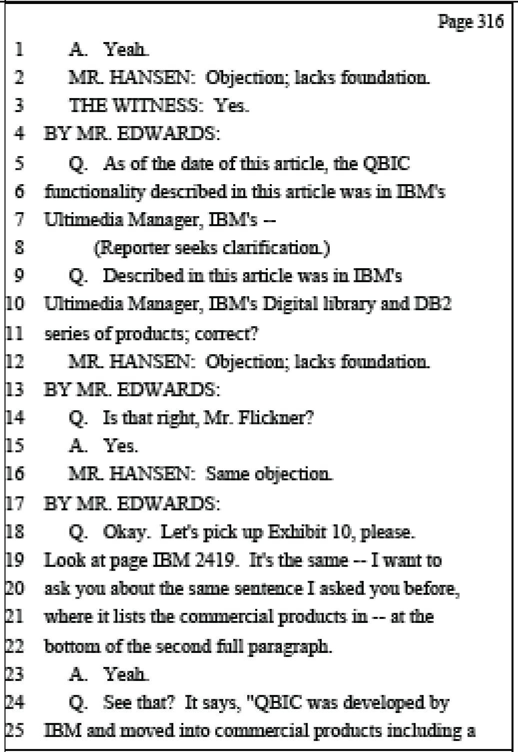 Page 316 A. Yeah. MR. HANSEN: Objection; lacks foundation. THE WITNESS: Yes. BY MR. EDWARDS: Q. As of the date of this article, the QBIC functionality described in this article was in IBM's Ultimedia Manager, IBM's -- (Reporter seeks clarification.) Q. Described in this article was in IBM's Ultimedia Manager, IBM's Digital library and DB2 series of products; correct? MR. HANSEN: Objection; lacks foundation. BY MR. EDWARDS: Q. Is that right, Mr. Flickner? A. Yes. MR. HANSEN: Same objection. BY MR. EDWARDS: Q. Okay. Let's pick up Exhibit 10, please. Look at page IBM 2419. It's the same -- I want to ask you about the same sentence I asked you before, where it lists the commercial products in -- at the bottom of the second full paragraph. A. Yeah. Q. See that? It says, "QBIC was developed by IBM and moved into commercial products including a 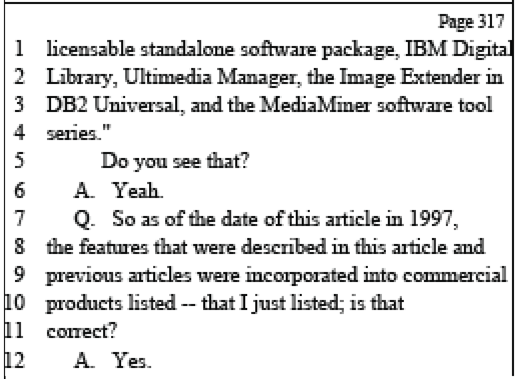 Page 317 licensable standalone software package, IBM Digital Library, Ultimedia Manager, the Image Extender in DB2 Universal, and the MediaMiner software tool series." Do you see that? A. Yeah. Q. So as of the date of this article in 1997, the features that were described in this article and previous articles were incorporated into commercial products listed -- that I just listed; is that correct? A. Yes.

Element	'036 Claim Recitation	Exemplary Citations in Reference
		QBIC Demo (IBM000747) – showing compatibility with images captured by a mobile phone camera 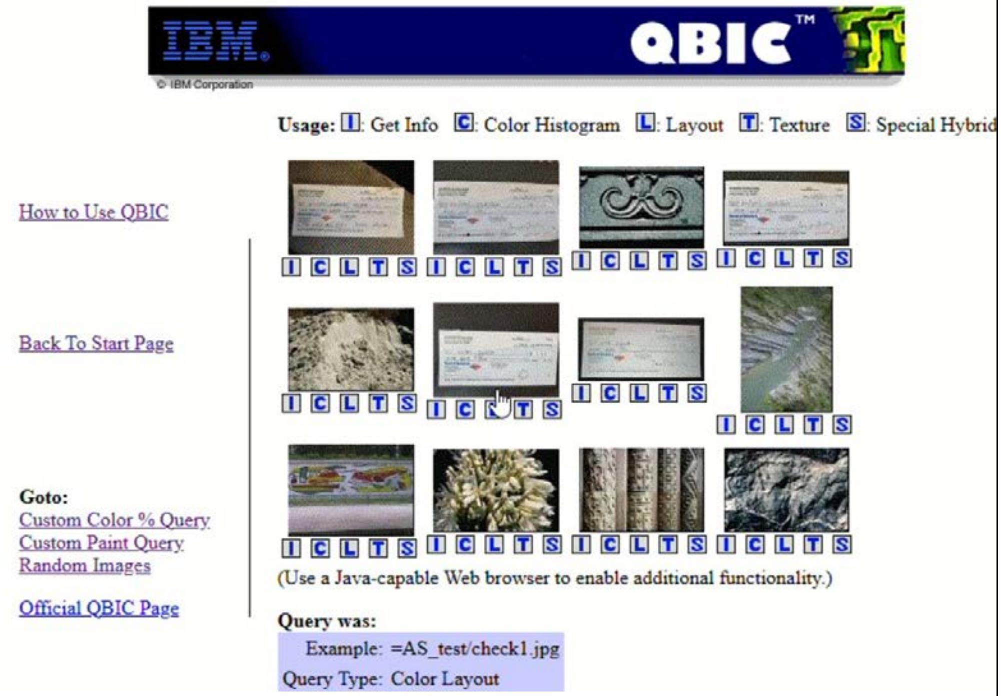 (Use a Java-capable Web browser to enable additional functionality.) Query was: Example: =AS_test/check1.jpg Query Type: Color Layout

Element	'036 Claim Recitation	Exemplary Citations in Reference
		

Element	'036 Claim Recitation	Exemplary Citations in Reference
		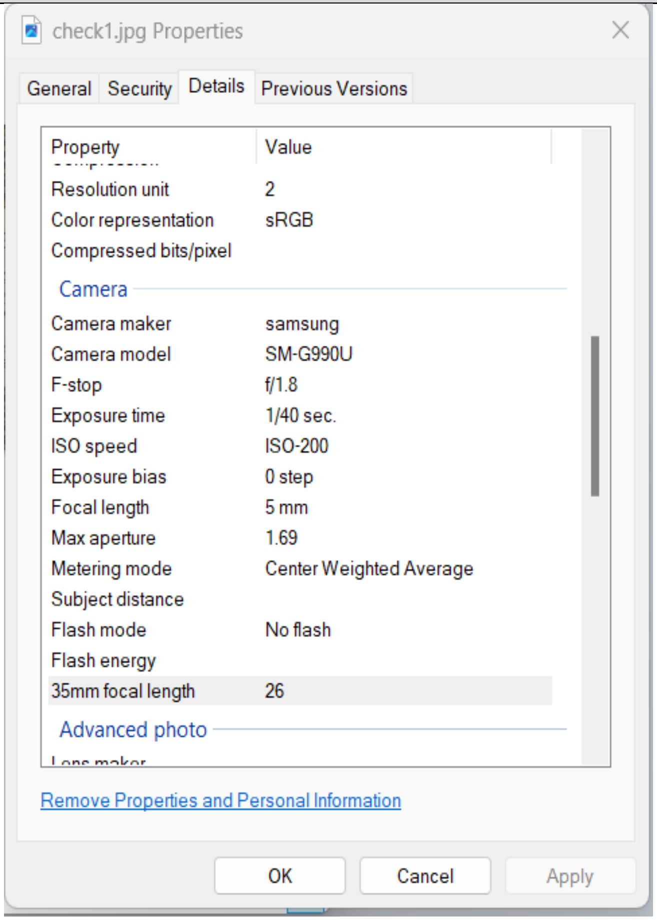

check1.jpg Properties

Property	Value
Resolution unit	2
Color representation	sRGB
Compressed bits/pixel	
Camera	
Camera maker	samsung
Camera model	SM-G990U
F-stop	f/1.8
Exposure time	1/40 sec.
ISO speed	ISO-200
Exposure bias	0 step
Focal length	5 mm
Max aperture	1.69
Metering mode	Center Weighted Average
Subject distance	
Flash mode	No flash
Flash energy	
35mm focal length	26
Advanced photo	

Element	'036 Claim Recitation	Exemplary Citations in Reference
		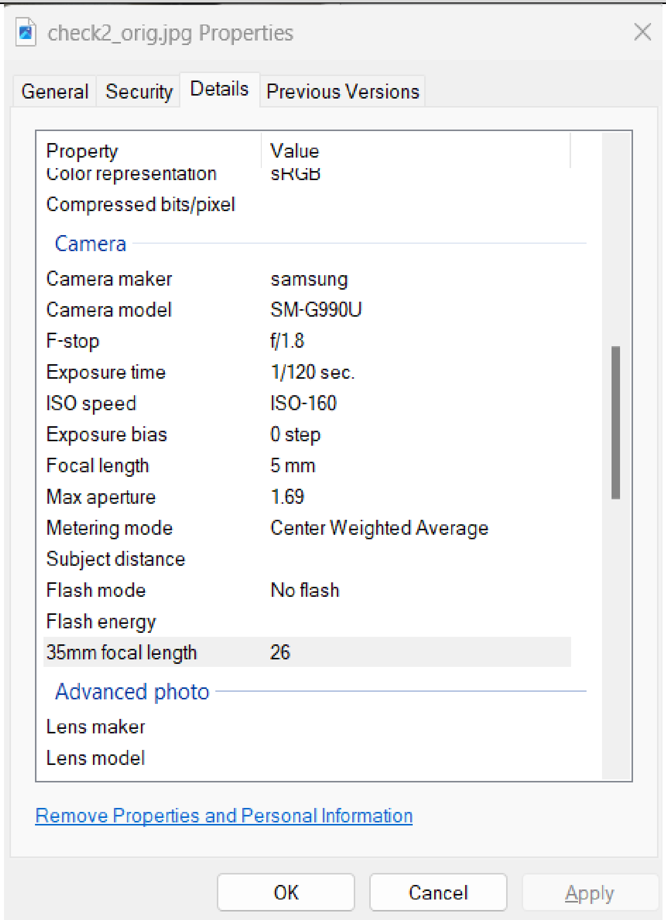 Property: Value Color representation: sRGB Compressed bits/pixel: Camera Camera maker: samsung Camera model: SM-G990U F-stop: f/1.8 Exposure time: 1/120 sec. ISO speed: ISO-160 Exposure bias: 0 step Focal length: 5 mm Max aperture: 1.69 Metering mode: Center Weighted Average Subject distance: Flash mode: No flash Flash energy: 35mm focal length: 26 Advanced photo Lens maker: Lens model:

Element	'036 Claim Recitation	Exemplary Citations in Reference
		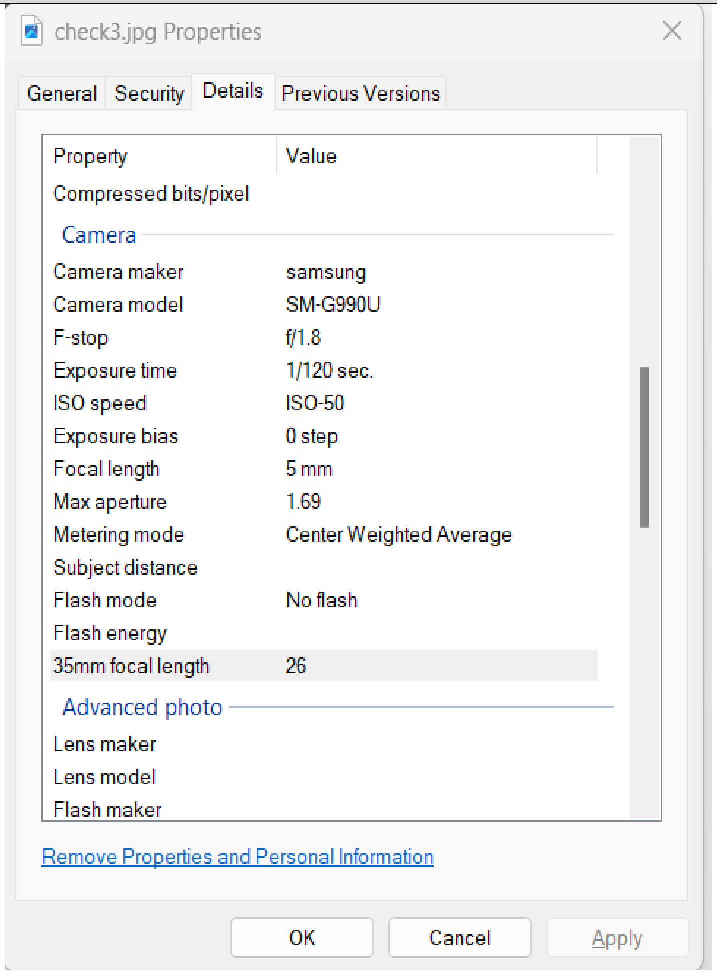 check3.jpg Properties Property Value Compressed bits/pixel Camera Camera maker samsung Camera model SM-G990U F-stop f/1.8 Exposure time 1/120 sec. ISO speed ISO-50 Exposure bias 0 step Focal length 5 mm Max aperture 1.69 Metering mode Center Weighted Average Subject distance Flash mode No flash Flash energy 35mm focal length 26 Advanced photo Lens maker Lens model Flash maker Remove Properties and Personal Information

Element	'036 Claim Recitation	Exemplary Citations in Reference
		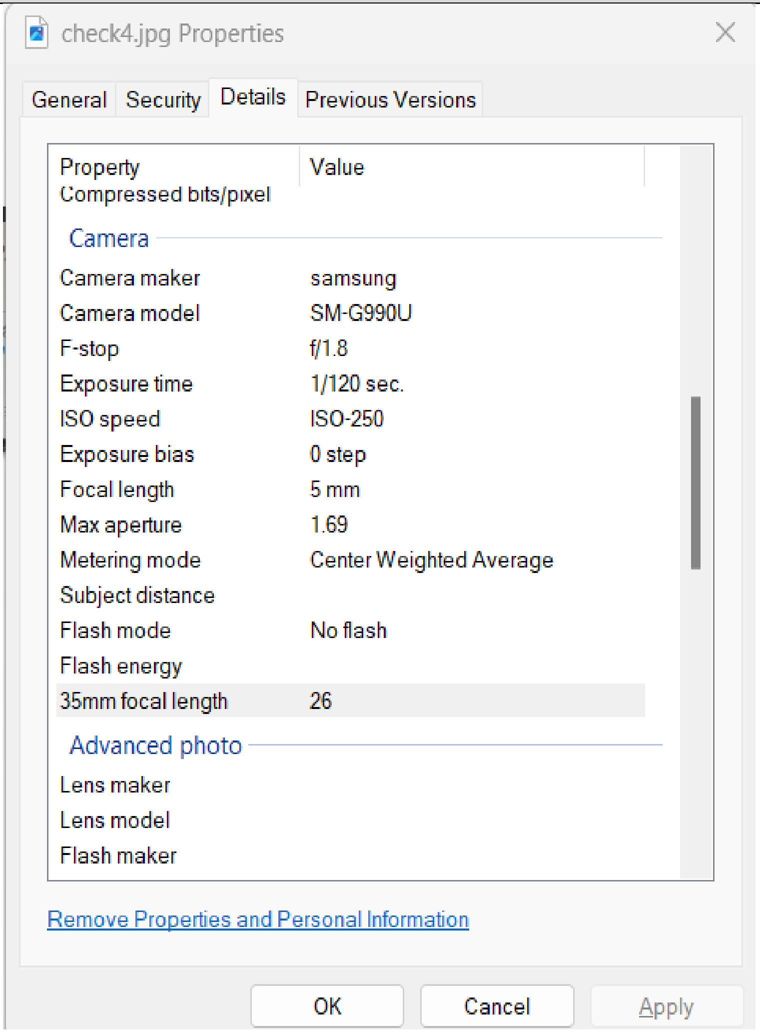 check4.jpg Properties General Security Details Previous Versions Property: Value Compressed bits/pixel: Camera Camera maker: samsung Camera model: SM-G990U F-stop: f/1.8 Exposure time: 1/120 sec. ISO speed: ISO-250 Exposure bias: 0 step Focal length: 5 mm Max aperture: 1.69 Metering mode: Center Weighted Average Subject distance: Flash mode: No flash Flash energy: 35mm focal length: 26 Advanced photo Lens maker: Lens model: Flash maker: Remove Properties and Personal Information OK Cancel Apply

Element	'036 Claim Recitation	Exemplary Citations in Reference
		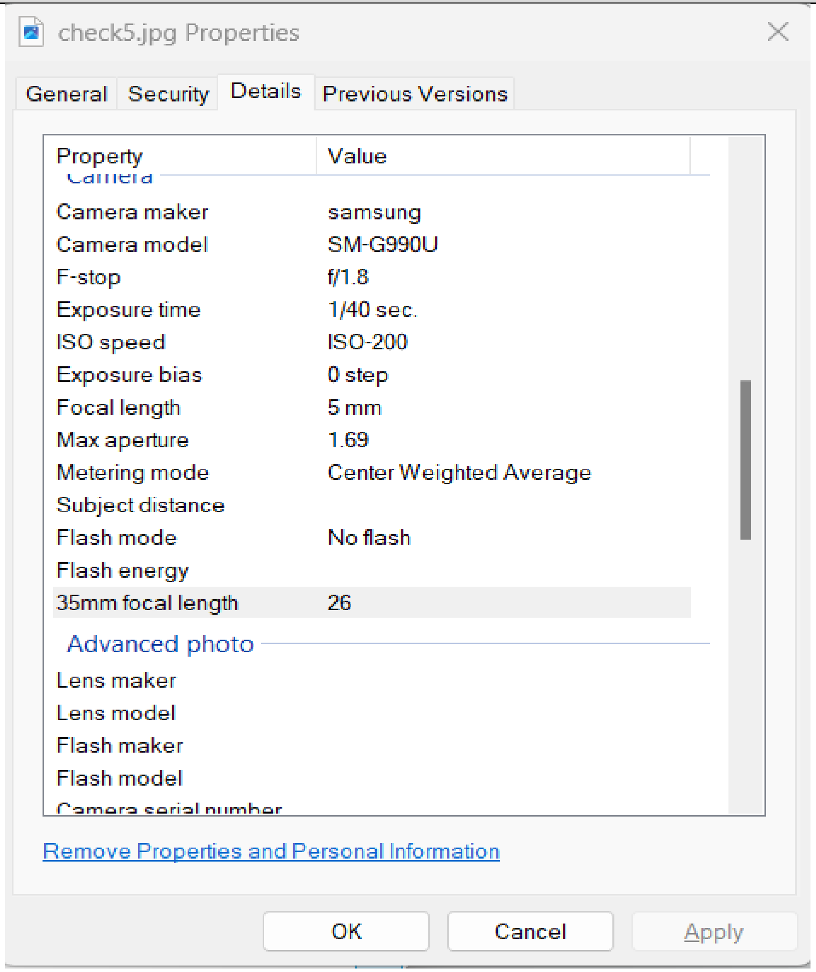 check5.jpg Properties Property Value Camera Camera maker samsung Camera model SM-G990U F-stop f/1.8 Exposure time 1/40 sec. ISO speed ISO-200 Exposure bias 0 step Focal length 5 mm Max aperture 1.69 Metering mode Center Weighted Average Subject distance Flash mode No flash Flash energy 35mm focal length 26 Advanced photo Lens maker Lens model Flash maker Flash model Camera serial number To the extent it is found that QBIC does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that QBIC does not anticipate this claim, QBIC renders it obvious, either alone, in combination with the

Element	'036 Claim Recitation	Exemplary Citations in Reference
		knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine QBIC with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1B.ii]	receives the digital representation from the mobile device; and	QBIC discloses an identification platform that receives the digital representation from the mobile device. Query by Image and Video Content: The QBIC System (IBM 0000893–902): "To achieve this functionality, QBIC has two main components: database population (the process of creating an image database) and database query. During the population, images and videos are processed to extract features describing their content-colors, textures, shapes, and camera and object motion-and the features are stored in a database. During the query, the user composes a query graphically. Features are generated from the graphical query and then input to a matching engine that finds images or videos from the database with similar features. Figure 2 shows the system architecture. . . . For both population and query, the QBIC data model has • still images or scenes (full images) that contain objects (subsets of an image), and • video shots that consist of sets of contiguous frames and contain motion objects." *Id*. at p. 3. Image, Audio, and Video Extenders Administration and Programming (IBM 000001–617): "To make an image available for searching by content, you catalog the image. When you catalog an image, the Image Extender analyzes the image, by computing the feature values for the image, and stores the values in a QBIC catalog. When you search for an image by content, your query identifies one or

Element	'036 Claim Recitation	Exemplary Citations in Reference
		more features for the search (such as average color), a source for each feature (such as an example image), and a target set of cataloged images." *Id*. at pp. 20–21. QBIC Demo (IBM000747): 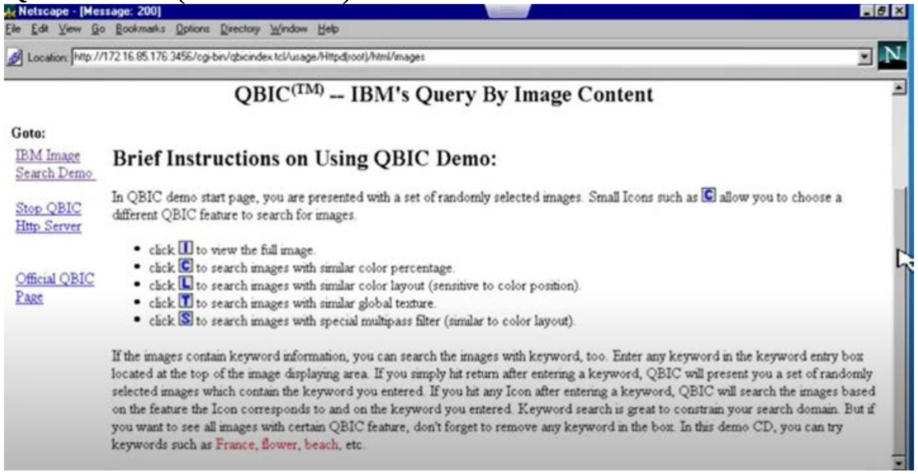 QBIC$^{(TM)}$ -- IBM's Query By Image Content **Brief Instructions on Using QBIC Demo:** In QBIC demo start page, you are presented with a set of randomly selected images. Small Icons such as C allow you to choose a different QBIC feature to search for images. - click I to view the full image. - click C to search images with similar color percentage. - click L to search images with similar color layout (sensitive to color position). - click T to search images with similar global texture. - click S to search images with special multipass filter (similar to color layout). If the images contain keyword information, you can search the images with keyword, too. Enter any keyword in the keyword entry box located at the top of the image displaying area. If you simply hit return after entering a keyword, QBIC will present you a set of randomly selected images which contain the keyword you entered. If you hit any Icon after entering a keyword, QBIC will search the images based on the feature the Icon corresponds to and on the keyword you entered. Keyword search is great to constrain your search domain. But if you want to see all images with certain QBIC feature, don't forget to remove any keyword in the box. In this demo CD, you can try keywords such as France, flower, beach, etc. QBIC Demo (IBM000747): 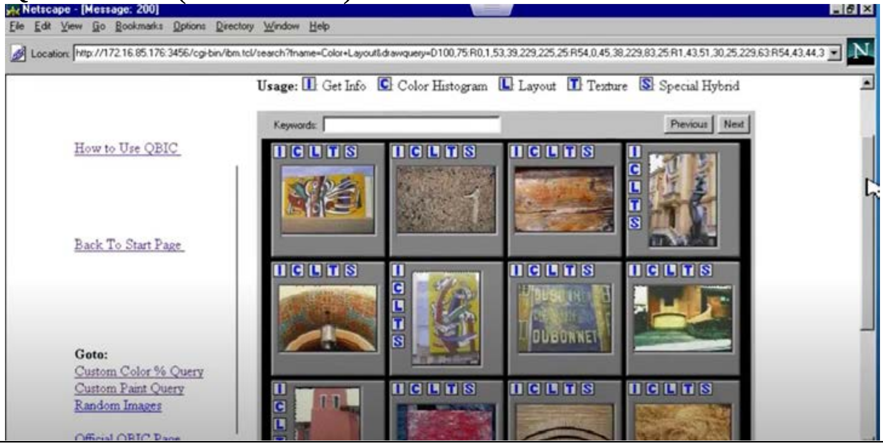

Element	'036 Claim Recitation	Exemplary Citations in Reference
		QBIC Demo (IBM000747): 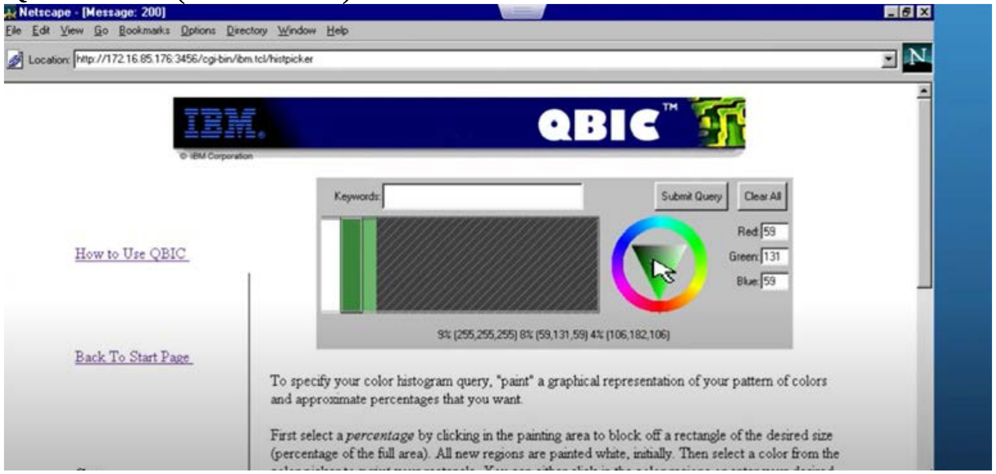 Location: http://172.16.85.176:3456/cgi-bin/ibm.tcl/histpicker Keywords: Red: 59 Green: 131 Blue: 59 9% (255,255,255) 8% (59,131,59) 4% (106,182,106) To specify your color histogram query, "paint" a graphical representation of your pattern of colors and approximate percentages that you want. First select a *percentage* by clicking in the painting area to block off a rectangle of the desired size (percentage of the full area). All new regions are painted white, initially. Then select a color from the QBIC Demo (IBM000747):

Element	'036 Claim Recitation	Exemplary Citations in Reference
		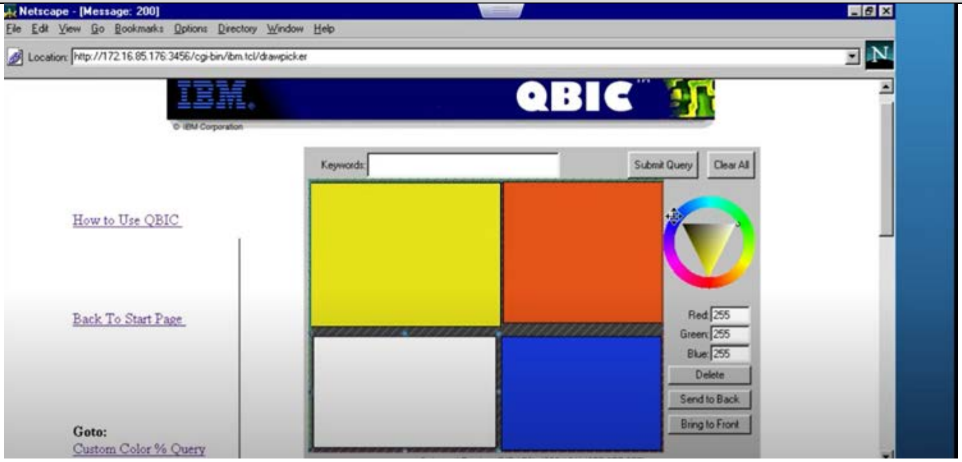 Netscape - [Message: 200] Location: http://172.16.85.176:3456/cgi-bin/ibm.tcl/drawpicker Keywords: Red: 255 Green: 255 Blue: 255 How to Use QBIC Back To Start Page Goto: Custom Color % Query QBIC Demo (IBM000747):  Netscape - [Message: 200] Location: http://172.16.85.176:3456/cgi-bin/ibm.tcl/search?fname=HybridColor&france10%2Ejpg.x=0 Goto: Custom Color % Query Custom Paint Query Random Images Official QBIC Page Query was: Example: =france10.jpg Query Type: QbHybridColorFeatureClass

Element	'036 Claim Recitation	Exemplary Citations in Reference
		Flickner 8/29/23 Depo at 53: 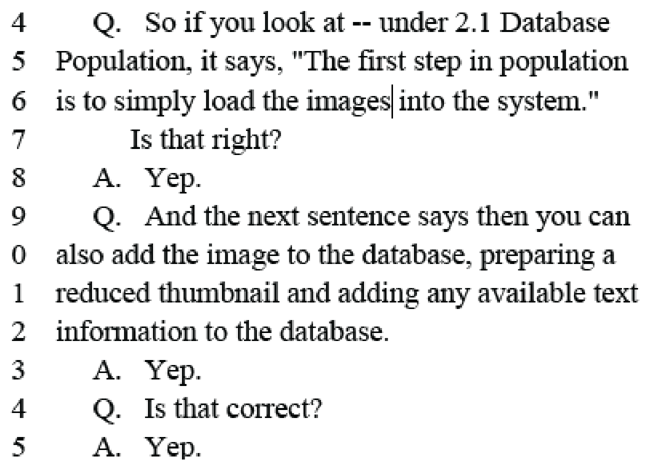 Q. So if you look at -- under 2.1 Database Population, it says, "The first step in population is to simply load the images into the system." Is that right? A. Yep. Q. And the next sentence says then you can also add the image to the database, preparing a reduced thumbnail and adding any available text information to the database. A. Yep. Q. Is that correct? A. Yep. Flickner 8/29/23 Depo at 60–61:

Element	’036 Claim Recitation	Exemplary Citations in Reference
		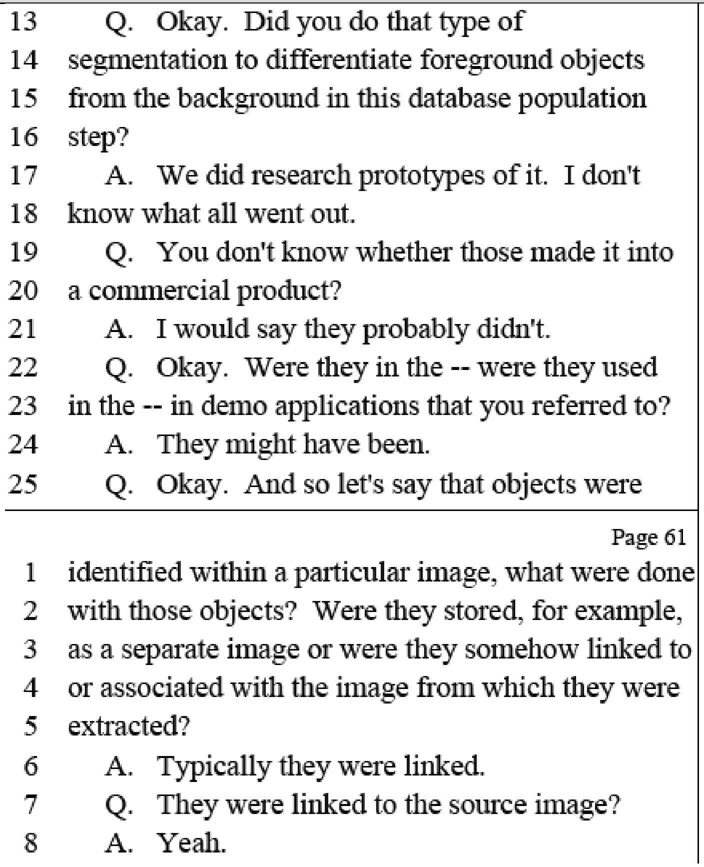 Q. Okay. Did you do that type of segmentation to differentiate foreground objects from the background in this database population step? A. We did research prototypes of it. I don't know what all went out. Q. You don't know whether those made it into a commercial product? A. I would say they probably didn't. Q. Okay. Were they in the -- were they used in the -- in demo applications that you referred to? A. They might have been. Q. Okay. And so let's say that objects were Page 61 identified within a particular image, what were done with those objects? Were they stored, for example, as a separate image or were they somehow linked to or associated with the image from which they were extracted? A. Typically they were linked. Q. They were linked to the source image? A. Yeah. Flickner 8/29/23 Depo at 32–33:

Element	'036 Claim Recitation	Exemplary Citations in Reference
		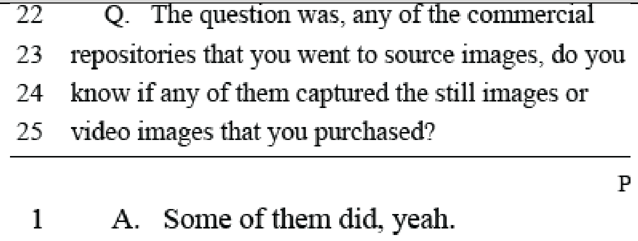 Q. The question was, any of the commercial repositories that you went to source images, do you know if any of them captured the still images or video images that you purchased? P A. Some of them did, yeah. Flickner 8/29/23 at 74–75:

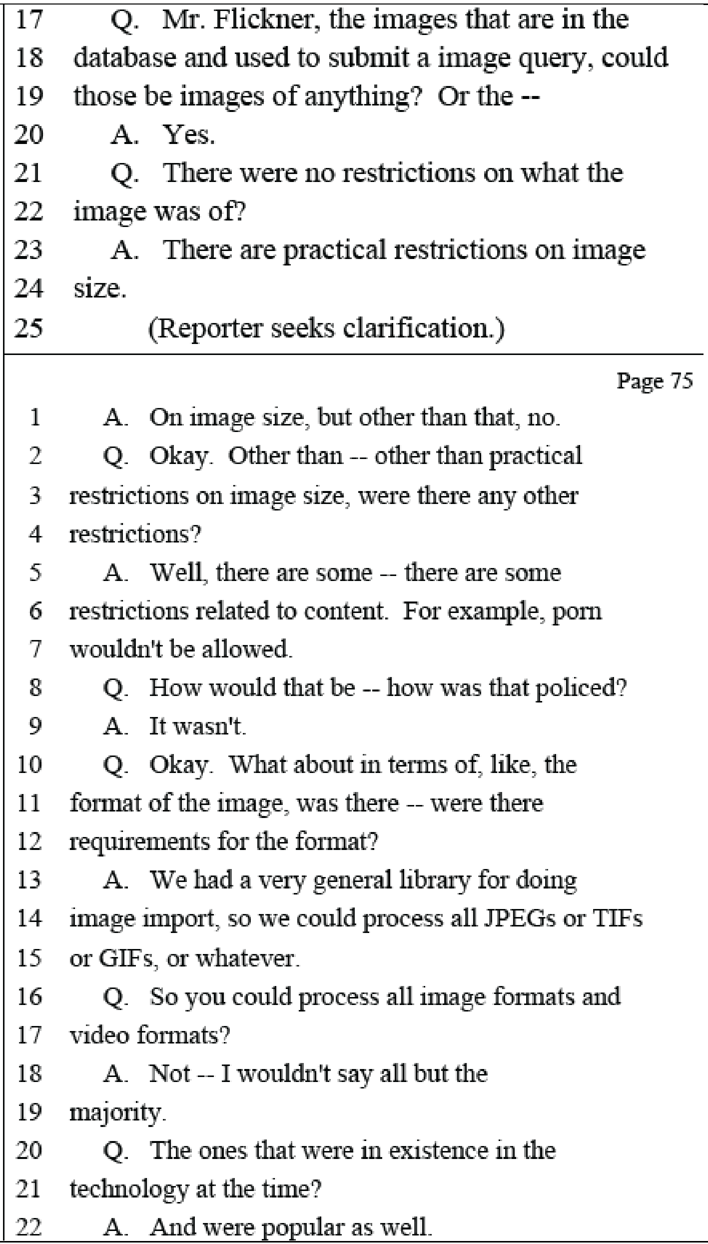

Q. Mr. Flickner, the images that are in the database and used to submit a image query, could those be images of anything? Or the --
A. Yes.
Q. There were no restrictions on what the image was of?
A. There are practical restrictions on image size.
(Reporter seeks clarification.)

Page 75

A. On image size, but other than that, no.
Q. Okay. Other than -- other than practical restrictions on image size, were there any other restrictions?
A. Well, there are some -- there are some restrictions related to content. For example, porn wouldn't be allowed.
Q. How would that be -- how was that policed?
A. It wasn't.
Q. Okay. What about in terms of, like, the format of the image, was there -- were there requirements for the format?
A. We had a very general library for doing image import, so we could process all JPEGs or TIFs or GIFs, or whatever.
Q. So you could process all image formats and video formats?
A. Not -- I wouldn't say all but the majority.
Q. The ones that were in existence in the technology at the time?
A. And were popular as well.

Element	'036 Claim Recitation	Exemplary Citations in Reference
		Id. at 114: 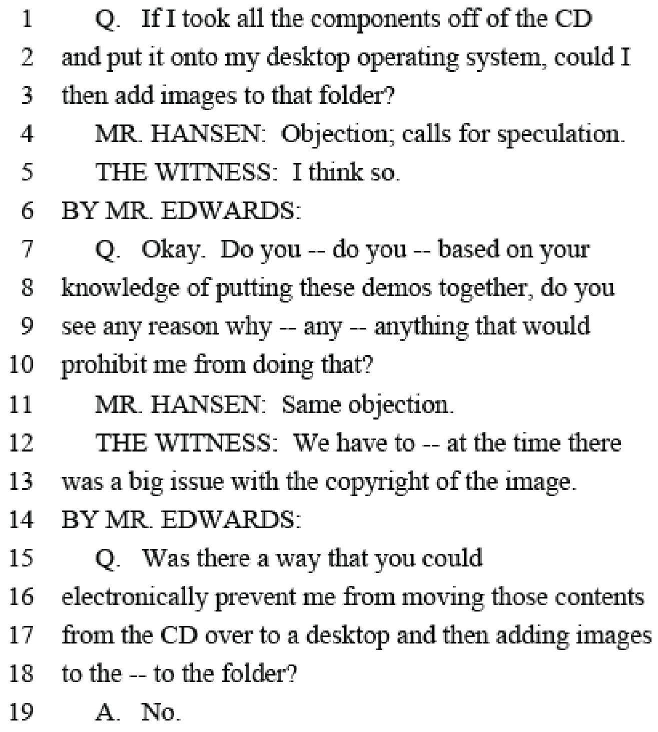 Q. If I took all the components off of the CD and put it onto my desktop operating system, could I then add images to that folder? MR. HANSEN: Objection; calls for speculation. THE WITNESS: I think so. BY MR. EDWARDS: Q. Okay. Do you -- do you -- based on your knowledge of putting these demos together, do you see any reason why -- any -- anything that would prohibit me from doing that? MR. HANSEN: Same objection. THE WITNESS: We have to -- at the time there was a big issue with the copyright of the image. BY MR. EDWARDS: Q. Was there a way that you could electronically prevent me from moving those contents from the CD over to a desktop and then adding images to the -- to the folder? A. No. *Id.* at 202:

Element	'036 Claim Recitation	Exemplary Citations in Reference
		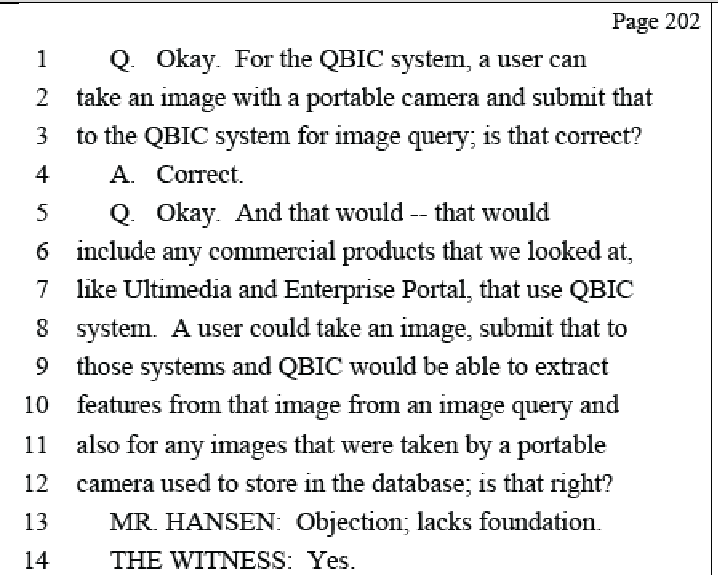 Page 202 Q. Okay. For the QBIC system, a user can take an image with a portable camera and submit that to the QBIC system for image query; is that correct? A. Correct. Q. Okay. And that would -- that would include any commercial products that we looked at, like Ultimedia and Enterprise Portal, that use QBIC system. A user could take an image, submit that to those systems and QBIC would be able to extract features from that image from an image query and also for any images that were taken by a portable camera used to store in the database; is that right? MR. HANSEN: Objection; lacks foundation. THE WITNESS: Yes. *Id.* at 300: 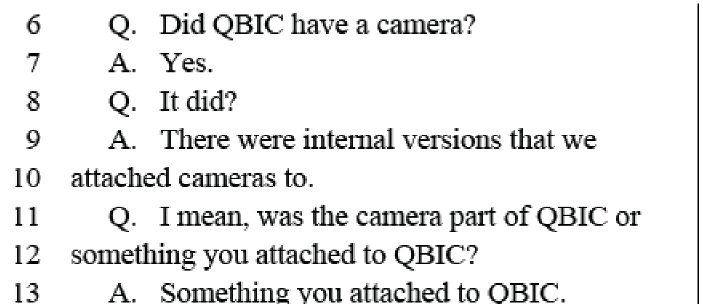 Q. Did QBIC have a camera? A. Yes. Q. It did? A. There were internal versions that we attached cameras to. Q. I mean, was the camera part of QBIC or something you attached to QBIC? A. Something you attached to QBIC. *Id.* at 312:

Element	'036 Claim Recitation	Exemplary Citations in Reference
		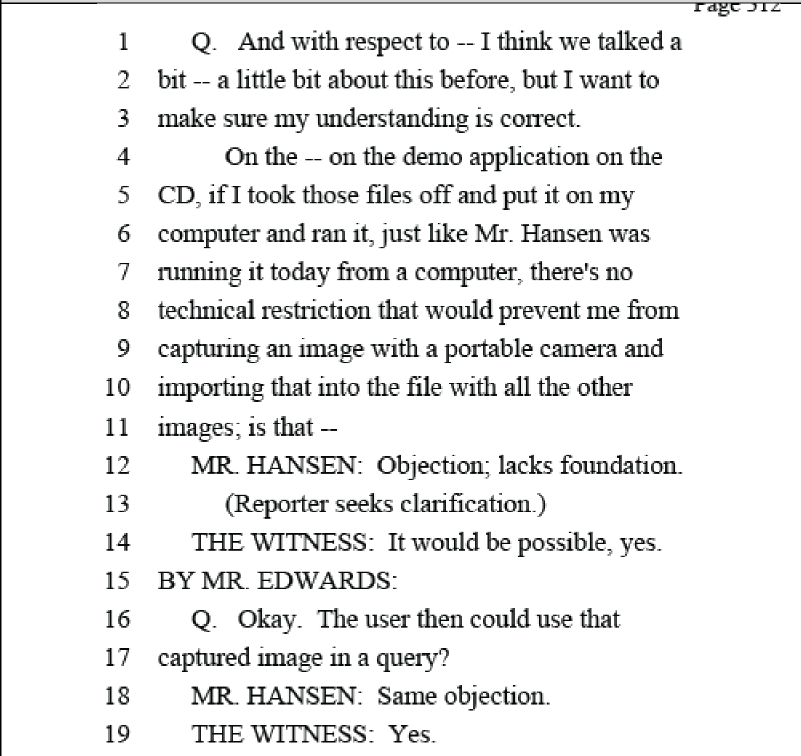Page 312 Q. And with respect to -- I think we talked a bit -- a little bit about this before, but I want to make sure my understanding is correct. On the -- on the demo application on the CD, if I took those files off and put it on my computer and ran it, just like Mr. Hansen was running it today from a computer, there's no technical restriction that would prevent me from capturing an image with a portable camera and importing that into the file with all the other images; is that -- MR. HANSEN: Objection; lacks foundation. (Reporter seeks clarification.) THE WITNESS: It would be possible, yes. BY MR. EDWARDS: Q. Okay. The user then could use that captured image in a query? MR. HANSEN: Same objection. THE WITNESS: Yes. *Id.* at 305: 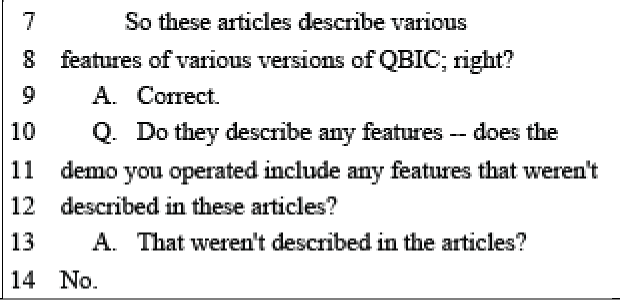So these articles describe various features of various versions of QBIC; right? A. Correct. Q. Do they describe any features -- does the demo you operated include any features that weren't described in these articles? A. That weren't described in the articles? No.

Element	'036 Claim Recitation	Exemplary Citations in Reference
		Id. at 35: 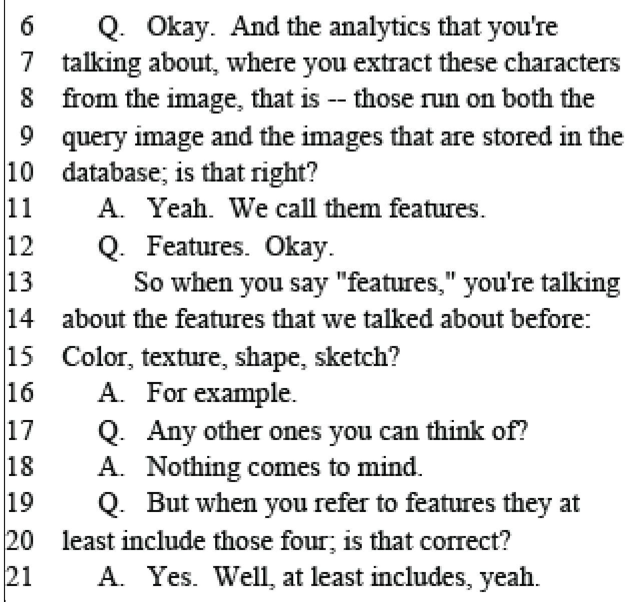 Q. Okay. And the analytics that you're talking about, where you extract these characters from the image, that is -- those run on both the query image and the images that are stored in the database; is that right? A. Yeah. We call them features. Q. Features. Okay. So when you say "features," you're talking about the features that we talked about before: Color, texture, shape, sketch? A. For example. Q. Any other ones you can think of? A. Nothing comes to mind. Q. But when you refer to features they at least include those four; is that correct? A. Yes. Well, at least includes, yeah. *Id.* at 109-110:

Element	'036 Claim Recitation	Exemplary Citations in Reference
		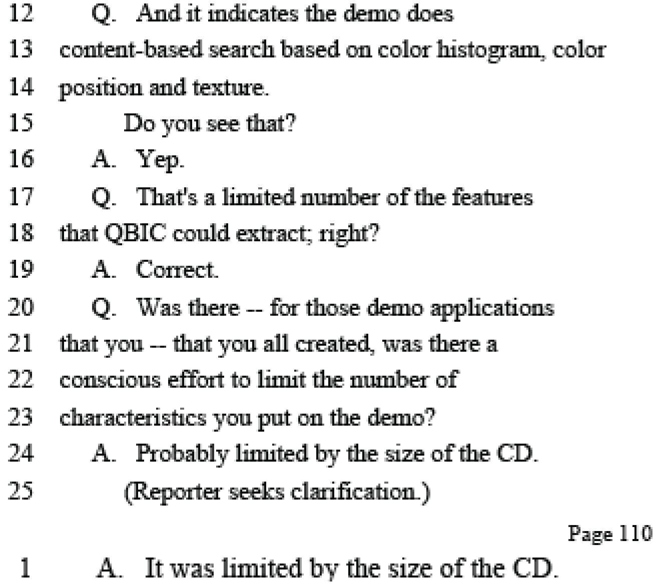 Q. And it indicates the demo does content-based search based on color histogram, color position and texture. Do you see that? A. Yep. Q. That's a limited number of the features that QBIC could extract; right? A. Correct. Q. Was there -- for those demo applications that you -- that you all created, was there a conscious effort to limit the number of characteristics you put on the demo? A. Probably limited by the size of the CD. (Reporter seeks clarification.) Page 110 A. It was limited by the size of the CD. *Id.* at 311: 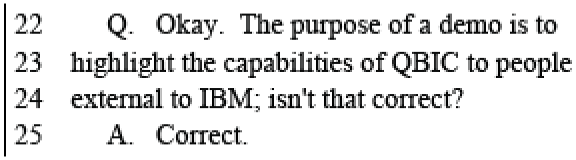 Q. Okay. The purpose of a demo is to highlight the capabilities of QBIC to people external to IBM; isn't that correct? A. Correct. *Id.* at 313-317:

Element	'036 Claim Recitation	Exemplary Citations in Reference
		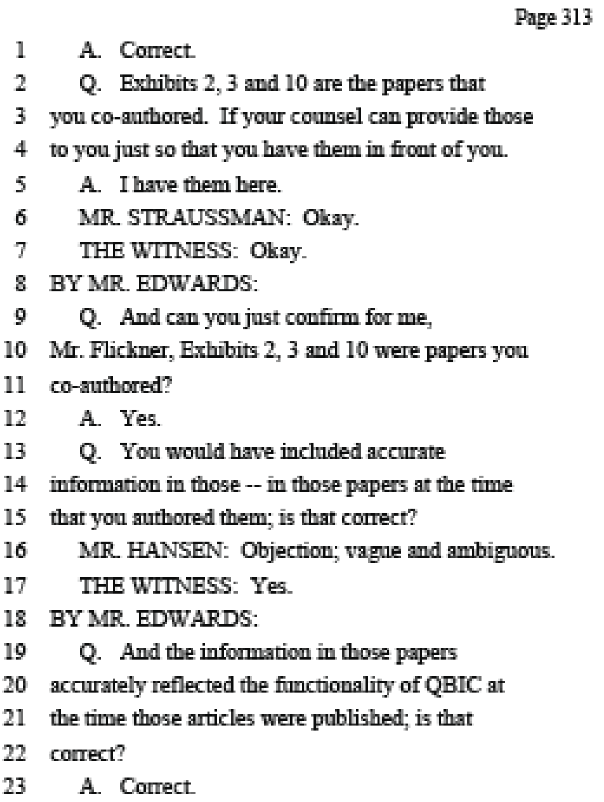 Page 313 A. Correct. Q. Exhibits 2, 3 and 10 are the papers that you co-authored. If your counsel can provide those to you just so that you have them in front of you. A. I have them here. MR. STRAUSSMAN: Okay. THE WITNESS: Okay. BY MR. EDWARDS: Q. And can you just confirm for me, Mr. Flickner, Exhibits 2, 3 and 10 were papers you co-authored? A. Yes. Q. You would have included accurate information in those -- in those papers at the time that you authored them; is that correct? MR. HANSEN: Objection; vague and ambiguous. THE WITNESS: Yes. BY MR. EDWARDS: Q. And the information in those papers accurately reflected the functionality of QBIC at the time those articles were published; is that correct? A. Correct.

		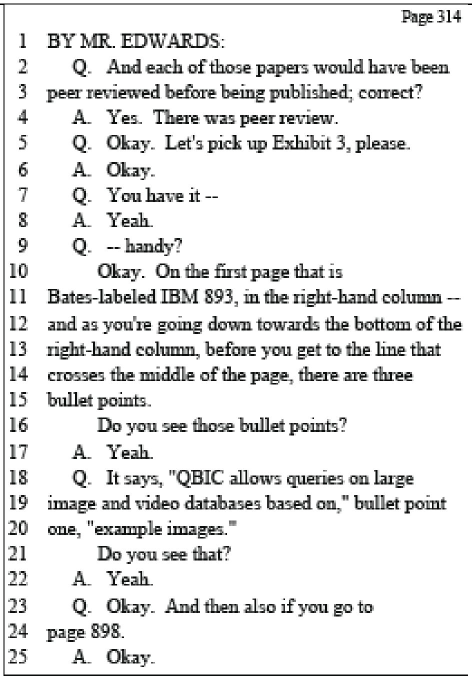 Page 314 BY MR. EDWARDS: Q. And each of those papers would have been peer reviewed before being published; correct? A. Yes. There was peer review. Q. Okay. Let's pick up Exhibit 3, please. A. Okay. Q. You have it -- A. Yeah. Q. -- handy? Okay. On the first page that is Bates-labeled IBM 893, in the right-hand column -- and as you're going down towards the bottom of the right-hand column, before you get to the line that crosses the middle of the page, there are three bullet points. Do you see those bullet points? A. Yeah. Q. It says, "QBIC allows queries on large image and video databases based on," bullet point one, "example images." Do you see that? A. Yeah. Q. Okay. And then also if you go to page 898. A. Okay. 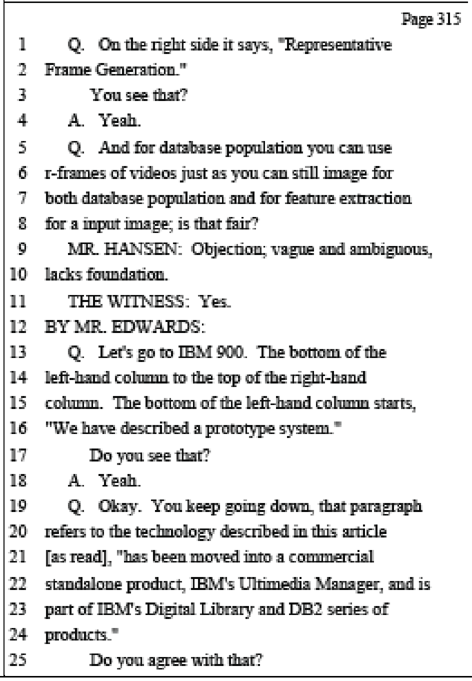 Page 315 Q. On the right side it says, "Representative Frame Generation." You see that? A. Yeah. Q. And for database population you can use r-frames of videos just as you can still image for both database population and for feature extraction for a input image; is that fair? MR. HANSEN: Objection; vague and ambiguous, lacks foundation. THE WITNESS: Yes. BY MR. EDWARDS: Q. Let's go to IBM 900. The bottom of the left-hand column to the top of the right-hand column. The bottom of the left-hand column starts, "We have described a prototype system." Do you see that? A. Yeah. Q. Okay. You keep going down, that paragraph refers to the technology described in this article [as read], "has been moved into a commercial standalone product, IBM's Ultimedia Manager, and is part of IBM's Digital Library and DB2 series of products." Do you agree with that?

Element	'036 Claim Recitation	Exemplary Citations in Reference
		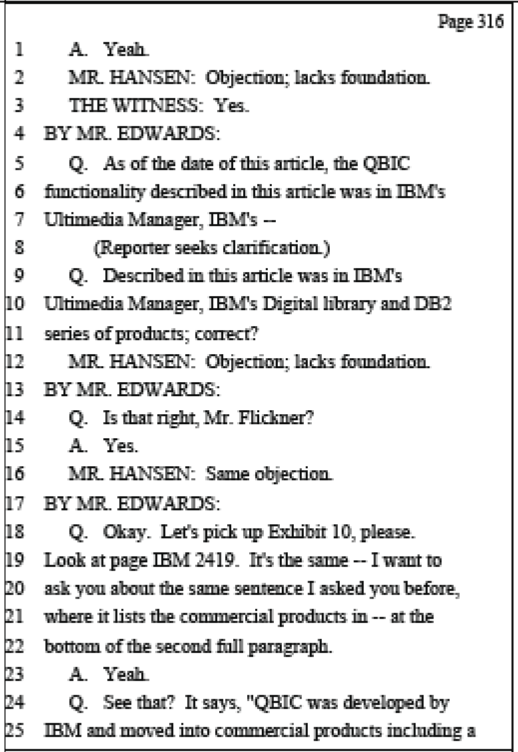 Page 316 A. Yeah. MR. HANSEN: Objection; lacks foundation. THE WITNESS: Yes. BY MR. EDWARDS: Q. As of the date of this article, the QBIC functionality described in this article was in IBM's Ultimedia Manager, IBM's -- (Reporter seeks clarification.) Q. Described in this article was in IBM's Ultimedia Manager, IBM's Digital library and DB2 series of products; correct? MR. HANSEN: Objection; lacks foundation. BY MR. EDWARDS: Q. Is that right, Mr. Flickner? A. Yes. MR. HANSEN: Same objection. BY MR. EDWARDS: Q. Okay. Let's pick up Exhibit 10, please. Look at page IBM 2419. It's the same -- I want to ask you about the same sentence I asked you before, where it lists the commercial products in -- at the bottom of the second full paragraph. A. Yeah. Q. See that? It says, "QBIC was developed by IBM and moved into commercial products including a 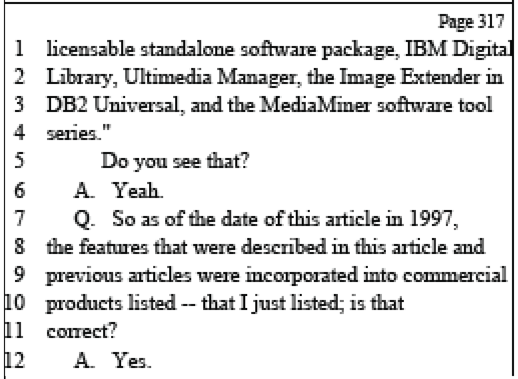 Page 317 licensable standalone software package, IBM Digital Library, Ultimedia Manager, the Image Extender in DB2 Universal, and the MediaMiner software tool series." Do you see that? A. Yeah. Q. So as of the date of this article in 1997, the features that were described in this article and previous articles were incorporated into commercial products listed -- that I just listed; is that correct? A. Yes.

Element	'036 Claim Recitation	Exemplary Citations in Reference
		See also Flickner Depo at, e.g., 24, 41-42, 62-63 (user could take a picture or video and submit it to the QBIC system). To the extent it is found that QBIC does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that QBIC does not anticipate this claim, QBIC renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine QBIC with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1B.iii]	recognizes the target as a known target from the target database based on comparing parameters derived from the digital representation to recognition parameters associated with the known targets; and	QBIC discloses the target as a known target from the target database based on comparing parameters derived from the digital representation to recognition parameters associated with the known targets. The QBIC Project: Querying Images by Content Using Color, Texture, and Shape (IBM 0002390–404) at 2390 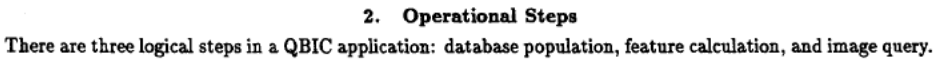 **2. Operational Steps** There are three logical steps in a QBIC application: database population, feature calculation, and image query. IBM0002390 at 2391:

Element	'036 Claim Recitation	Exemplary Citations in Reference
		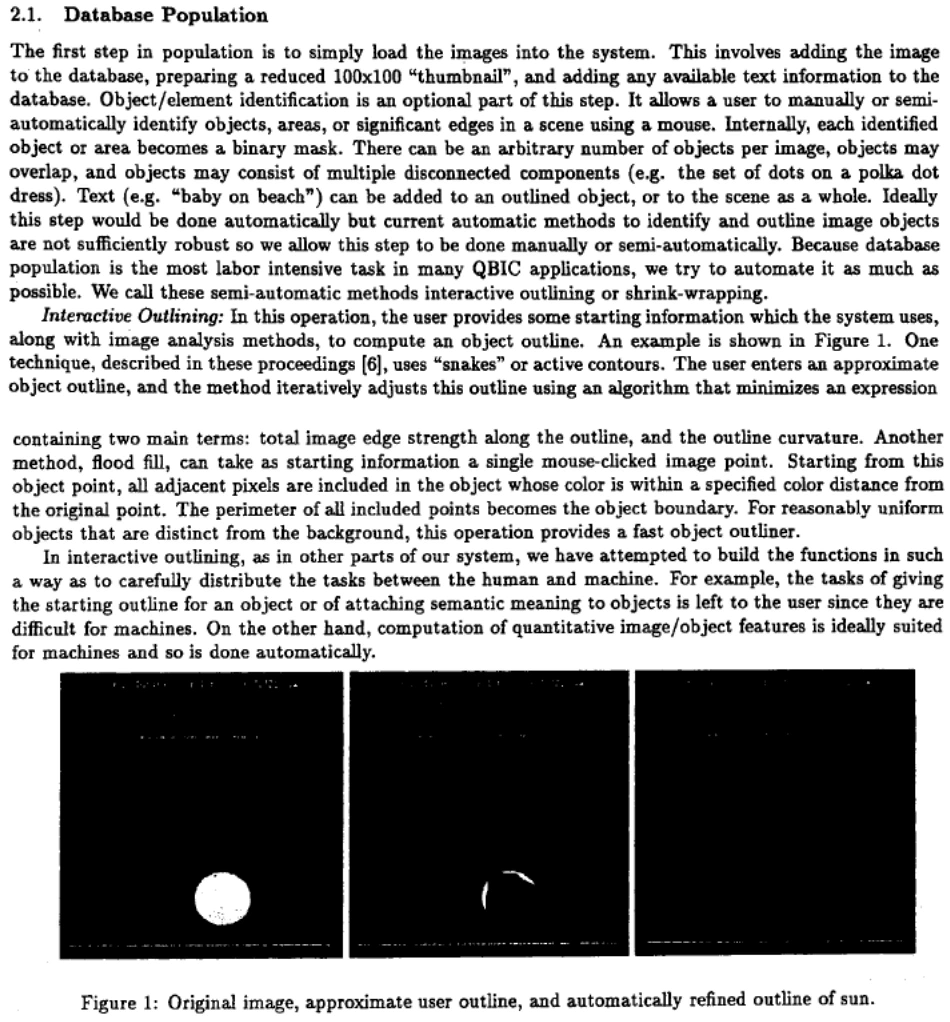**2.1. Database Population** The first step in population is to simply load the images into the system. This involves adding the image to the database, preparing a reduced 100x100 "thumbnail", and adding any available text information to the database. Object/element identification is an optional part of this step. It allows a user to manually or semi-automatically identify objects, areas, or significant edges in a scene using a mouse. Internally, each identified object or area becomes a binary mask. There can be an arbitrary number of objects per image, objects may overlap, and objects may consist of multiple disconnected components (e.g. the set of dots on a polka dot dress). Text (e.g. "baby on beach") can be added to an outlined object, or to the scene as a whole. Ideally this step would be done automatically but current automatic methods to identify and outline image objects are not sufficiently robust so we allow this step to be done manually or semi-automatically. Because database population is the most labor intensive task in many QBIC applications, we try to automate it as much as possible. We call these semi-automatic methods interactive outlining or shrink-wrapping. *Interactive Outlining:* In this operation, the user provides some starting information which the system uses, along with image analysis methods, to compute an object outline. An example is shown in Figure 1. One technique, described in these proceedings [6], uses "snakes" or active contours. The user enters an approximate object outline, and the method iteratively adjusts this outline using an algorithm that minimizes an expression containing two main terms: total image edge strength along the outline, and the outline curvature. Another method, flood fill, can take as starting information a single mouse-clicked image point. Starting from this object point, all adjacent pixels are included in the object whose color is within a specified color distance from the original point. The perimeter of all included points becomes the object boundary. For reasonably uniform objects that are distinct from the background, this operation provides a fast object outliner. In interactive outlining, as in other parts of our system, we have attempted to build the functions in such a way as to carefully distribute the tasks between the human and machine. For example, the tasks of giving the starting outline for an object or of attaching semantic meaning to objects is left to the user since they are difficult for machines. On the other hand, computation of quantitative image/object features is ideally suited for machines and so is done automatically. Figure 1: Original image, approximate user outline, and automatically refined outline of sun. IBM0002390 at 2391-92:

Element	'036 Claim Recitation	Exemplary Citations in Reference
		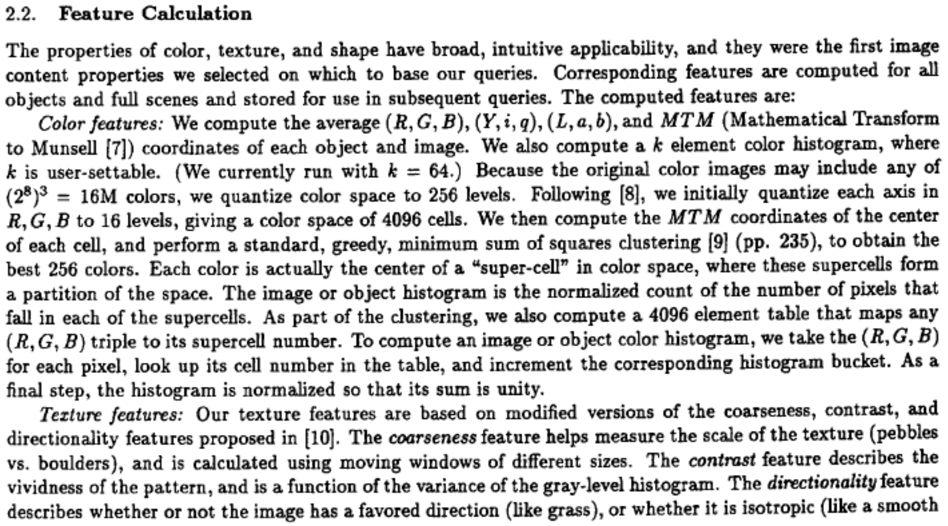 2.2. **Feature Calculation** The properties of color, texture, and shape have broad, intuitive applicability, and they were the first image content properties we selected on which to base our queries. Corresponding features are computed for all objects and full scenes and stored for use in subsequent queries. The computed features are: *Color features:* We compute the average (R, G, B), (Y, i, q), (L, a, b), and MTM (Mathematical Transform to Munsell [7]) coordinates of each object and image. We also compute a k element color histogram, where k is user-settable. (We currently run with $k = 64$.) Because the original color images may include any of $(2^8)^3 = 16$M colors, we quantize color space to 256 levels. Following [8], we initially quantize each axis in R, G, B to 16 levels, giving a color space of 4096 cells. We then compute the MTM coordinates of the center of each cell, and perform a standard, greedy, minimum sum of squares clustering [9] (pp. 235), to obtain the best 256 colors. Each color is actually the center of a "super-cell" in color space, where these supercells form a partition of the space. The image or object histogram is the normalized count of the number of pixels that fall in each of the supercells. As part of the clustering, we also compute a 4096 element table that maps any (R, G, B) triple to its supercell number. To compute an image or object color histogram, we take the (R, G, B) for each pixel, look up its cell number in the table, and increment the corresponding histogram bucket. As a final step, the histogram is normalized so that its sum is unity. *Texture features:* Our texture features are based on modified versions of the coarseness, contrast, and directionality features proposed in [10]. The *coarseness* feature helps measure the scale of the texture (pebbles vs. boulders), and is calculated using moving windows of different sizes. The *contrast* feature describes the vividness of the pattern, and is a function of the variance of the gray-level histogram. The *directionality* feature describes whether or not the image has a favored direction (like grass), or whether it is isotropic (like a smooth

Element	'036 Claim Recitation	Exemplary Citations in Reference
		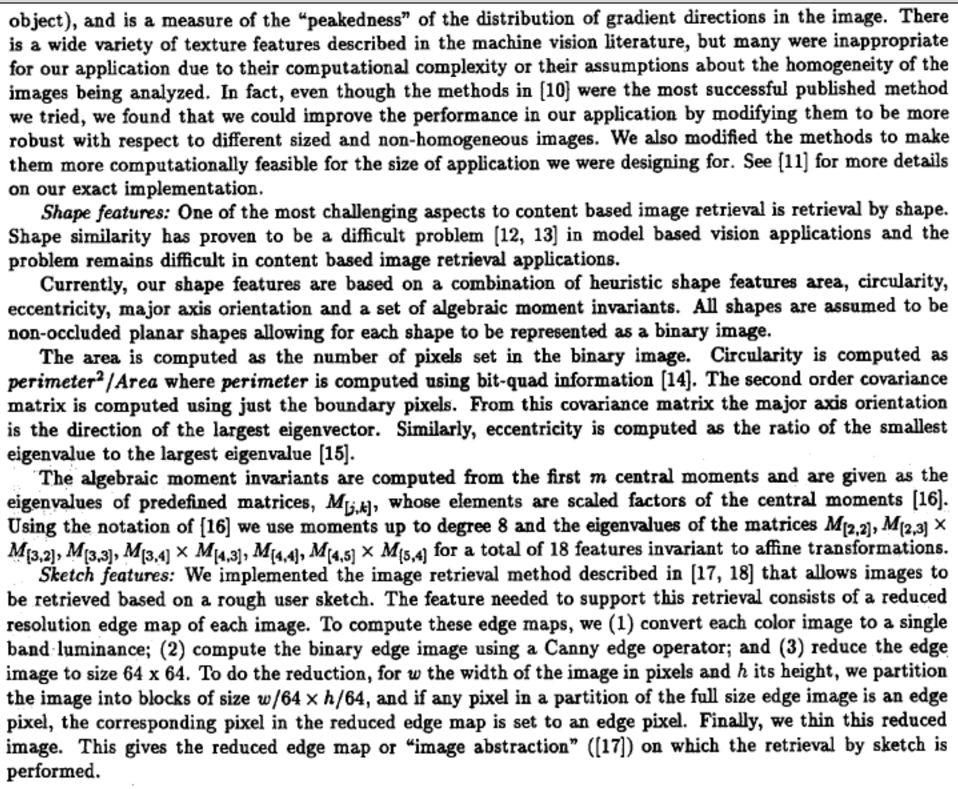 object), and is a measure of the "peakedness" of the distribution of gradient directions in the image. There is a wide variety of texture features described in the machine vision literature, but many were inappropriate for our application due to their computational complexity or their assumptions about the homogeneity of the images being analyzed. In fact, even though the methods in [10] were the most successful published method we tried, we found that we could improve the performance in our application by modifying them to be more robust with respect to different sized and non-homogeneous images. We also modified the methods to make them more computationally feasible for the size of application we were designing for. See [11] for more details on our exact implementation. *Shape features:* One of the most challenging aspects to content based image retrieval is retrieval by shape. Shape similarity has proven to be a difficult problem [12, 13] in model based vision applications and the problem remains difficult in content based image retrieval applications. Currently, our shape features are based on a combination of heuristic shape features area, circularity, eccentricity, major axis orientation and a set of algebraic moment invariants. All shapes are assumed to be non-occluded planar shapes allowing for each shape to be represented as a binary image. The area is computed as the number of pixels set in the binary image. Circularity is computed as $perimeter^2/Area$ where $perimeter$ is computed using bit-quad information [14]. The second order covariance matrix is computed using just the boundary pixels. From this covariance matrix the major axis orientation is the direction of the largest eigenvector. Similarly, eccentricity is computed as the ratio of the smallest eigenvalue to the largest eigenvalue [15]. The algebraic moment invariants are computed from the first m central moments and are given as the eigenvalues of predefined matrices, $M_{[j,k]}$, whose elements are scaled factors of the central moments [16]. Using the notation of [16] we use moments up to degree 8 and the eigenvalues of the matrices $M_{[2,2]}, M_{[2,3]} \times M_{[3,2]}, M_{[3,3]}, M_{[3,4]} \times M_{[4,3]}, M_{[4,4]}, M_{[4,5]} \times M_{[5,4]}$ for a total of 18 features invariant to affine transformations. *Sketch features:* We implemented the image retrieval method described in [17, 18] that allows images to be retrieved based on a rough user sketch. The feature needed to support this retrieval consists of a reduced resolution edge map of each image. To compute these edge maps, we (1) convert each color image to a single band luminance; (2) compute the binary edge image using a Canny edge operator; and (3) reduce the edge image to size 64 x 64. To do the reduction, for w the width of the image in pixels and h its height, we partition the image into blocks of size $w/64 \times h/64$, and if any pixel in a partition of the full size edge image is an edge pixel, the corresponding pixel in the reduced edge map is set to an edge pixel. Finally, we thin this reduced image. This gives the reduced edge map or "image abstraction" ([17]) on which the retrieval by sketch is performed. IBM0002390 at 2392-93:

Element	'036 Claim Recitation	Exemplary Citations in Reference
		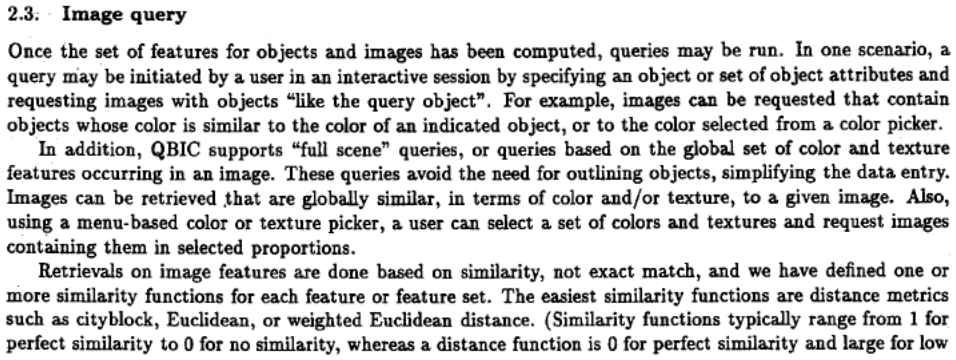 **2.3. Image query** Once the set of features for objects and images has been computed, queries may be run. In one scenario, a query may be initiated by a user in an interactive session by specifying an object or set of object attributes and requesting images with objects "like the query object". For example, images can be requested that contain objects whose color is similar to the color of an indicated object, or to the color selected from a color picker. In addition, QBIC supports "full scene" queries, or queries based on the global set of color and texture features occurring in an image. These queries avoid the need for outlining objects, simplifying the data entry. Images can be retrieved that are globally similar, in terms of color and/or texture, to a given image. Also, using a menu-based color or texture picker, a user can select a set of colors and textures and request images containing them in selected proportions. Retrievals on image features are done based on similarity, not exact match, and we have defined one or more similarity functions for each feature or feature set. The easiest similarity functions are distance metrics such as cityblock, Euclidean, or weighted Euclidean distance. (Similarity functions typically range from 1 for perfect similarity to 0 for no similarity, whereas a distance function is 0 for perfect similarity and large for low

Element	'036 Claim Recitation	Exemplary Citations in Reference
		similarity, so we must convert as necessary to obtain consistent measures.) Different similarity measures may be used, and they may vary from application to application. The matching methods and similarity functions for the different feature types are described below: *Color:* For average color, the distance between a query object and database object is weighted Euclidean distance. The weights are user adjustable so that we may, for example, downweight the L component of an (L, a, b) query. Best results, as judged by our subjective evaluation in on-going experiments, are obtained by representing the average color in the MTM color space, and by inversely weighting each component by its standard deviation over the samples in the database. Improved results are obtained when color matching is done not on average color, but on the distribution of colors that occur in an object or image. In this way, a red and blue object matches other red and blue objects better than to a purple object. To match color histograms, we use a method similar to the one described in [8]. Let X be the query histogram and Y the histogram of an item in the database, both normalized. We compute the element by element difference histogram Z. Then the similarity between X and Y is given by $\|Z\| = Z^T A Z$, where A is a symmetric color similarity matrix with $$a(i,j) \equiv 1 - d(c_i, c_j)/d_{max}$$ where c_i and c_j are the ith and jth colors in the color histograms, and $d(c_i, c_j)$ is the MTM color distance, and d_{max} is the maximum distance between any two colors ([8] used Luv distance.) This metric gives the (weighted) length of the difference vector between X and Y, weighted by A, which accounts for the perceptual distance between different pairs of colors. The result is a measure that accounts for both the difference between the amounts in any given color (e.g. a particular shade of red), as well as differences between similar colors (e.g. red and orange). Based on our experience, we also found it useful for the user to be able to request images with $x\%$ of color 1, $y\%$ of color 2, etc. where both the percentages and the colors are explicitly specified by the user. This allows searches such as for a beach scene (25% white (for the sand) and 50% blue (for the sky and water)), and for many scenes for which the user has seen before, and for which he/she can remember approximate color percentages. The color picker that allows this type of query is shown in Figure 2. 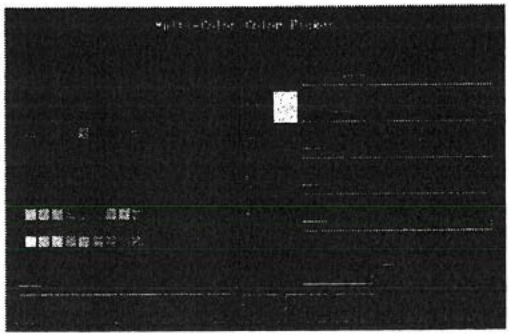 Figure 2: Multi-Color Color Picker with two selected colors. The palette on the left displays the set of selectable colors. Once a color is selected, it appears in a box on the right, and the amount of the color can be specified by the sliders. In this case, 25% light blue and 15% white have been selected. *Texture:* Texture distance is computed as weighted Euclidean distance in the three dimensional texture space. The most common weighting factors are the inverse variances for each component, computed over the

Element	'036 Claim Recitation	Exemplary Citations in Reference
		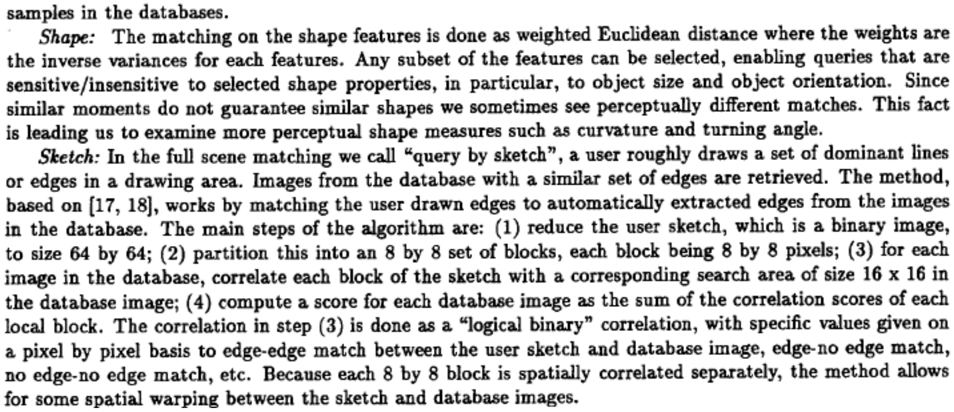 samples in the databases. *Shape:* The matching on the shape features is done as weighted Euclidean distance where the weights are the inverse variances for each features. Any subset of the features can be selected, enabling queries that are sensitive/insensitive to selected shape properties, in particular, to object size and object orientation. Since similar moments do not guarantee similar shapes we sometimes see perceptually different matches. This fact is leading us to examine more perceptual shape measures such as curvature and turning angle. *Sketch:* In the full scene matching we call "query by sketch", a user roughly draws a set of dominant lines or edges in a drawing area. Images from the database with a similar set of edges are retrieved. The method, based on [17, 18], works by matching the user drawn edges to automatically extracted edges from the images in the database. The main steps of the algorithm are: (1) reduce the user sketch, which is a binary image, to size 64 by 64; (2) partition this into an 8 by 8 set of blocks, each block being 8 by 8 pixels; (3) for each image in the database, correlate each block of the sketch with a corresponding search area of size 16 x 16 in the database image; (4) compute a score for each database image as the sum of the correlation scores of each local block. The correlation in step (3) is done as a "logical binary" correlation, with specific values given on a pixel by pixel basis to edge-edge match between the user sketch and database image, edge-no edge match, no edge-no edge match, etc. Because each 8 by 8 block is spatially correlated separately, the method allows for some spatial warping between the sketch and database images. IBM0002390 at 2393-95: QBIC (Query by Image Content) (BOFA00001391–392) at 1:

Element	'036 Claim Recitation	Exemplary Citations in Reference
		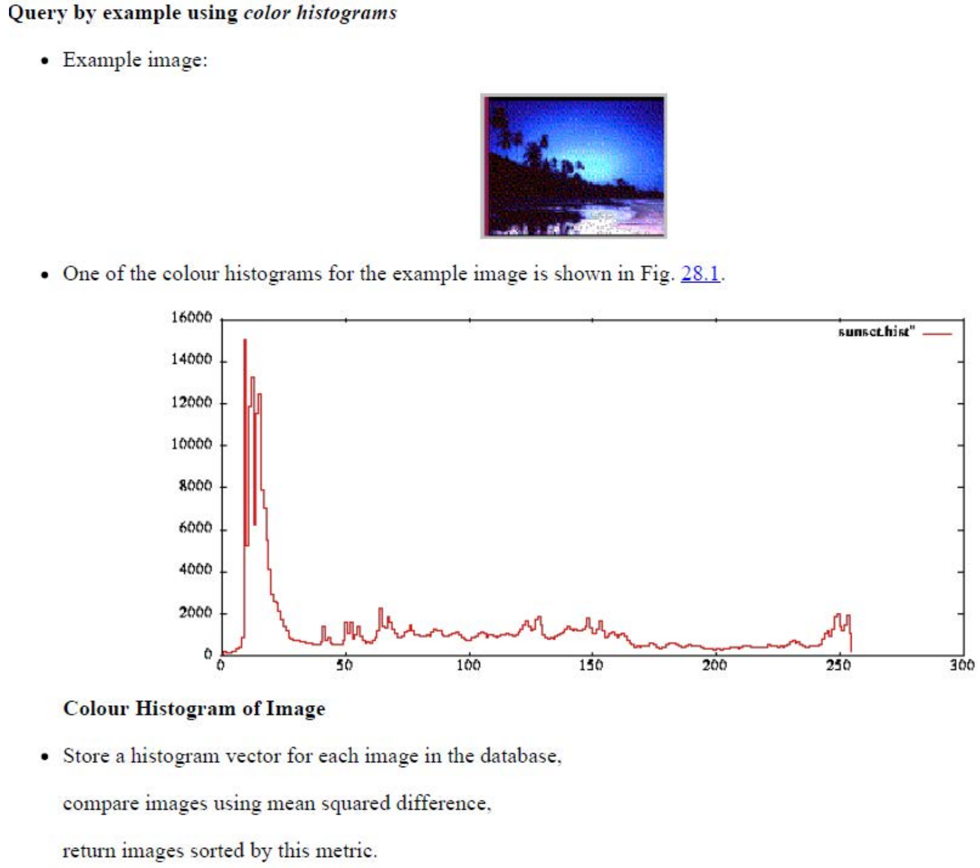 **Query by example using *color histograms*** - Example image: - One of the colour histograms for the example image is shown in Fig. 28.1. **Colour Histogram of Image** - Store a histogram vector for each image in the database. compare images using mean squared difference. return images sorted by this metric. Query by Image and Video Content: The QBIC System (IBM 0000893–902): "To achieve this functionality, QBIC has two main components: database population (the process of creating an image database) and database query. During the population, images and videos are processed to extract features describing their content-colors, textures, shapes, and camera and object motion-and the features are stored in a database. During the query, the user composes a query graphically. Features are generated from the graphical query and then input to a matching

Element	'036 Claim Recitation	Exemplary Citations in Reference
		engine that finds images or videos from the database with similar features. Figure 2 shows the system architecture." *Id.* at p. 3. 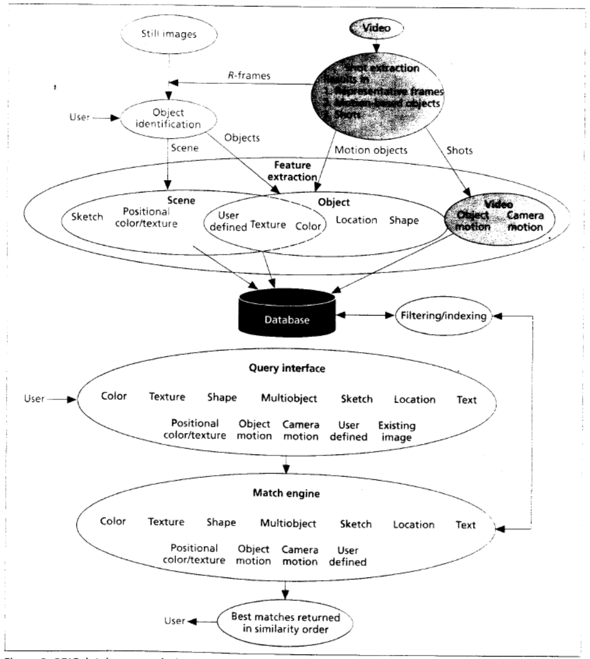 Figure 2. QBIC database population (top) and query (bottom) architecture. *Id*. at Fig. 2. "For both population and query, the QBIC data model has • still images or scenes (full images) that contain objects (subsets of an image), and • video shots that consist of sets of contiguous frames and contain motion objects. For still images, the QBIC data model distinguishes between 'scenes' (or images) and 'objects.' A scene is an image or single representative frame of video. An object is a part of a scene-for example, the fox in Figure 3-or a moving entity in a video. For still image database population, features are extracted from images and objects and stored in a database as shown in the top left part of Figure 2. Videos are broken into clips called shots. Representative frames, or r-frames, are

Element	'036 Claim Recitation	Exemplary Citations in Reference
		generated for each extracted shot. R-frames are treated as still images, and features are extracted and stored in the database. Further processing of shots generates motion objects-for example, a car moving across the screen. Queries are allowed on objects ('Find images with a red, round object'), scenes ('Find images that have approximately 30-percent red and IS-percent blue colors'), shots ('Find all shots panning from left to right'), or any combination ('Find images that have 30 percent red and contain a blue textured object')." *Id.* at p. 3. Image, Audio, and Video Extenders Administration and Programming (IBM 000001–617): "To make an image available for searching by content, you catalog the image. When you catalog an image, the Image Extender analyzes the image, by computing the feature values for the image, and stores the values in a QBIC catalog. When you search for an image by content, your query identifies one or more features for the search (such as average color), a source for each feature (such as an example image), and a target set of cataloged images. The Image Extender computes the feature value of the source and compares it to the cataloged feature values for the target images. It then computes a score that indicates how similar the feature values of the target images are to the source. You can have the Image Extender return the images whose features are most similar to the source. The Image Extender will return the handle of each image and the image score. You can also have the Image Extender return only the score of a single image." *Id.* at pp. 20–21. QBIC Demo (IBM000747):

Element	'036 Claim Recitation	Exemplary Citations in Reference
		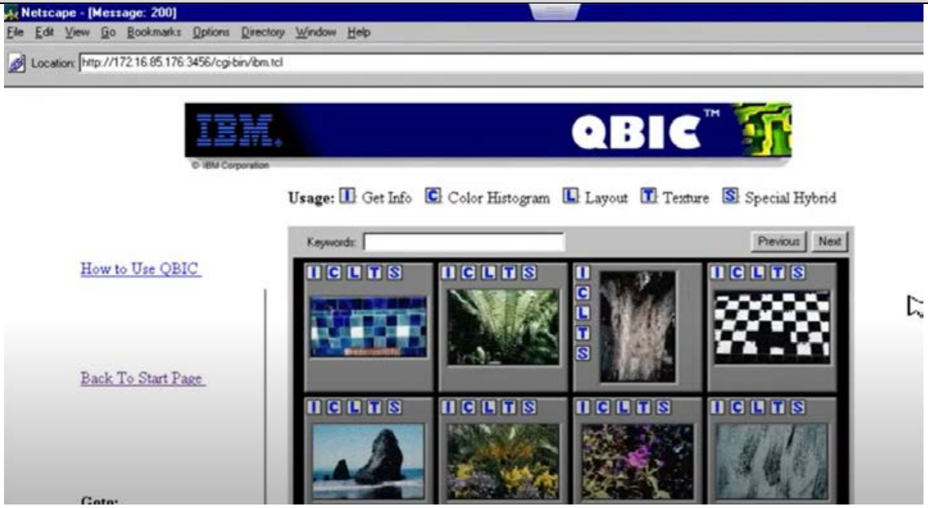 QBIC Demo (IBM000747): 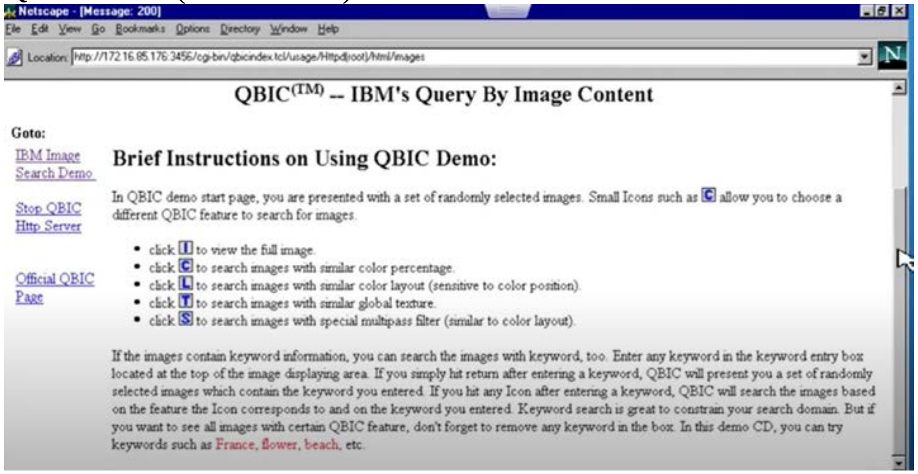 QBIC(TM) -- IBM's Query By Image Content Goto: IBM Image Search Demo Stop QBIC Http Server Official QBIC Page **Brief Instructions on Using QBIC Demo:** In QBIC demo start page, you are presented with a set of randomly selected images. Small Icons such as C allow you to choose a different QBIC feature to search for images. - click I to view the full image. - click C to search images with similar color percentage. - click L to search images with similar color layout (sensitive to color position). - click T to search images with similar global texture. - click S to search images with special multipass filter (similar to color layout). If the images contain keyword information, you can search the images with keyword, too. Enter any keyword in the keyword entry box located at the top of the image displaying area. If you simply hit return after entering a keyword, QBIC will present you a set of randomly selected images which contain the keyword you entered. If you hit any Icon after entering a keyword, QBIC will search the images based on the feature the Icon corresponds to and on the keyword you entered. Keyword search is great to constrain your search domain. But if you want to see all images with certain QBIC feature, don't forget to remove any keyword in the box. In this demo CD, you can try keywords such as France, flower, beach, etc. QBIC Demo (IBM000747):

Element	'036 Claim Recitation	Exemplary Citations in Reference
		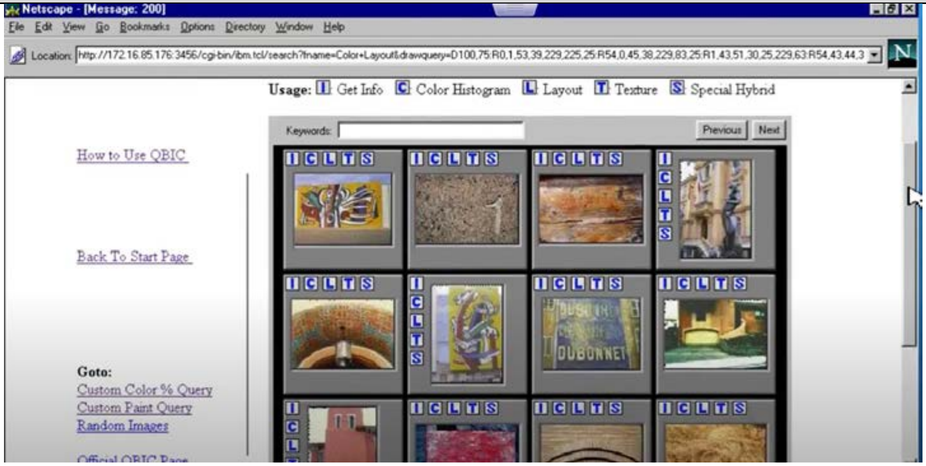 Location: http://172.16.85.176:3456/cgi-bin/ibm.tcl/search?fname=Color+Layout&drawquery=D100,75:R0,1,53,39,229,225,25:R54,0,45,38,229,83,25:R1,43,51,30,25,229,63:R54,43,44,3 QBIC Demo (IBM000747): 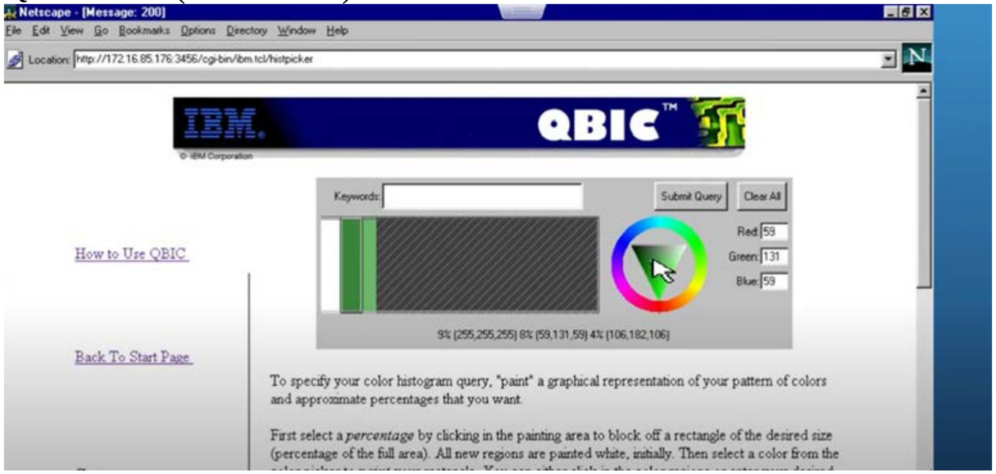 Location: http://172.16.85.176:3456/cgi-bin/ibm.tcl/histpicker Red: 59 Green: 131 Blue: 59 9% (255,255,255) 8% (59,131,59) 4% (106,182,106) To specify your color histogram query, "paint" a graphical representation of your pattern of colors and approximate percentages that you want. First select a *percentage* by clicking in the painting area to block off a rectangle of the desired size (percentage of the full area). All new regions are painted white, initially. Then select a color from the QBIC Demo (IBM000747):

Element	'036 Claim Recitation	Exemplary Citations in Reference
		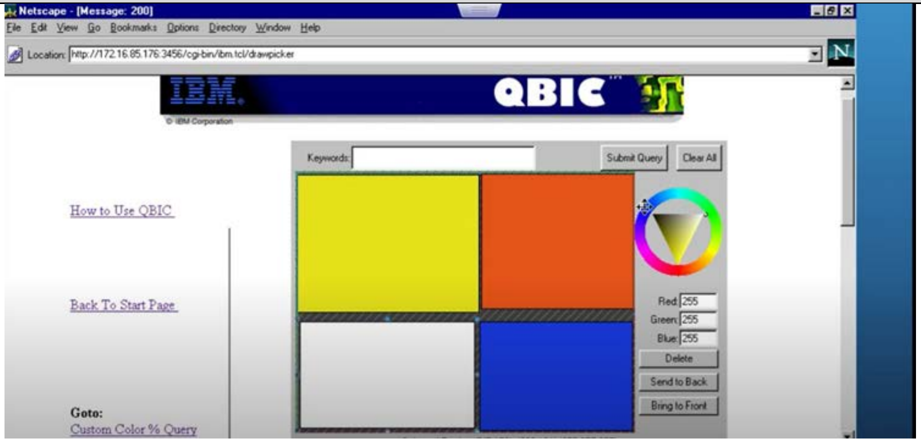 Location: http://172.16.85.176:3456/cgi-bin/ibm.tcl/drawpicker Keywords: Red: 255 Green: 255 Blue: 255 Goto: Custom Color % Query  Location: http://172.16.85.176:3456/cgi-bin/ibm.tcl/search?fname=HybridColor&france10%2Ejpg.x=0 Goto: Custom Color % Query Custom Paint Query Random Images Official QBIC Page Query was: Example: =france10.jpg Query Type: QbHybridColorFeatureClass Flickner Depo at 21:

Element	'036 Claim Recitation	Exemplary Citations in Reference
		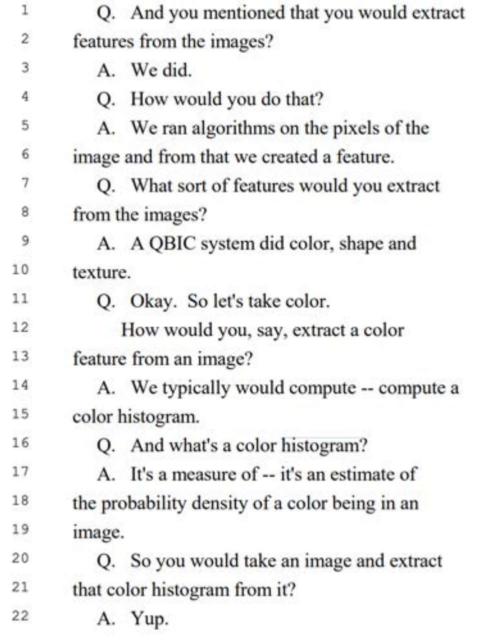 Q. And you mentioned that you would extract features from the images? A. We did. Q. How would you do that? A. We ran algorithms on the pixels of the image and from that we created a feature. Q. What sort of features would you extract from the images? A. A QBIC system did color, shape and texture. Q. Okay. So let's take color. How would you, say, extract a color feature from an image? A. We typically would compute -- compute a color histogram. Q. And what's a color histogram? A. It's a measure of -- it's an estimate of the probability density of a color being in an image. Q. So you would take an image and extract that color histogram from it? A. Yup. Flickner Depo at 88: 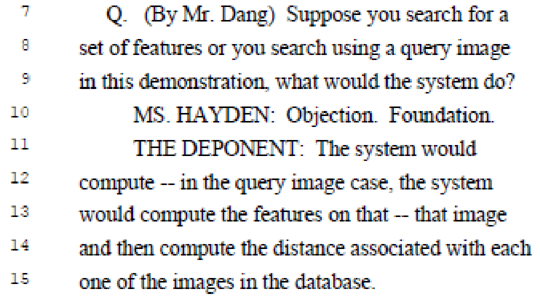 Q. (By Mr. Dang) Suppose you search for a set of features or you search using a query image in this demonstration, what would the system do? MS. HAYDEN: Objection. Foundation. THE DEPONENT: The system would compute -- in the query image case, the system would compute the features on that -- that image and then compute the distance associated with each one of the images in the database.

Element	'036 Claim Recitation	Exemplary Citations in Reference
		Flickner Depo at 25-26: 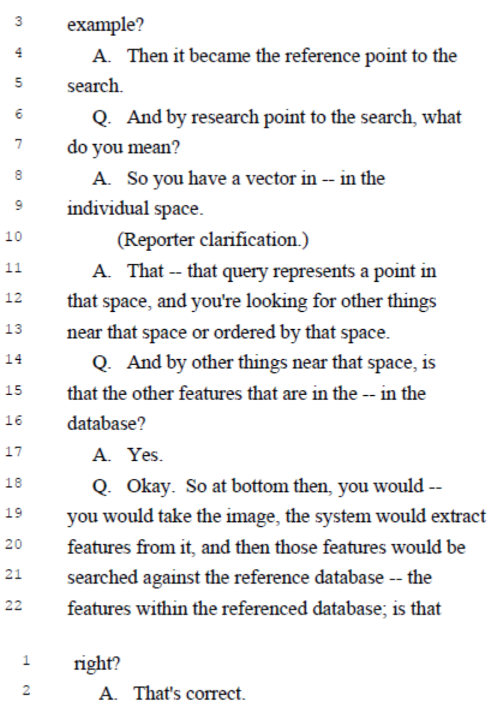 Q. Okay. And then what would it do with those features that it had extracted from that example? A. Then it became the reference point to the search. Q. And by research point to the search, what do you mean? A. So you have a vector in -- in the individual space. (Reporter clarification.) A. That -- that query represents a point in that space, and you're looking for other things near that space or ordered by that space. Q. And by other things near that space, is that the other features that are in the -- in the database? A. Yes. Q. Okay. So at bottom then, you would -- you would take the image, the system would extract features from it, and then those features would be searched against the reference database -- the features within the referenced database; is that right? A. That's correct.

Element	'036 Claim Recitation	Exemplary Citations in Reference
		Flickner Depo at 75-76 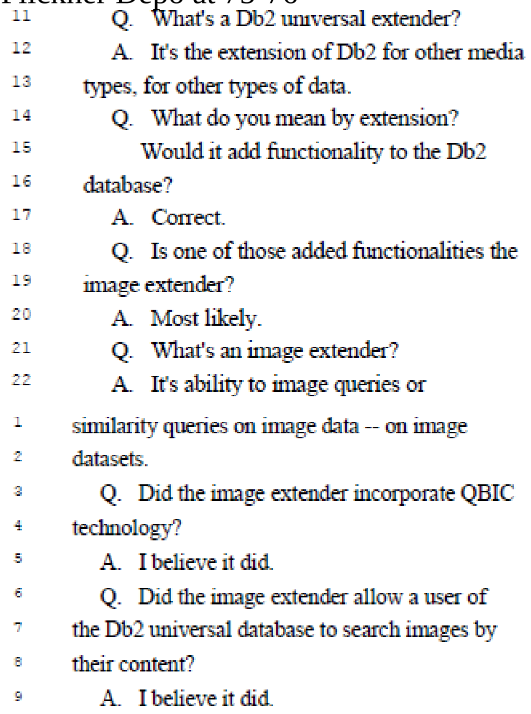 Q. What's a Db2 universal extender? A. It's the extension of Db2 for other media types, for other types of data. Q. What do you mean by extension? Would it add functionality to the Db2 database? A. Correct. Q. Is one of those added functionalities the image extender? A. Most likely. Q. What's an image extender? A. It's ability to image queries or similarity queries on image data -- on image datasets. Q. Did the image extender incorporate QBIC technology? A. I believe it did. Q. Did the image extender allow a user of the Db2 universal database to search images by their content? A. I believe it did. Flickner Depo at 40-42:

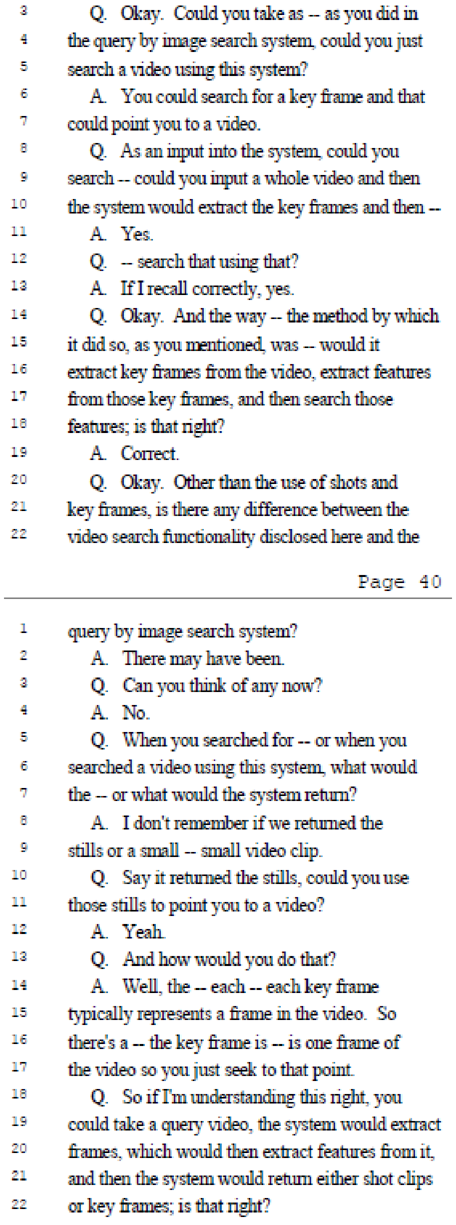

Q. Okay. Could you take as -- as you did in the query by image search system, could you just search a video using this system?
A. You could search for a key frame and that could point you to a video.
Q. As an input into the system, could you search -- could you input a whole video and then the system would extract the key frames and then --
A. Yes.
Q. -- search that using that?
A. If I recall correctly, yes.
Q. Okay. And the way -- the method by which it did so, as you mentioned, was -- would it extract key frames from the video, extract features from those key frames, and then search those features; is that right?
A. Correct.
Q. Okay. Other than the use of shots and key frames, is there any difference between the video search functionality disclosed here and the

Page 40

query by image search system?
A. There may have been.
Q. Can you think of any now?
A. No.
Q. When you searched for -- or when you searched a video using this system, what would the -- or what would the system return?
A. I don't remember if we returned the stills or a small -- small video clip.
Q. Say it returned the stills, could you use those stills to point you to a video?
A. Yeah.
Q. And how would you do that?
A. Well, the -- each -- each key frame typically represents a frame in the video. So there's a -- the key frame is -- is one frame of the video so you just seek to that point.
Q. So if I'm understanding this right, you could take a query video, the system would extract frames, which would then extract features from it, and then the system would return either shot clips or key frames; is that right?

Element	'036 Claim Recitation	Exemplary Citations in Reference
		 A. Right. Flickner Depo at 87: 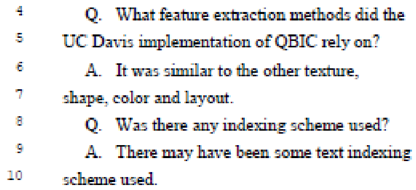 Q. What feature extraction methods did the UC Davis implementation of QBIC rely on? A. It was similar to the other texture, shape, color and layout. Q. Was there any indexing scheme used? A. There may have been some text indexing scheme used. Flickner 8/29/23 Depo at 33–37:

Element	'036 Claim Recitation	Exemplary Citations in Reference
		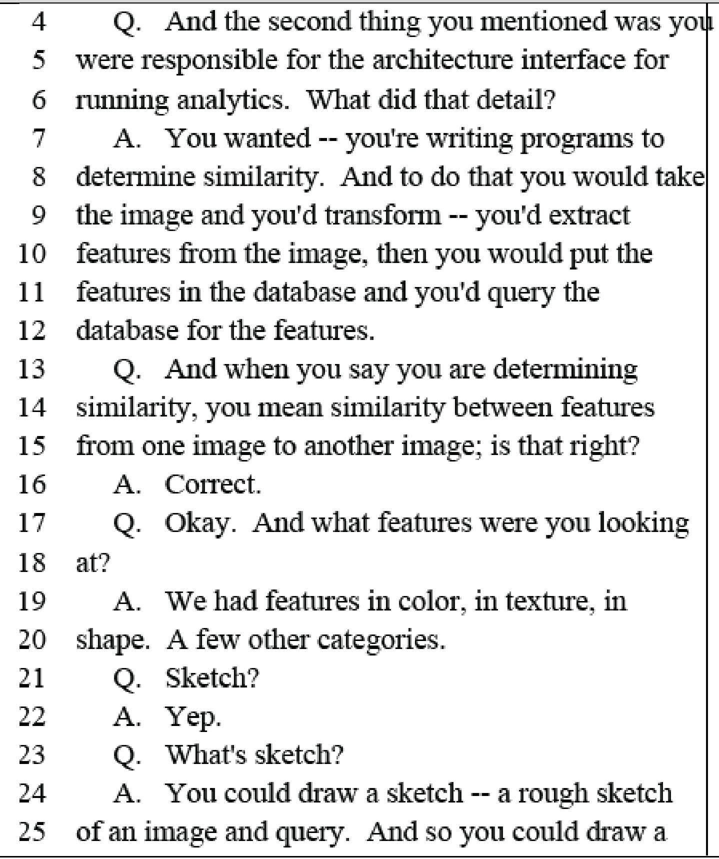 Q. And the second thing you mentioned was you were responsible for the architecture interface for running analytics. What did that detail? A. You wanted -- you're writing programs to determine similarity. And to do that you would take the image and you'd transform -- you'd extract features from the image, then you would put the features in the database and you'd query the database for the features. Q. And when you say you are determining similarity, you mean similarity between features from one image to another image; is that right? A. Correct. Q. Okay. And what features were you looking at? A. We had features in color, in texture, in shape. A few other categories. Q. Sketch? A. Yep. Q. What's sketch? A. You could draw a sketch -- a rough sketch of an image and query. And so you could draw a

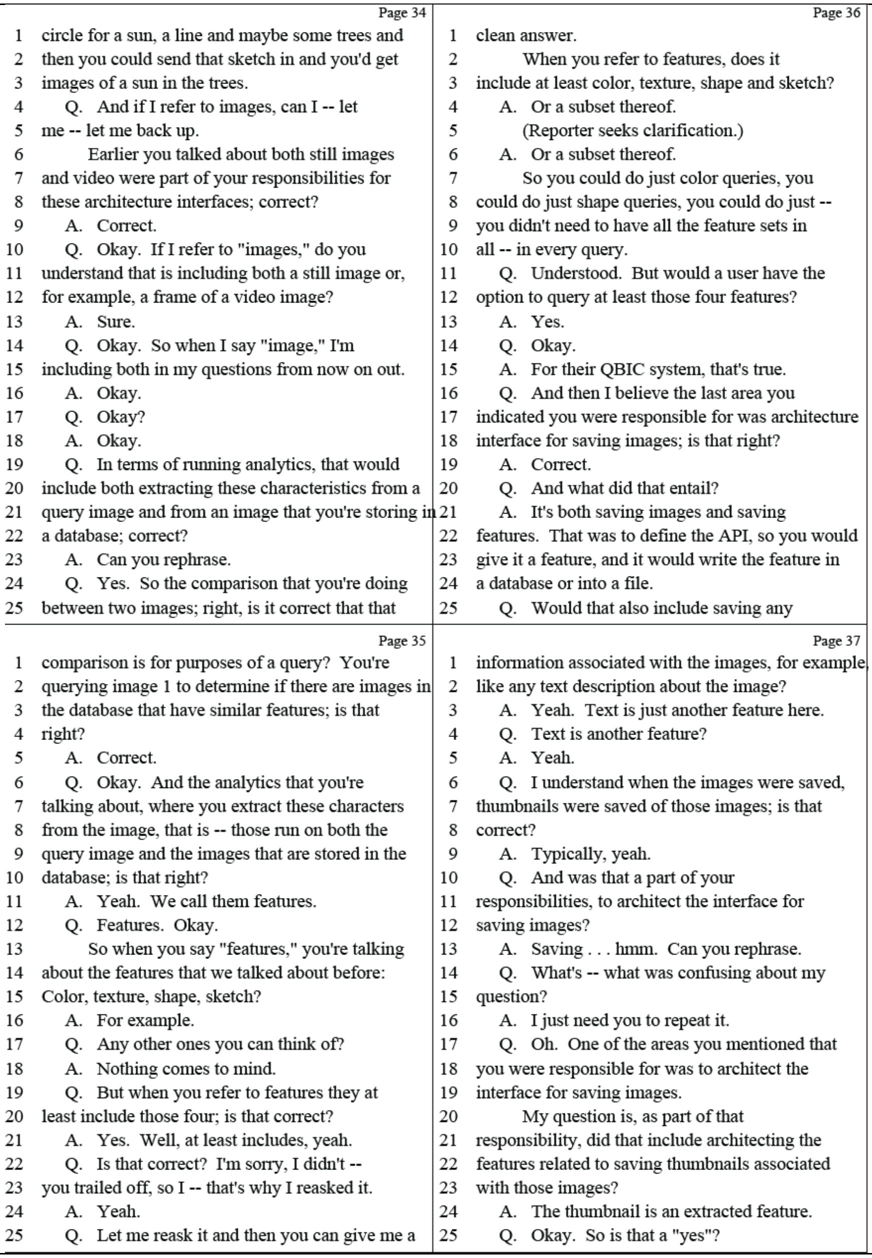

Page 34

circle for a sun, a line and maybe some trees and then you could send that sketch in and you'd get images of a sun in the trees.
Q. And if I refer to images, can I -- let me -- let me back up.
Earlier you talked about both still images and video were part of your responsibilities for these architecture interfaces; correct?
A. Correct.
Q. Okay. If I refer to "images," do you understand that is including both a still image or, for example, a frame of a video image?
A. Sure.
Q. Okay. So when I say "image," I'm including both in my questions from now on out.
A. Okay.
Q. Okay?
A. Okay.
Q. In terms of running analytics, that would include both extracting these characteristics from a query image and from an image that you're storing in a database; correct?
A. Can you rephrase.
Q. Yes. So the comparison that you're doing between two images; right, is it correct that that

Page 35

comparison is for purposes of a query? You're querying image 1 to determine if there are images in the database that have similar features; is that right?
A. Correct.
Q. Okay. And the analytics that you're talking about, where you extract these characters from the image, that is -- those run on both the query image and the images that are stored in the database; is that right?
A. Yeah. We call them features.
Q. Features. Okay.
So when you say "features," you're talking about the features that we talked about before: Color, texture, shape, sketch?
A. For example.
Q. Any other ones you can think of?
A. Nothing comes to mind.
Q. But when you refer to features they at least include those four; is that correct?
A. Yes. Well, at least includes, yeah.
Q. Is that correct? I'm sorry, I didn't -- you trailed off, so I -- that's why I reasked it.
A. Yeah.
Q. Let me reask it and then you can give me a

Page 36

clean answer.
When you refer to features, does it include at least color, texture, shape and sketch?
A. Or a subset thereof.
(Reporter seeks clarification.)
A. Or a subset thereof.
So you could do just color queries, you could do just shape queries, you could do just -- you didn't need to have all the feature sets in all -- in every query.
Q. Understood. But would a user have the option to query at least those four features?
A. Yes.
Q. Okay.
A. For their QBIC system, that's true.
Q. And then I believe the last area you indicated you were responsible for was architecture interface for saving images; is that right?
A. Correct.
Q. And what did that entail?
A. It's both saving images and saving features. That was to define the API, so you would give it a feature, and it would write the feature in a database or into a file.
Q. Would that also include saving any

Page 37

information associated with the images, for example, like any text description about the image?
A. Yeah. Text is just another feature here.
Q. Text is another feature?
A. Yeah.
Q. I understand when the images were saved, thumbnails were saved of those images; is that correct?
A. Typically, yeah.
Q. And was that a part of your responsibilities, to architect the interface for saving images?
A. Saving . . . hmm. Can you rephrase.
Q. What's -- what was confusing about my question?
A. I just need you to repeat it.
Q. Oh. One of the areas you mentioned that you were responsible for was to architect the interface for saving images.
My question is, as part of that responsibility, did that include architecting the features related to saving thumbnails associated with those images?
A. The thumbnail is an extracted feature.
Q. Okay. So is that a "yes"?

Element	'036 Claim Recitation	Exemplary Citations in Reference
		Id. at 52–53: 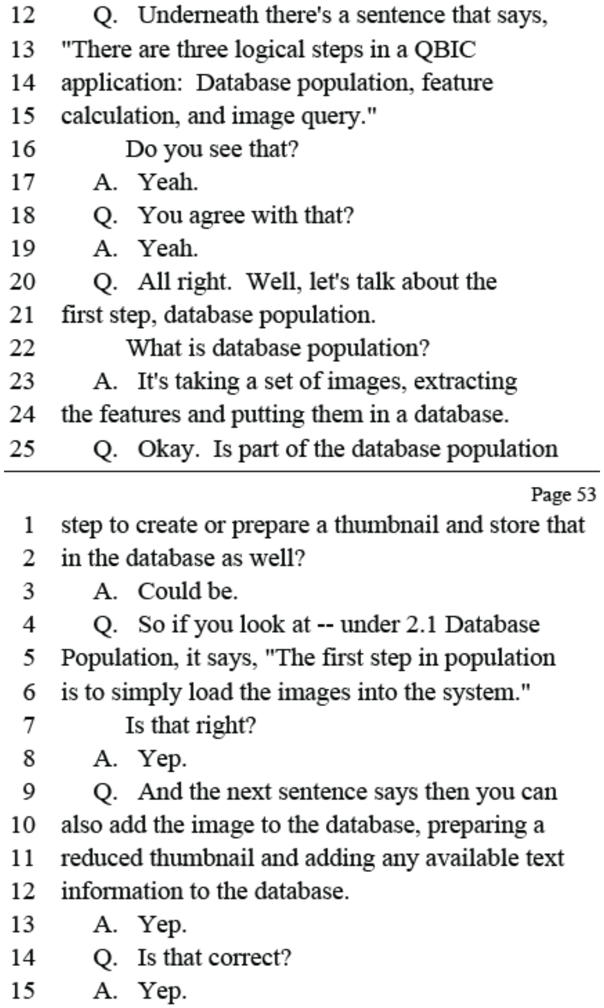 Q. Underneath there's a sentence that says, "There are three logical steps in a QBIC application: Database population, feature calculation, and image query." Do you see that? A. Yeah. Q. You agree with that? A. Yeah. Q. All right. Well, let's talk about the first step, database population. What is database population? A. It's taking a set of images, extracting the features and putting them in a database. Q. Okay. Is part of the database population Page 53 step to create or prepare a thumbnail and store that in the database as well? A. Could be. Q. So if you look at -- under 2.1 Database Population, it says, "The first step in population is to simply load the images into the system." Is that right? A. Yep. Q. And the next sentence says then you can also add the image to the database, preparing a reduced thumbnail and adding any available text information to the database. A. Yep. Q. Is that correct? A. Yep. *Id.* at 77–78:

Element	'036 Claim Recitation	Exemplary Citations in Reference
		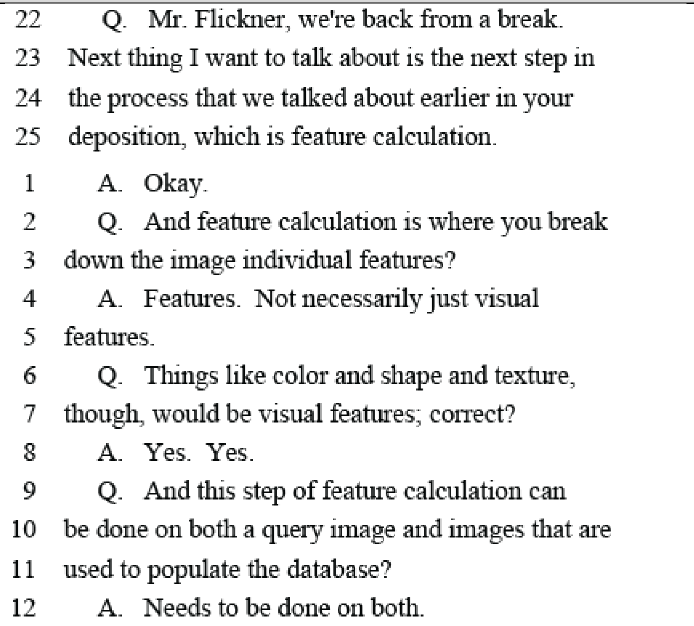 Q. Mr. Flickner, we're back from a break. Next thing I want to talk about is the next step in the process that we talked about earlier in your deposition, which is feature calculation. A. Okay. Q. And feature calculation is where you break down the image individual features? A. Features. Not necessarily just visual features. Q. Things like color and shape and texture, though, would be visual features; correct? A. Yes. Yes. Q. And this step of feature calculation can be done on both a query image and images that are used to populate the database? A. Needs to be done on both. *Id.* at 85–99: 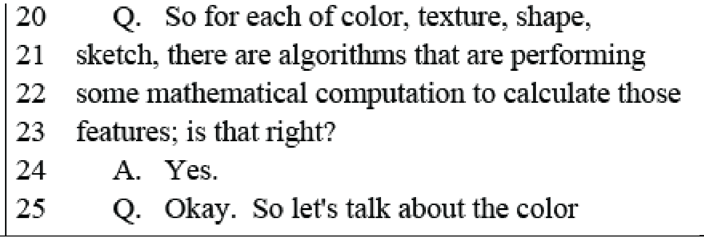 Q. So for each of color, texture, shape, sketch, there are algorithms that are performing some mathematical computation to calculate those features; is that right? A. Yes. Q. Okay. So let's talk about the color

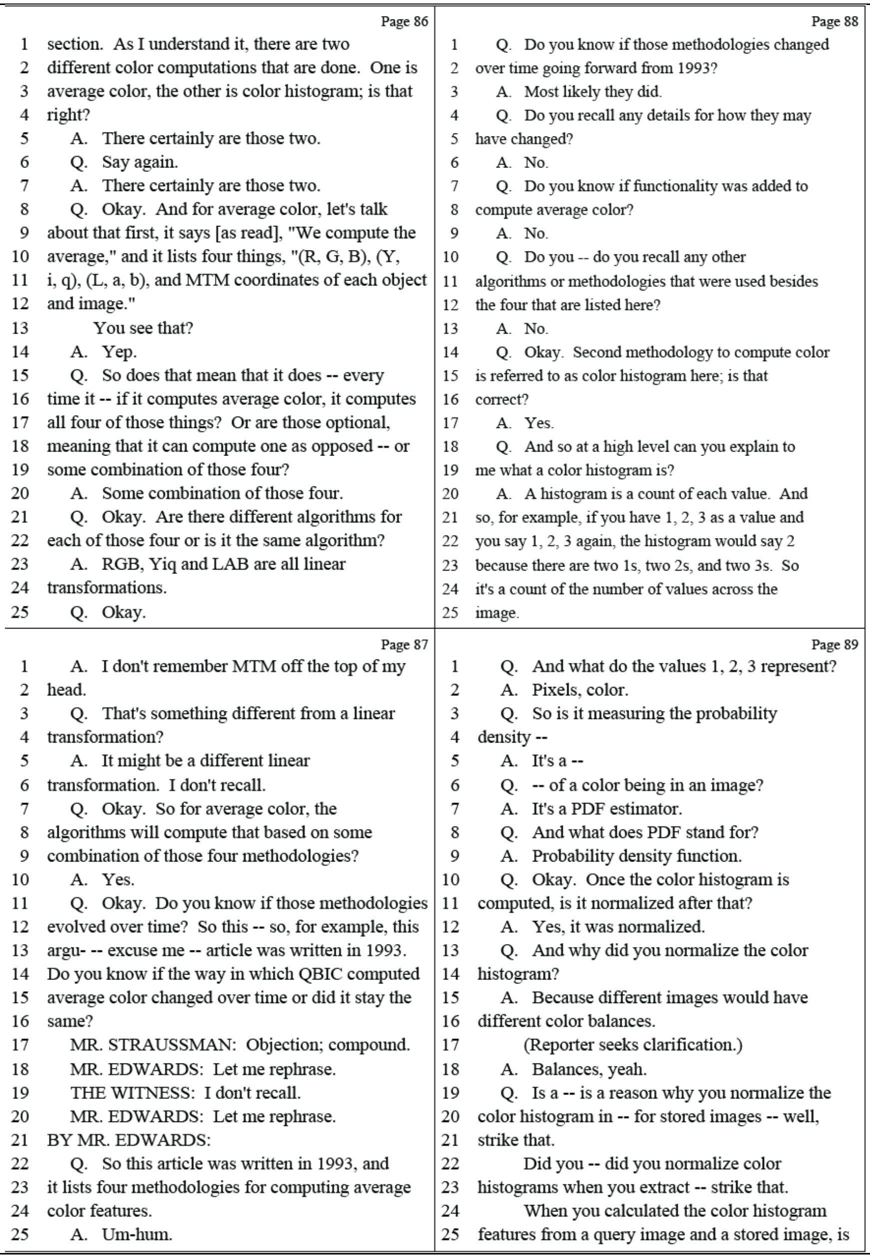

Page 86

section. As I understand it, there are two different color computations that are done. One is average color, the other is color histogram; is that right?
A. There certainly are those two.
Q. Say again.
A. There certainly are those two.
Q. Okay. And for average color, let's talk about that first, it says [as read], "We compute the average," and it lists four things, "(R, G, B), (Y, i, q), (L, a, b), and MTM coordinates of each object and image."
You see that?
A. Yep.
Q. So does that mean that it does -- every time it -- if it computes average color, it computes all four of those things? Or are those optional, meaning that it can compute one as opposed -- or some combination of those four?
A. Some combination of those four.
Q. Okay. Are there different algorithms for each of those four or is it the same algorithm?
A. RGB, Yiq and LAB are all linear transformations.
Q. Okay.

Page 87

A. I don't remember MTM off the top of my head.
Q. That's something different from a linear transformation?
A. It might be a different linear transformation. I don't recall.
Q. Okay. So for average color, the algorithms will compute that based on some combination of those four methodologies?
A. Yes.
Q. Okay. Do you know if those methodologies evolved over time? So this -- so, for example, this argu- -- excuse me -- article was written in 1993. Do you know if the way in which QBIC computed average color changed over time or did it stay the same?
MR. STRAUSSMAN: Objection; compound.
MR. EDWARDS: Let me rephrase.
THE WITNESS: I don't recall.
MR. EDWARDS: Let me rephrase.
BY MR. EDWARDS:
Q. So this article was written in 1993, and it lists four methodologies for computing average color features.
A. Um-hum.

Page 88

Q. Do you know if those methodologies changed over time going forward from 1993?
A. Most likely they did.
Q. Do you recall any details for how they may have changed?
A. No.
Q. Do you know if functionality was added to compute average color?
A. No.
Q. Do you -- do you recall any other algorithms or methodologies that were used besides the four that are listed here?
A. No.
Q. Okay. Second methodology to compute color is referred to as color histogram here; is that correct?
A. Yes.
Q. And so at a high level can you explain to me what a color histogram is?
A. A histogram is a count of each value. And so, for example, if you have 1, 2, 3 as a value and you say 1, 2, 3 again, the histogram would say 2 because there are two 1s, two 2s, and two 3s. So it's a count of the number of values across the image.

Page 89

Q. And what do the values 1, 2, 3 represent?
A. Pixels, color.
Q. So is it measuring the probability density --
A. It's a --
Q. -- of a color being in an image?
A. It's a PDF estimator.
Q. And what does PDF stand for?
A. Probability density function.
Q. Okay. Once the color histogram is computed, is it normalized after that?
A. Yes, it was normalized.
Q. And why did you normalize the color histogram?
A. Because different images would have different color balances.
(Reporter seeks clarification.)
A. Balances, yeah.
Q. Is a -- is a reason why you normalize the color histogram in -- for stored images -- well, strike that.
Did you -- did you normalize color histograms when you extract -- strike that.
When you calculated the color histogram features from a query image and a stored image, is

		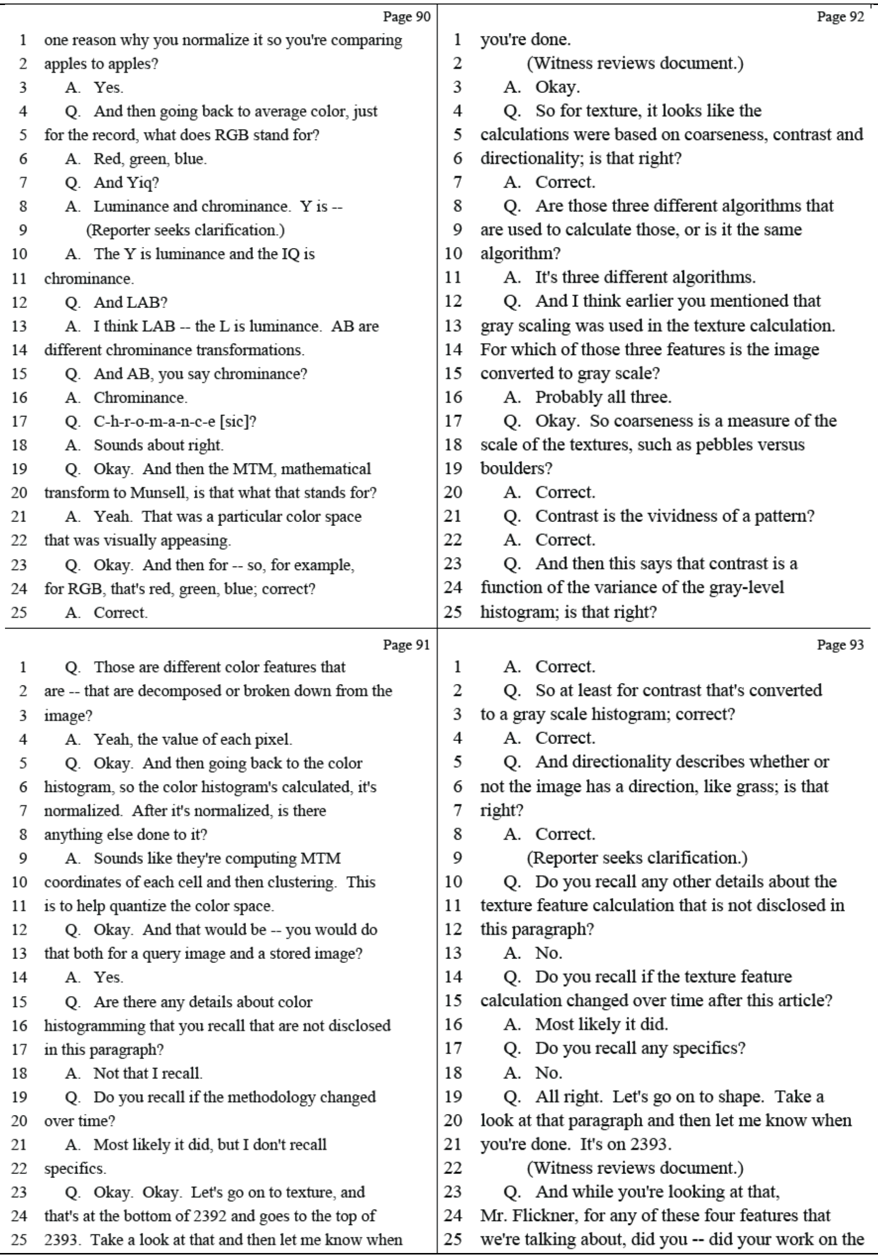

Page 90

one reason why you normalize it so you're comparing apples to apples?
A. Yes.
Q. And then going back to average color, just for the record, what does RGB stand for?
A. Red, green, blue.
Q. And Yiq?
A. Luminance and chrominance. Y is --
(Reporter seeks clarification.)
A. The Y is luminance and the IQ is chrominance.
Q. And LAB?
A. I think LAB -- the L is luminance. AB are different chrominance transformations.
Q. And AB, you say chrominance?
A. Chrominance.
Q. C-h-r-o-m-a-n-c-e [sic]?
A. Sounds about right.
Q. Okay. And then the MTM, mathematical transform to Munsell, is that what that stands for?
A. Yeah. That was a particular color space that was visually appeasing.
Q. Okay. And then for -- so, for example, for RGB, that's red, green, blue; correct?
A. Correct.

Page 91

Q. Those are different color features that are -- that are decomposed or broken down from the image?
A. Yeah, the value of each pixel.
Q. Okay. And then going back to the color histogram, so the color histogram's calculated, it's normalized. After it's normalized, is there anything else done to it?
A. Sounds like they're computing MTM coordinates of each cell and then clustering. This is to help quantize the color space.
Q. Okay. And that would be -- you would do that both for a query image and a stored image?
A. Yes.
Q. Are there any details about color histogramming that you recall that are not disclosed in this paragraph?
A. Not that I recall.
Q. Do you recall if the methodology changed over time?
A. Most likely it did, but I don't recall specifics.
Q. Okay. Okay. Let's go on to texture, and that's at the bottom of 2392 and goes to the top of 2393. Take a look at that and then let me know when

Page 92

you're done.
(Witness reviews document.)
A. Okay.
Q. So for texture, it looks like the calculations were based on coarseness, contrast and directionality; is that right?
A. Correct.
Q. Are those three different algorithms that are used to calculate those, or is it the same algorithm?
A. It's three different algorithms.
Q. And I think earlier you mentioned that gray scaling was used in the texture calculation. For which of those three features is the image converted to gray scale?
A. Probably all three.
Q. Okay. So coarseness is a measure of the scale of the textures, such as pebbles versus boulders?
A. Correct.
Q. Contrast is the vividness of a pattern?
A. Correct.
Q. And then this says that contrast is a function of the variance of the gray-level histogram; is that right?

Page 93

A. Correct.
Q. So at least for contrast that's converted to a gray scale histogram; correct?
A. Correct.
Q. And directionality describes whether or not the image has a direction, like grass; is that right?
A. Correct.
(Reporter seeks clarification.)
Q. Do you recall any other details about the texture feature calculation that is not disclosed in this paragraph?
A. No.
Q. Do you recall if the texture feature calculation changed over time after this article?
A. Most likely it did.
Q. Do you recall any specifics?
A. No.
Q. All right. Let's go on to shape. Take a look at that paragraph and then let me know when you're done. It's on 2393.
(Witness reviews document.)
Q. And while you're looking at that, Mr. Flickner, for any of these four features that we're talking about, did you -- did your work on the

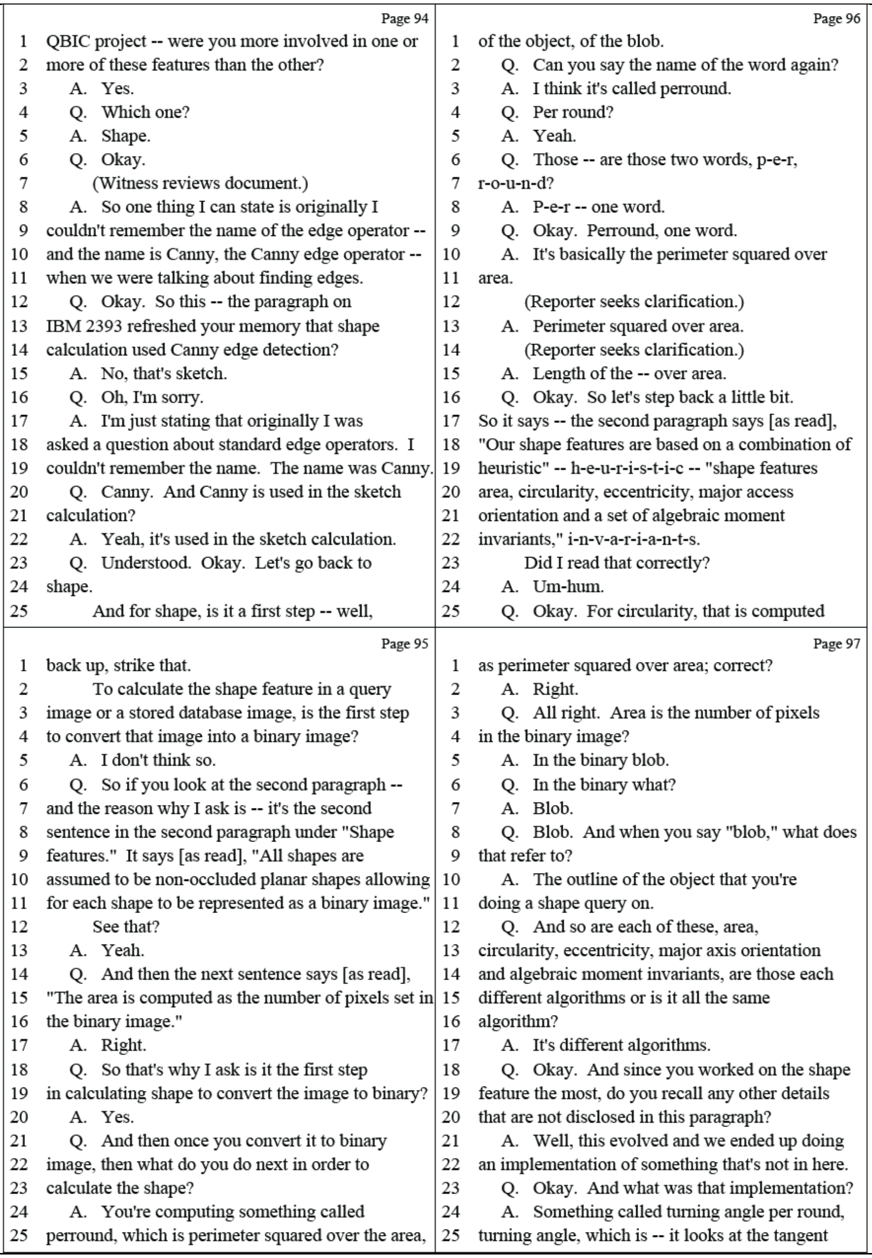

Page 94

QBIC project -- were you more involved in one or more of these features than the other?
A. Yes.
Q. Which one?
A. Shape.
Q. Okay.
(Witness reviews document.)
A. So one thing I can state is originally I couldn't remember the name of the edge operator -- and the name is Canny, the Canny edge operator -- when we were talking about finding edges.
Q. Okay. So this -- the paragraph on IBM 2393 refreshed your memory that shape calculation used Canny edge detection?
A. No, that's sketch.
Q. Oh, I'm sorry.
A. I'm just stating that originally I was asked a question about standard edge operators. I couldn't remember the name. The name was Canny.
Q. Canny. And Canny is used in the sketch calculation?
A. Yeah, it's used in the sketch calculation.
Q. Understood. Okay. Let's go back to shape.
And for shape, is it a first step -- well,

Page 95

back up, strike that.
To calculate the shape feature in a query image or a stored database image, is the first step to convert that image into a binary image?
A. I don't think so.
Q. So if you look at the second paragraph -- and the reason why I ask is -- it's the second sentence in the second paragraph under "Shape features." It says [as read], "All shapes are assumed to be non-occluded planar shapes allowing for each shape to be represented as a binary image."
See that?
A. Yeah.
Q. And then the next sentence says [as read], "The area is computed as the number of pixels set in the binary image."
A. Right.
Q. So that's why I ask is it the first step in calculating shape to convert the image to binary?
A. Yes.
Q. And then once you convert it to binary image, then what do you do next in order to calculate the shape?
A. You're computing something called perround, which is perimeter squared over the area,

Page 96

of the object, of the blob.
Q. Can you say the name of the word again?
A. I think it's called perround.
Q. Per round?
A. Yeah.
Q. Those -- are those two words, p-e-r, r-o-u-n-d?
A. P-e-r -- one word.
Q. Okay. Perround, one word.
A. It's basically the perimeter squared over area.
(Reporter seeks clarification.)
A. Perimeter squared over area.
(Reporter seeks clarification.)
A. Length of the -- over area.
Q. Okay. So let's step back a little bit. So it says -- the second paragraph says [as read], "Our shape features are based on a combination of heuristic" -- h-e-u-r-i-s-t-i-c -- "shape features area, circularity, eccentricity, major access orientation and a set of algebraic moment invariants," i-n-v-a-r-i-a-n-t-s.
Did I read that correctly?
A. Um-hum.
Q. Okay. For circularity, that is computed

Page 97

as perimeter squared over area; correct?
A. Right.
Q. All right. Area is the number of pixels in the binary image?
A. In the binary blob.
Q. In the binary what?
A. Blob.
Q. Blob. And when you say "blob," what does that refer to?
A. The outline of the object that you're doing a shape query on.
Q. And so are each of these, area, circularity, eccentricity, major axis orientation and algebraic moment invariants, are those each different algorithms or is it all the same algorithm?
A. It's different algorithms.
Q. Okay. And since you worked on the shape feature the most, do you recall any other details that are not disclosed in this paragraph?
A. Well, this evolved and we ended up doing an implementation of something that's not in here.
Q. Okay. And what was that implementation?
A. Something called turning angle per round, turning angle, which is -- it looks at the tangent

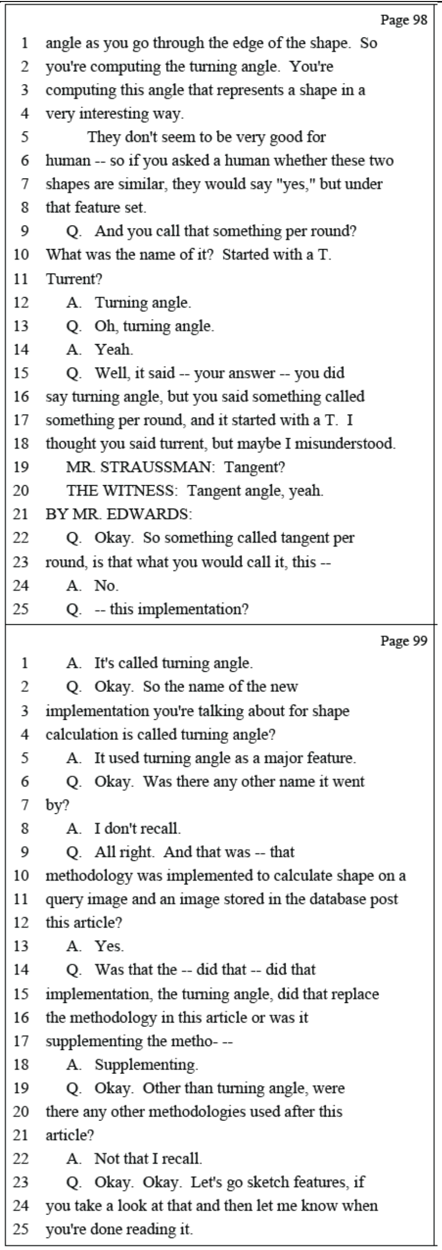

Page 98

angle as you go through the edge of the shape. So you're computing the turning angle. You're computing this angle that represents a shape in a very interesting way.

They don't seem to be very good for human -- so if you asked a human whether these two shapes are similar, they would say "yes," but under that feature set.

Q. And you call that something per round? What was the name of it? Started with a T. Turrent?

A. Turning angle.

Q. Oh, turning angle.

A. Yeah.

Q. Well, it said -- your answer -- you did say turning angle, but you said something called something per round, and it started with a T. I thought you said turrent, but maybe I misunderstood.

MR. STRAUSSMAN: Tangent?

THE WITNESS: Tangent angle, yeah.

BY MR. EDWARDS:

Q. Okay. So something called tangent per round, is that what you would call it, this --

A. No.

Q. -- this implementation?

Page 99

A. It's called turning angle.

Q. Okay. So the name of the new implementation you're talking about for shape calculation is called turning angle?

A. It used turning angle as a major feature.

Q. Okay. Was there any other name it went by?

A. I don't recall.

Q. All right. And that was -- that methodology was implemented to calculate shape on a query image and an image stored in the database post this article?

A. Yes.

Q. Was that the -- did that -- did that implementation, the turning angle, did that replace the methodology in this article or was it supplementing the metho- --

A. Supplementing.

Q. Okay. Other than turning angle, were there any other methodologies used after this article?

A. Not that I recall.

Q. Okay. Okay. Let's go sketch features, if you take a look at that and then let me know when you're done reading it.

Element	’036 Claim Recitation	Exemplary Citations in Reference
		Id. at 119–128 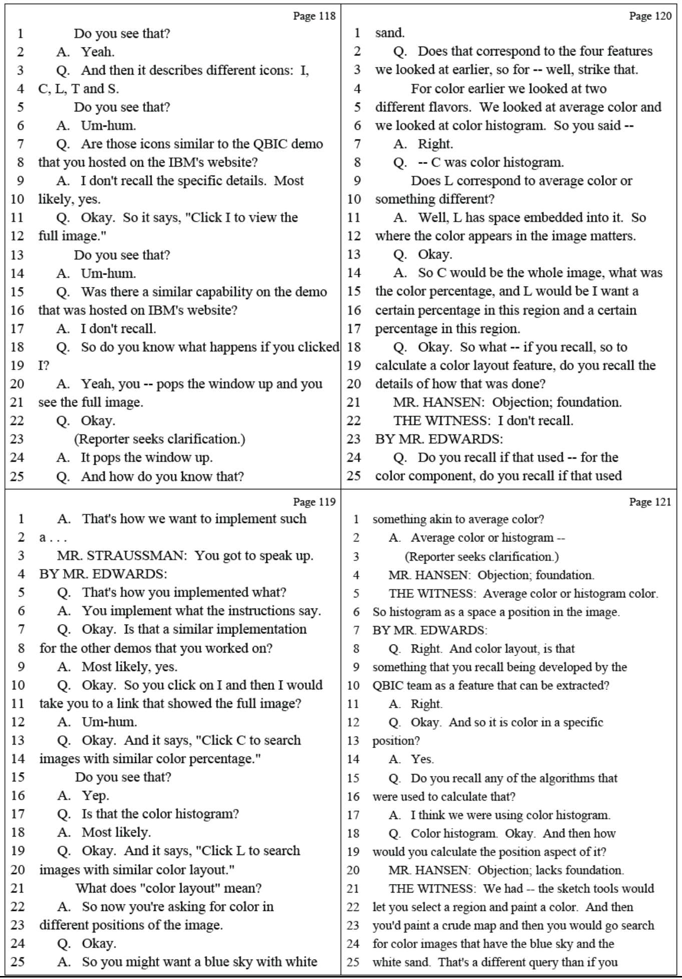 Page 118 Do you see that? A. Yeah. Q. And then it describes different icons: I, C, L, T and S. Do you see that? A. Um-hum. Q. Are those icons similar to the QBIC demo that you hosted on the IBM's website? A. I don't recall the specific details. Most likely, yes. Q. Okay. So it says, "Click I to view the full image." Do you see that? A. Um-hum. Q. Was there a similar capability on the demo that was hosted on IBM's website? A. I don't recall. Q. So do you know what happens if you clicked I? A. Yeah, you -- pops the window up and you see the full image. Q. Okay. (Reporter seeks clarification.) A. It pops the window up. Q. And how do you know that? Page 119 A. That's how we want to implement such a . . . MR. STRAUSSMAN: You got to speak up. BY MR. EDWARDS: Q. That's how you implemented what? A. You implement what the instructions say. Q. Okay. Is that a similar implementation for the other demos that you worked on? A. Most likely, yes. Q. Okay. So you click on I and then I would take you to a link that showed the full image? A. Um-hum. Q. Okay. And it says, "Click C to search images with similar color percentage." Do you see that? A. Yep. Q. Is that the color histogram? A. Most likely. Q. Okay. And it says, "Click L to search images with similar color layout." What does "color layout" mean? A. So now you're asking for color in different positions of the image. Q. Okay. A. So you might want a blue sky with white Page 120 sand. Q. Does that correspond to the four features we looked at earlier, so for -- well, strike that. For color earlier we looked at two different flavors. We looked at average color and we looked at color histogram. So you said -- A. Right. Q. -- C was color histogram. Does L correspond to average color or something different? A. Well, L has space embedded into it. So where the color appears in the image matters. Q. Okay. A. So C would be the whole image, what was the color percentage, and L would be I want a certain percentage in this region and a certain percentage in this region. Q. Okay. So what -- if you recall, so to calculate a color layout feature, do you recall the details of how that was done? MR. HANSEN: Objection; foundation. THE WITNESS: I don't recall. BY MR. EDWARDS: Q. Do you recall if that used -- for the color component, do you recall if that used Page 121 something akin to average color? A. Average color or histogram -- (Reporter seeks clarification.) MR. HANSEN: Objection; foundation. THE WITNESS: Average color or histogram color. So histogram as a space a position in the image. BY MR. EDWARDS: Q. Right. And color layout, is that something that you recall being developed by the QBIC team as a feature that can be extracted? A. Right. Q. Okay. And so it is color in a specific position? A. Yes. Q. Do you recall any of the algorithms that were used to calculate that? A. I think we were using color histogram. Q. Color histogram. Okay. And then how would you calculate the position aspect of it? MR. HANSEN: Objection; lacks foundation. THE WITNESS: We had -- the sketch tools would let you select a region and paint a color. And then you'd paint a crude map and then you would go search for color images that have the blue sky and the white sand. That's a different query than if you

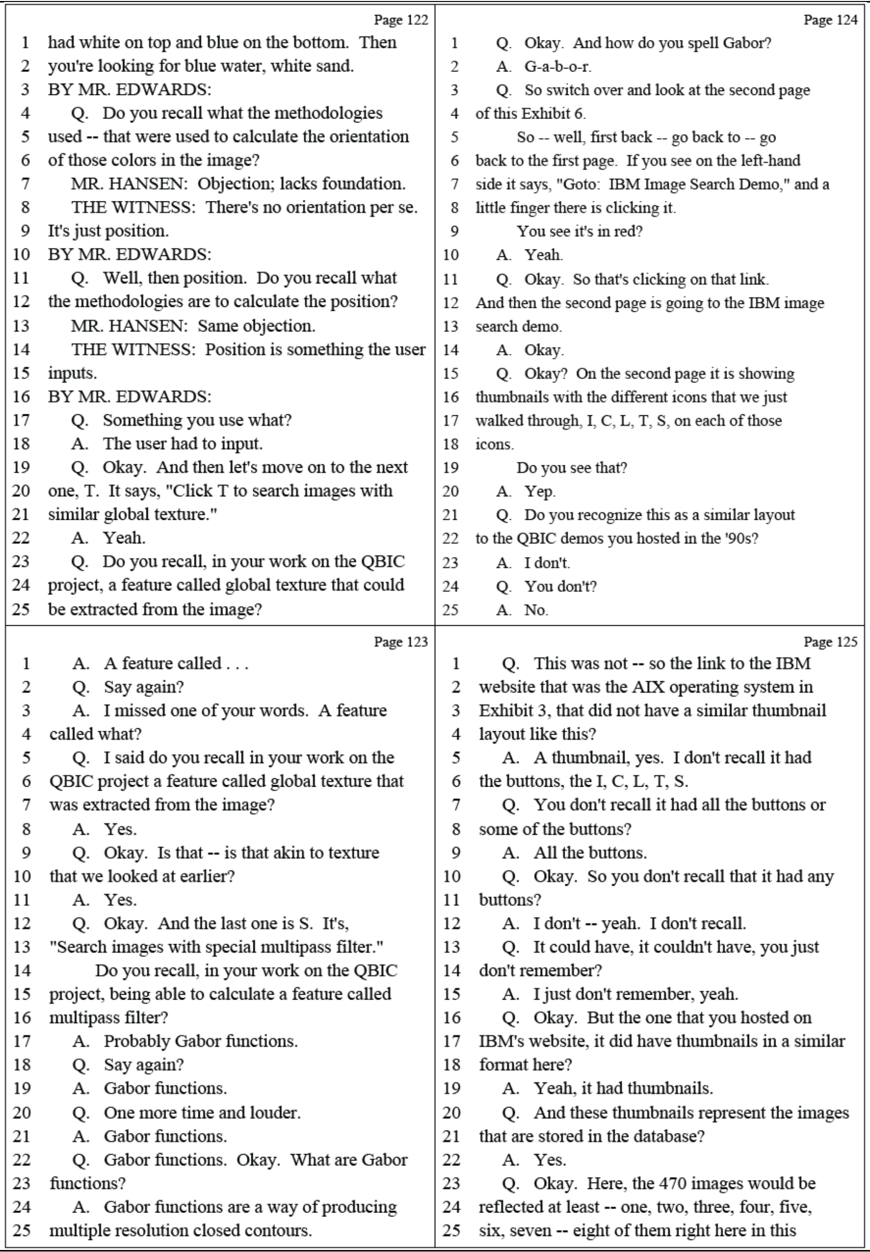

Page 122

had white on top and blue on the bottom. Then you're looking for blue water, white sand.
BY MR. EDWARDS:
Q. Do you recall what the methodologies used -- that were used to calculate the orientation of those colors in the image?
MR. HANSEN: Objection; lacks foundation.
THE WITNESS: There's no orientation per se. It's just position.
BY MR. EDWARDS:
Q. Well, then position. Do you recall what the methodologies are to calculate the position?
MR. HANSEN: Same objection.
THE WITNESS: Position is something the user inputs.
BY MR. EDWARDS:
Q. Something you use what?
A. The user had to input.
Q. Okay. And then let's move on to the next one, T. It says, "Click T to search images with similar global texture."
A. Yeah.
Q. Do you recall, in your work on the QBIC project, a feature called global texture that could be extracted from the image?

Page 123

A. A feature called . . .
Q. Say again?
A. I missed one of your words. A feature called what?
Q. I said do you recall in your work on the QBIC project a feature called global texture that was extracted from the image?
A. Yes.
Q. Okay. Is that -- is that akin to texture that we looked at earlier?
A. Yes.
Q. Okay. And the last one is S. It's, "Search images with special multipass filter."
Do you recall, in your work on the QBIC project, being able to calculate a feature called multipass filter?
A. Probably Gabor functions.
Q. Say again?
A. Gabor functions.
Q. One more time and louder.
A. Gabor functions.
Q. Gabor functions. Okay. What are Gabor functions?
A. Gabor functions are a way of producing multiple resolution closed contours.

Page 124

Q. Okay. And how do you spell Gabor?
A. G-a-b-o-r.
Q. So switch over and look at the second page of this Exhibit 6.
So -- well, first back -- go back to -- go back to the first page. If you see on the left-hand side it says, "Goto: IBM Image Search Demo," and a little finger there is clicking it.
You see it's in red?
A. Yeah.
Q. Okay. So that's clicking on that link. And then the second page is going to the IBM image search demo.
A. Okay.
Q. Okay? On the second page it is showing thumbnails with the different icons that we just walked through, I, C, L, T, S, on each of those icons.
Do you see that?
A. Yep.
Q. Do you recognize this as a similar layout to the QBIC demos you hosted in the '90s?
A. I don't.
Q. You don't?
A. No.

Page 125

Q. This was not -- so the link to the IBM website that was the AIX operating system in Exhibit 3, that did not have a similar thumbnail layout like this?
A. A thumbnail, yes. I don't recall it had the buttons, the I, C, L, T, S.
Q. You don't recall it had all the buttons or some of the buttons?
A. All the buttons.
Q. Okay. So you don't recall that it had any buttons?
A. I don't -- yeah. I don't recall.
Q. It could have, it couldn't have, you just don't remember?
A. I just don't remember, yeah.
Q. Okay. But the one that you hosted on IBM's website, it did have thumbnails in a similar format here?
A. Yeah, it had thumbnails.
Q. And these thumbnails represent the images that are stored in the database?
A. Yes.
Q. Okay. Here, the 470 images would be reflected at least -- one, two, three, four, five, six, seven -- eight of them right here in this

Element	'036 Claim Recitation	Exemplary Citations in Reference

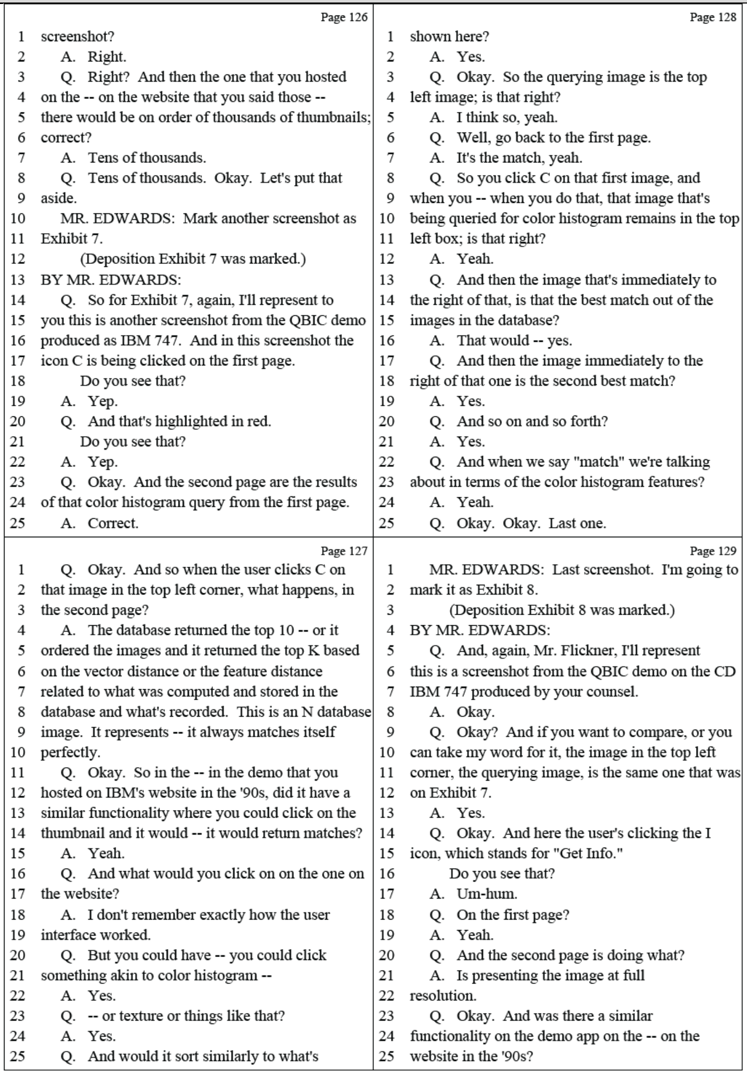

Page 126

screenshot?
A. Right.
Q. Right? And then the one that you hosted on the -- on the website that you said those -- there would be on order of thousands of thumbnails; correct?
A. Tens of thousands.
Q. Tens of thousands. Okay. Let's put that aside.
MR. EDWARDS: Mark another screenshot as Exhibit 7.
(Deposition Exhibit 7 was marked.)
BY MR. EDWARDS:
Q. So for Exhibit 7, again, I'll represent to you this is another screenshot from the QBIC demo produced as IBM 747. And in this screenshot the icon C is being clicked on the first page.
Do you see that?
A. Yep.
Q. And that's highlighted in red.
Do you see that?
A. Yep.
Q. Okay. And the second page are the results of that color histogram query from the first page.
A. Correct.

Page 127

Q. Okay. And so when the user clicks C on that image in the top left corner, what happens, in the second page?
A. The database returned the top 10 -- or it ordered the images and it returned the top K based on the vector distance or the feature distance related to what was computed and stored in the database and what's recorded. This is an N database image. It represents -- it always matches itself perfectly.
Q. Okay. So in the -- in the demo that you hosted on IBM's website in the '90s, did it have a similar functionality where you could click on the thumbnail and it would -- it would return matches?
A. Yeah.
Q. And what would you click on on the one on the website?
A. I don't remember exactly how the user interface worked.
Q. But you could have -- you could click something akin to color histogram --
A. Yes.
Q. -- or texture or things like that?
A. Yes.
Q. And would it sort similarly to what's

Page 128

shown here?
A. Yes.
Q. Okay. So the querying image is the top left image; is that right?
A. I think so, yeah.
Q. Well, go back to the first page.
A. It's the match, yeah.
Q. So you click C on that first image, and when you -- when you do that, that image that's being queried for color histogram remains in the top left box; is that right?
A. Yeah.
Q. And then the image that's immediately to the right of that, is that the best match out of the images in the database?
A. That would -- yes.
Q. And then the image immediately to the right of that one is the second best match?
A. Yes.
Q. And so on and so forth?
A. Yes.
Q. And when we say "match" we're talking about in terms of the color histogram features?
A. Yeah.
Q. Okay. Okay. Last one.

Page 129

MR. EDWARDS: Last screenshot. I'm going to mark it as Exhibit 8.
(Deposition Exhibit 8 was marked.)
BY MR. EDWARDS:
Q. And, again, Mr. Flickner, I'll represent this is a screenshot from the QBIC demo on the CD IBM 747 produced by your counsel.
A. Okay.
Q. Okay? And if you want to compare, or you can take my word for it, the image in the top left corner, the querying image, is the same one that was on Exhibit 7.
A. Yes.
Q. Okay. And here the user's clicking the I icon, which stands for "Get Info."
Do you see that?
A. Um-hum.
Q. On the first page?
A. Yeah.
Q. And the second page is doing what?
A. Is presenting the image at full resolution.
Q. Okay. And was there a similar functionality on the demo app on the -- on the website in the '90s?

Id. at 135 – 136:

]

Element	'036 Claim Recitation	Exemplary Citations in Reference
		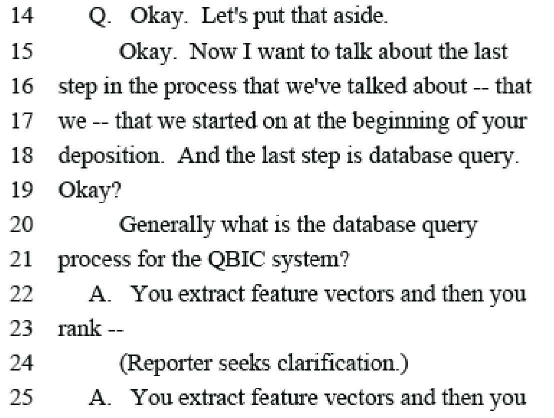 Q. Okay. Let's put that aside. Okay. Now I want to talk about the last step in the process that we've talked about -- that we -- that we started on at the beginning of your deposition. And the last step is database query. Okay? Generally what is the database query process for the QBIC system? A. You extract feature vectors and then you rank -- (Reporter seeks clarification.) A. You extract feature vectors and then you 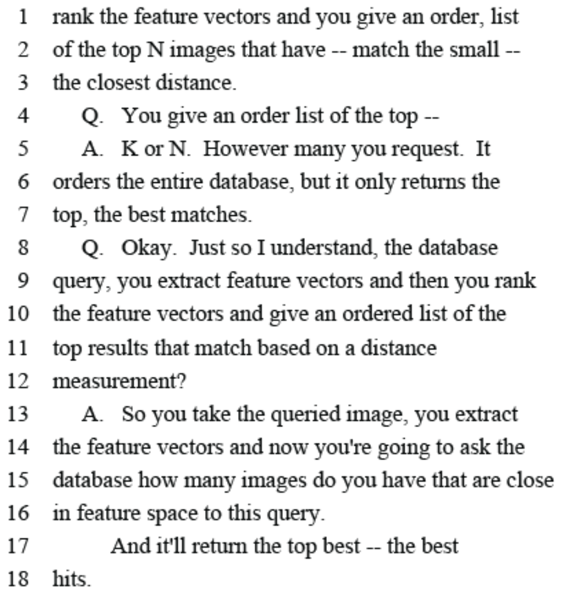 rank the feature vectors and you give an order, list of the top N images that have -- match the small -- the closest distance. Q. You give an order list of the top -- A. K or N. However many you request. It orders the entire database, but it only returns the top, the best matches. Q. Okay. Just so I understand, the database query, you extract feature vectors and then you rank the feature vectors and give an ordered list of the top results that match based on a distance measurement? A. So you take the queried image, you extract the feature vectors and now you're going to ask the database how many images do you have that are close in feature space to this query. And it'll return the top best -- the best hits. *Id.* at 139–165:

Element	'036 Claim Recitation	Exemplary Citations in Reference
		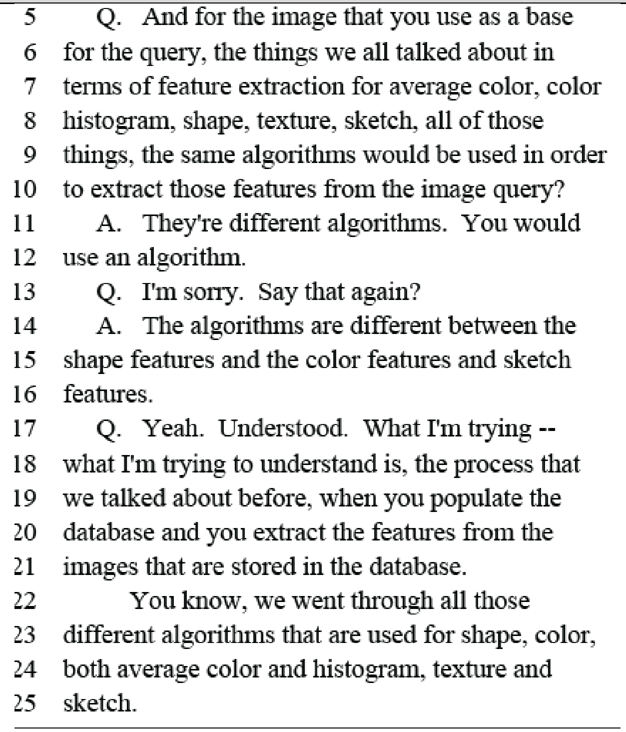 Q. And for the image that you use as a base for the query, the things we all talked about in terms of feature extraction for average color, color histogram, shape, texture, sketch, all of those things, the same algorithms would be used in order to extract those features from the image query? A. They're different algorithms. You would use an algorithm. Q. I'm sorry. Say that again? A. The algorithms are different between the shape features and the color features and sketch features. Q. Yeah. Understood. What I'm trying -- what I'm trying to understand is, the process that we talked about before, when you populate the database and you extract the features from the images that are stored in the database. You know, we went through all those different algorithms that are used for shape, color, both average color and histogram, texture and sketch.

Element	'036 Claim Recitation	Exemplary Citations in Reference
		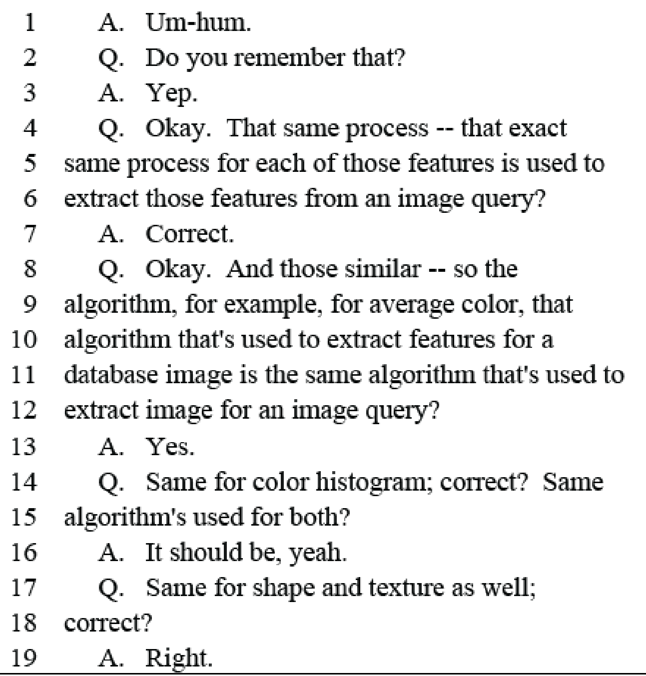 A. Um-hum. Q. Do you remember that? A. Yep. Q. Okay. That same process -- that exact same process for each of those features is used to extract those features from an image query? A. Correct. Q. Okay. And those similar -- so the algorithm, for example, for average color, that algorithm that's used to extract features for a database image is the same algorithm that's used to extract image for an image query? A. Yes. Q. Same for color histogram; correct? Same algorithm's used for both? A. It should be, yeah. Q. Same for shape and texture as well; correct? A. Right.

Element	'036 Claim Recitation	Exemplary Citations in Reference
		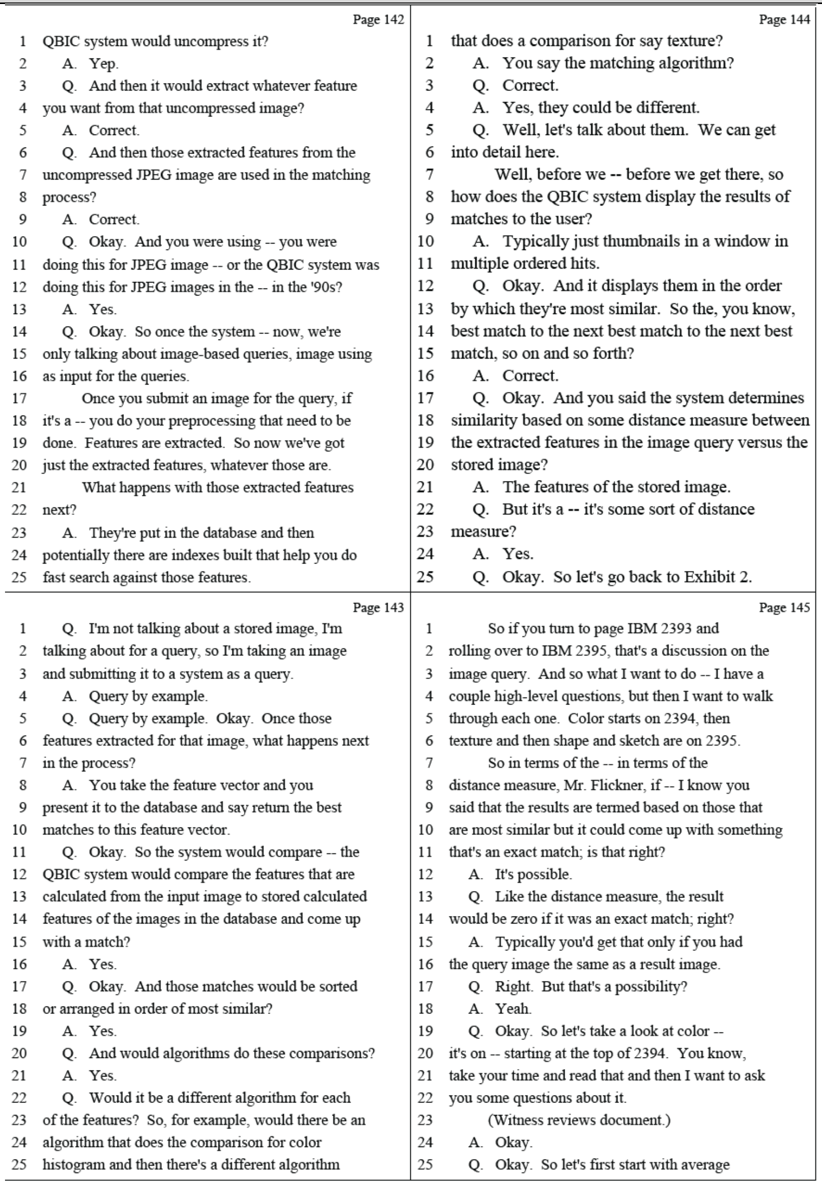 Page 142 QBIC system would uncompress it? A. Yep. Q. And then it would extract whatever feature you want from that uncompressed image? A. Correct. Q. And then those extracted features from the uncompressed JPEG image are used in the matching process? A. Correct. Q. Okay. And you were using -- you were doing this for JPEG image -- or the QBIC system was doing this for JPEG images in the -- in the '90s? A. Yes. Q. Okay. So once the system -- now, we're only talking about image-based queries, image using as input for the queries. Once you submit an image for the query, if it's a -- you do your preprocessing that need to be done. Features are extracted. So now we've got just the extracted features, whatever those are. What happens with those extracted features next? A. They're put in the database and then potentially there are indexes built that help you do fast search against those features. Page 143 Q. I'm not talking about a stored image, I'm talking about for a query, so I'm taking an image and submitting it to a system as a query. A. Query by example. Q. Query by example. Okay. Once those features extracted for that image, what happens next in the process? A. You take the feature vector and you present it to the database and say return the best matches to this feature vector. Q. Okay. So the system would compare -- the QBIC system would compare the features that are calculated from the input image to stored calculated features of the images in the database and come up with a match? A. Yes. Q. Okay. And those matches would be sorted or arranged in order of most similar? A. Yes. Q. And would algorithms do these comparisons? A. Yes. Q. Would it be a different algorithm for each of the features? So, for example, would there be an algorithm that does the comparison for color histogram and then there's a different algorithm Page 144 that does a comparison for say texture? A. You say the matching algorithm? Q. Correct. A. Yes, they could be different. Q. Well, let's talk about them. We can get into detail here. Well, before we -- before we get there, so how does the QBIC system display the results of matches to the user? A. Typically just thumbnails in a window in multiple ordered hits. Q. Okay. And it displays them in the order by which they're most similar. So the, you know, best match to the next best match to the next best match, so on and so forth? A. Correct. Q. Okay. And you said the system determines similarity based on some distance measure between the extracted features in the image query versus the stored image? A. The features of the stored image. Q. But it's a -- it's some sort of distance measure? A. Yes. Q. Okay. So let's go back to Exhibit 2. Page 145 So if you turn to page IBM 2393 and rolling over to IBM 2395, that's a discussion on the image query. And so what I want to do -- I have a couple high-level questions, but then I want to walk through each one. Color starts on 2394, then texture and then shape and sketch are on 2395. So in terms of the -- in terms of the distance measure, Mr. Flickner, if -- I know you said that the results are termed based on those that are most similar but it could come up with something that's an exact match; is that right? A. It's possible. Q. Like the distance measure, the result would be zero if it was an exact match; right? A. Typically you'd get that only if you had the query image the same as a result image. Q. Right. But that's a possibility? A. Yeah. Q. Okay. So let's take a look at color -- it's on -- starting at the top of 2394. You know, take your time and read that and then I want to ask you some questions about it. (Witness reviews document.) A. Okay. Q. Okay. So let's first start with average

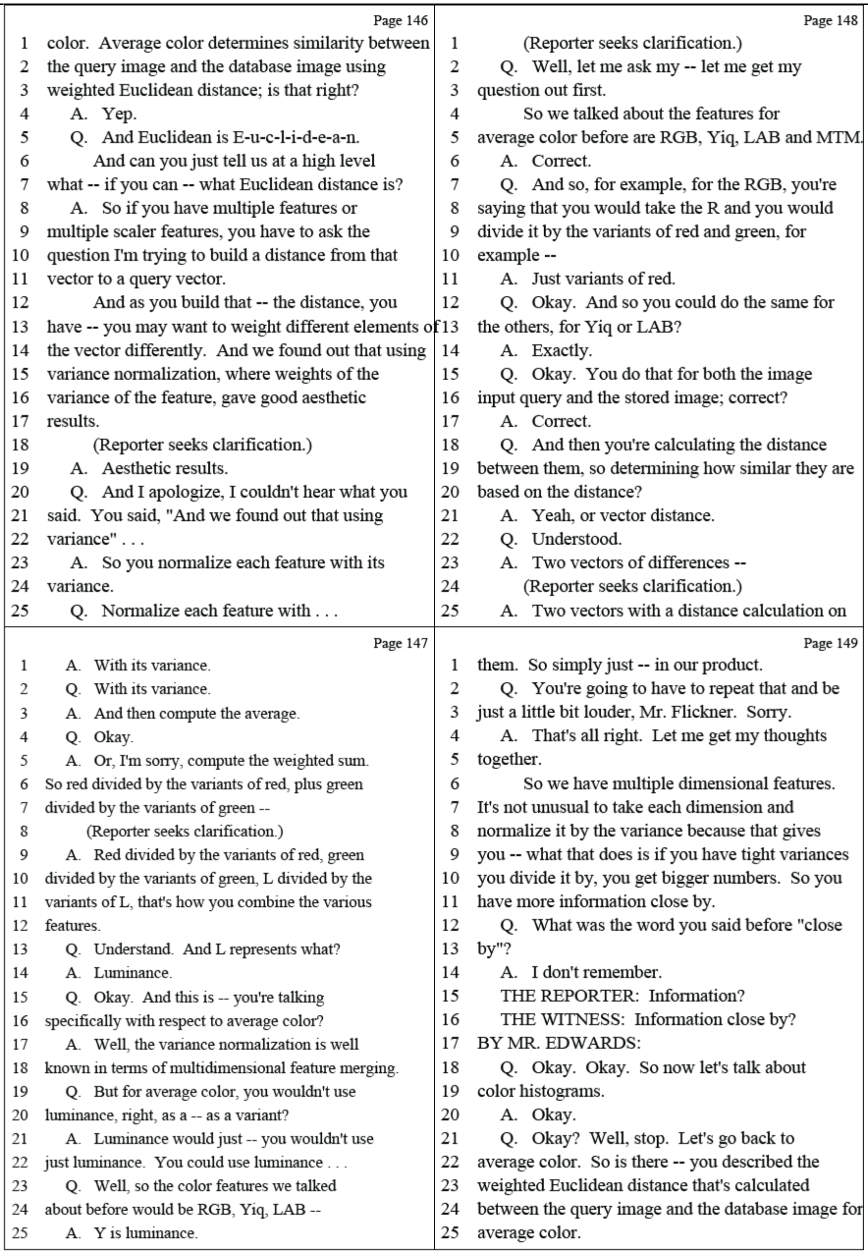

Page 146

color. Average color determines similarity between the query image and the database image using weighted Euclidean distance; is that right?

A. Yep.

Q. And Euclidean is E-u-c-l-i-d-e-a-n.

And can you just tell us at a high level what -- if you can -- what Euclidean distance is?

A. So if you have multiple features or multiple scaler features, you have to ask the question I'm trying to build a distance from that vector to a query vector.

And as you build that -- the distance, you have -- you may want to weight different elements of the vector differently. And we found out that using variance normalization, where weights of the variance of the feature, gave good aesthetic results.

(Reporter seeks clarification.)

A. Aesthetic results.

Q. And I apologize, I couldn't hear what you said. You said, "And we found out that using variance" . . .

A. So you normalize each feature with its variance.

Q. Normalize each feature with . . .

Page 147

A. With its variance.

Q. With its variance.

A. And then compute the average.

Q. Okay.

A. Or, I'm sorry, compute the weighted sum. So red divided by the variants of red, plus green divided by the variants of green --

(Reporter seeks clarification.)

A. Red divided by the variants of red, green divided by the variants of green, L divided by the variants of L, that's how you combine the various features.

Q. Understand. And L represents what?

A. Luminance.

Q. Okay. And this is -- you're talking specifically with respect to average color?

A. Well, the variance normalization is well known in terms of multidimensional feature merging.

Q. But for average color, you wouldn't use luminance, right, as a -- as a variant?

A. Luminance would just -- you wouldn't use just luminance. You could use luminance . . .

Q. Well, so the color features we talked about before would be RGB, Yiq, LAB --

A. Y is luminance.

Page 148

(Reporter seeks clarification.)

Q. Well, let me ask my -- let me get my question out first.

So we talked about the features for average color before are RGB, Yiq, LAB and MTM.

A. Correct.

Q. And so, for example, for the RGB, you're saying that you would take the R and you would divide it by the variants of red and green, for example --

A. Just variants of red.

Q. Okay. And so you could do the same for the others, for Yiq or LAB?

A. Exactly.

Q. Okay. You do that for both the image input query and the stored image; correct?

A. Correct.

Q. And then you're calculating the distance between them, so determining how similar they are based on the distance?

A. Yeah, or vector distance.

Q. Understood.

A. Two vectors of differences --

(Reporter seeks clarification.)

A. Two vectors with a distance calculation on

Page 149

them. So simply just -- in our product.

Q. You're going to have to repeat that and be just a little bit louder, Mr. Flickner. Sorry.

A. That's all right. Let me get my thoughts together.

So we have multiple dimensional features. It's not unusual to take each dimension and normalize it by the variance because that gives you -- what that does is if you have tight variances you divide it by, you get bigger numbers. So you have more information close by.

Q. What was the word you said before "close by"?

A. I don't remember.

THE REPORTER: Information?

THE WITNESS: Information close by?

BY MR. EDWARDS:

Q. Okay. Okay. So now let's talk about color histograms.

A. Okay.

Q. Okay? Well, stop. Let's go back to average color. So is there -- you described the weighted Euclidean distance that's calculated between the query image and the database image for average color.

		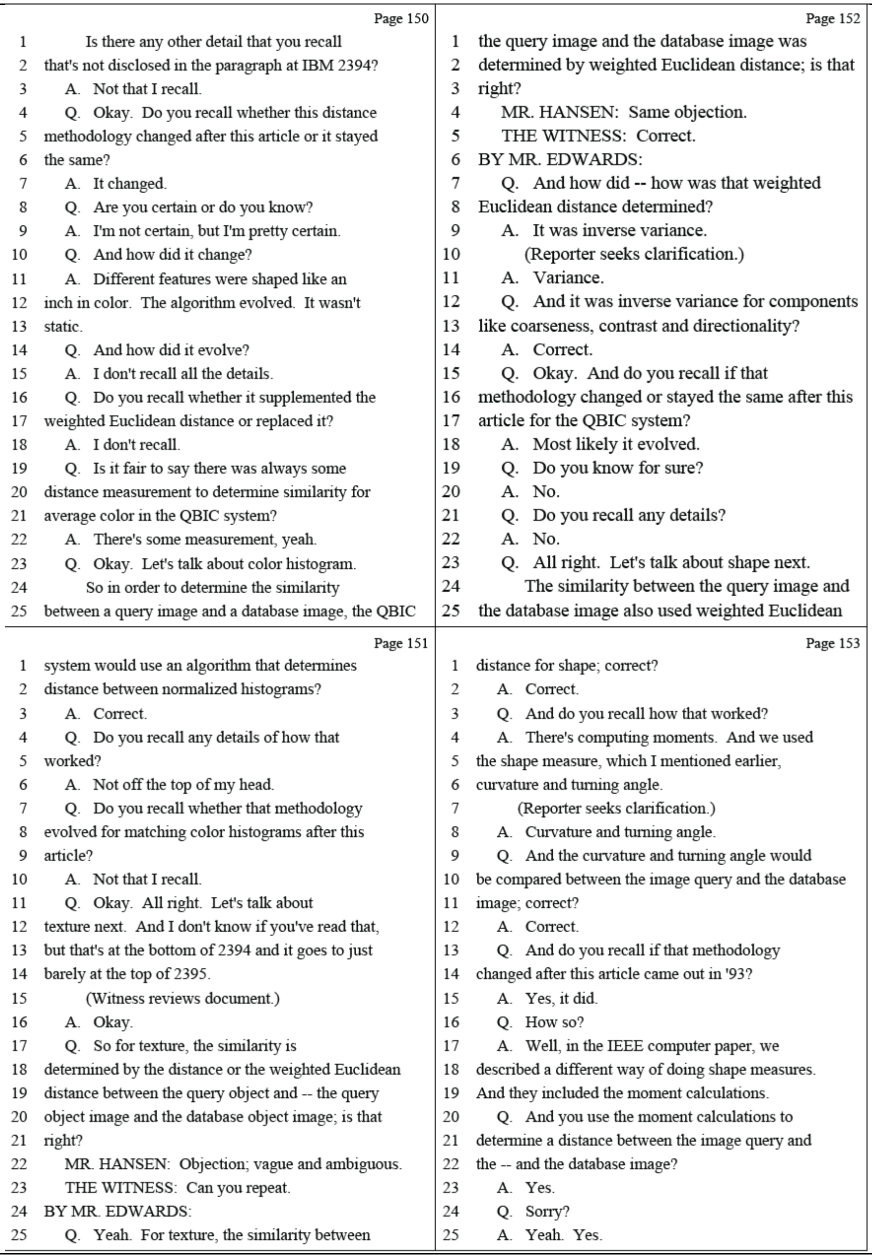

Page 150

Is there any other detail that you recall that's not disclosed in the paragraph at IBM 2394?

A. Not that I recall.

Q. Okay. Do you recall whether this distance methodology changed after this article or it stayed the same?

A. It changed.

Q. Are you certain or do you know?

A. I'm not certain, but I'm pretty certain.

Q. And how did it change?

A. Different features were shaped like an inch in color. The algorithm evolved. It wasn't static.

Q. And how did it evolve?

A. I don't recall all the details.

Q. Do you recall whether it supplemented the weighted Euclidean distance or replaced it?

A. I don't recall.

Q. Is it fair to say there was always some distance measurement to determine similarity for average color in the QBIC system?

A. There's some measurement, yeah.

Q. Okay. Let's talk about color histogram.

So in order to determine the similarity between a query image and a database image, the QBIC

Page 151

system would use an algorithm that determines distance between normalized histograms?

A. Correct.

Q. Do you recall any details of how that worked?

A. Not off the top of my head.

Q. Do you recall whether that methodology evolved for matching color histograms after this article?

A. Not that I recall.

Q. Okay. All right. Let's talk about texture next. And I don't know if you've read that, but that's at the bottom of 2394 and it goes to just barely at the top of 2395.

(Witness reviews document.)

A. Okay.

Q. So for texture, the similarity is determined by the distance or the weighted Euclidean distance between the query object and -- the query object image and the database object image; is that right?

MR. HANSEN: Objection; vague and ambiguous.

THE WITNESS: Can you repeat.

BY MR. EDWARDS:

Q. Yeah. For texture, the similarity between

Page 152

the query image and the database image was determined by weighted Euclidean distance; is that right?

MR. HANSEN: Same objection.

THE WITNESS: Correct.

BY MR. EDWARDS:

Q. And how did -- how was that weighted Euclidean distance determined?

A. It was inverse variance.

(Reporter seeks clarification.)

A. Variance.

Q. And it was inverse variance for components like coarseness, contrast and directionality?

A. Correct.

Q. Okay. And do you recall if that methodology changed or stayed the same after this article for the QBIC system?

A. Most likely it evolved.

Q. Do you know for sure?

A. No.

Q. Do you recall any details?

A. No.

Q. All right. Let's talk about shape next.

The similarity between the query image and the database image also used weighted Euclidean

Page 153

distance for shape; correct?

A. Correct.

Q. And do you recall how that worked?

A. There's computing moments. And we used the shape measure, which I mentioned earlier, curvature and turning angle.

(Reporter seeks clarification.)

A. Curvature and turning angle.

Q. And the curvature and turning angle would be compared between the image query and the database image; correct?

A. Correct.

Q. And do you recall if that methodology changed after this article came out in '93?

A. Yes, it did.

Q. How so?

A. Well, in the IEEE computer paper, we described a different way of doing shape measures. And they included the moment calculations.

Q. And you use the moment calculations to determine a distance between the image query and the -- and the database image?

A. Yes.

Q. Sorry?

A. Yeah. Yes.

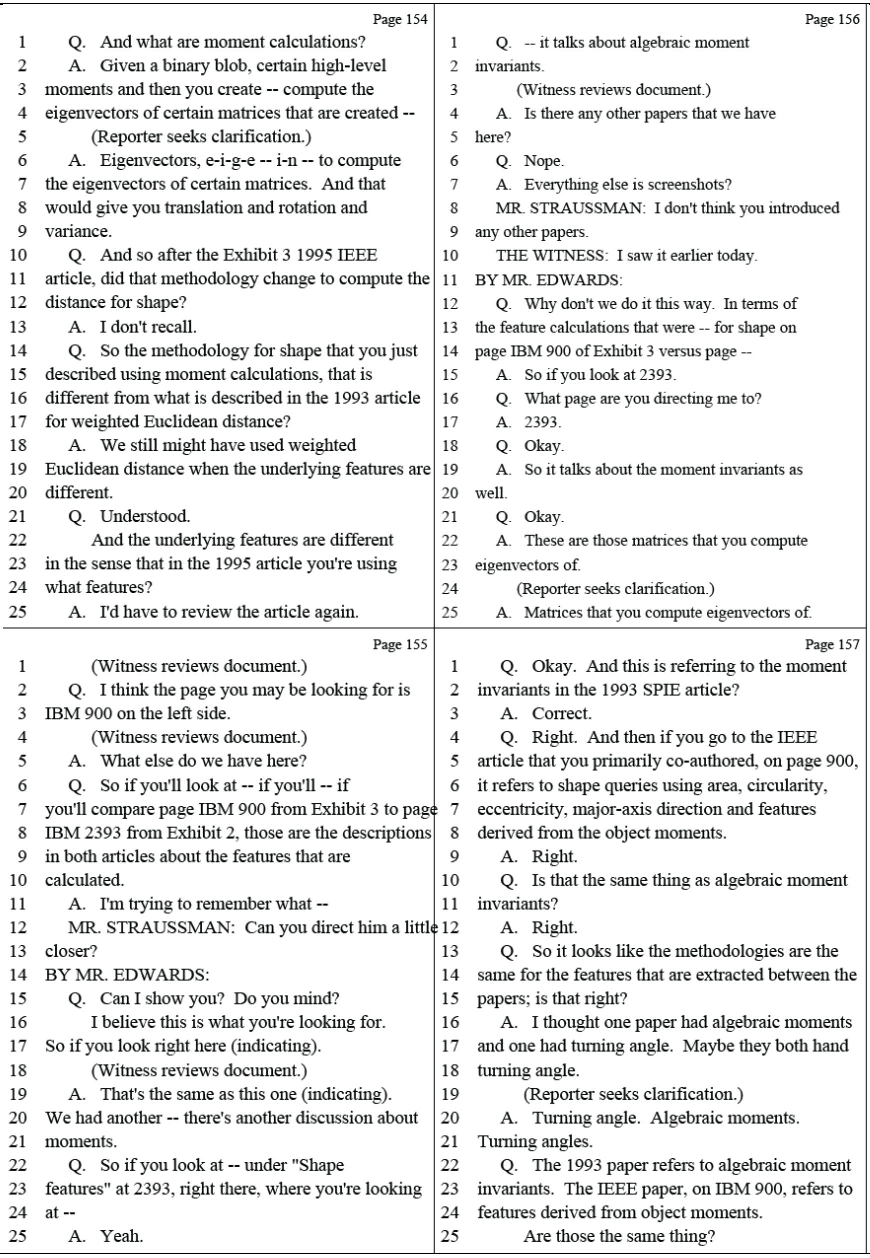

Page 154

Q. And what are moment calculations?
A. Given a binary blob, certain high-level moments and then you create -- compute the eigenvectors of certain matrices that are created --
(Reporter seeks clarification.)
A. Eigenvectors, e-i-g-e -- i-n -- to compute the eigenvectors of certain matrices. And that would give you translation and rotation and variance.
Q. And so after the Exhibit 3 1995 IEEE article, did that methodology change to compute the distance for shape?
A. I don't recall.
Q. So the methodology for shape that you just described using moment calculations, that is different from what is described in the 1993 article for weighted Euclidean distance?
A. We still might have used weighted Euclidean distance when the underlying features are different.
Q. Understood.
And the underlying features are different in the sense that in the 1995 article you're using what features?
A. I'd have to review the article again.

Page 155

(Witness reviews document.)
Q. I think the page you may be looking for is IBM 900 on the left side.
(Witness reviews document.)
A. What else do we have here?
Q. So if you'll look at -- if you'll -- if you'll compare page IBM 900 from Exhibit 3 to page IBM 2393 from Exhibit 2, those are the descriptions in both articles about the features that are calculated.
A. I'm trying to remember what --
MR. STRAUSSMAN: Can you direct him a little closer?
BY MR. EDWARDS:
Q. Can I show you? Do you mind?
I believe this is what you're looking for. So if you look right here (indicating).
(Witness reviews document.)
A. That's the same as this one (indicating). We had another -- there's another discussion about moments.
Q. So if you look at -- under "Shape features" at 2393, right there, where you're looking at --
A. Yeah.

Page 156

Q. -- it talks about algebraic moment invariants.
(Witness reviews document.)
A. Is there any other papers that we have here?
Q. Nope.
A. Everything else is screenshots?
MR. STRAUSSMAN: I don't think you introduced any other papers.
THE WITNESS: I saw it earlier today.
BY MR. EDWARDS:
Q. Why don't we do it this way. In terms of the feature calculations that were -- for shape on page IBM 900 of Exhibit 3 versus page --
A. So if you look at 2393.
Q. What page are you directing me to?
A. 2393.
Q. Okay.
A. So it talks about the moment invariants as well.
Q. Okay.
A. These are those matrices that you compute eigenvectors of.
(Reporter seeks clarification.)
A. Matrices that you compute eigenvectors of.

Page 157

Q. Okay. And this is referring to the moment invariants in the 1993 SPIE article?
A. Correct.
Q. Right. And then if you go to the IEEE article that you primarily co-authored, on page 900, it refers to shape queries using area, circularity, eccentricity, major-axis direction and features derived from the object moments.
A. Right.
Q. Is that the same thing as algebraic moment invariants?
A. Right.
Q. So it looks like the methodologies are the same for the features that are extracted between the papers; is that right?
A. I thought one paper had algebraic moments and one had turning angle. Maybe they both hand turning angle.
(Reporter seeks clarification.)
A. Turning angle. Algebraic moments. Turning angles.
Q. The 1993 paper refers to algebraic moment invariants. The IEEE paper, on IBM 900, refers to features derived from object moments.
Are those the same thing?

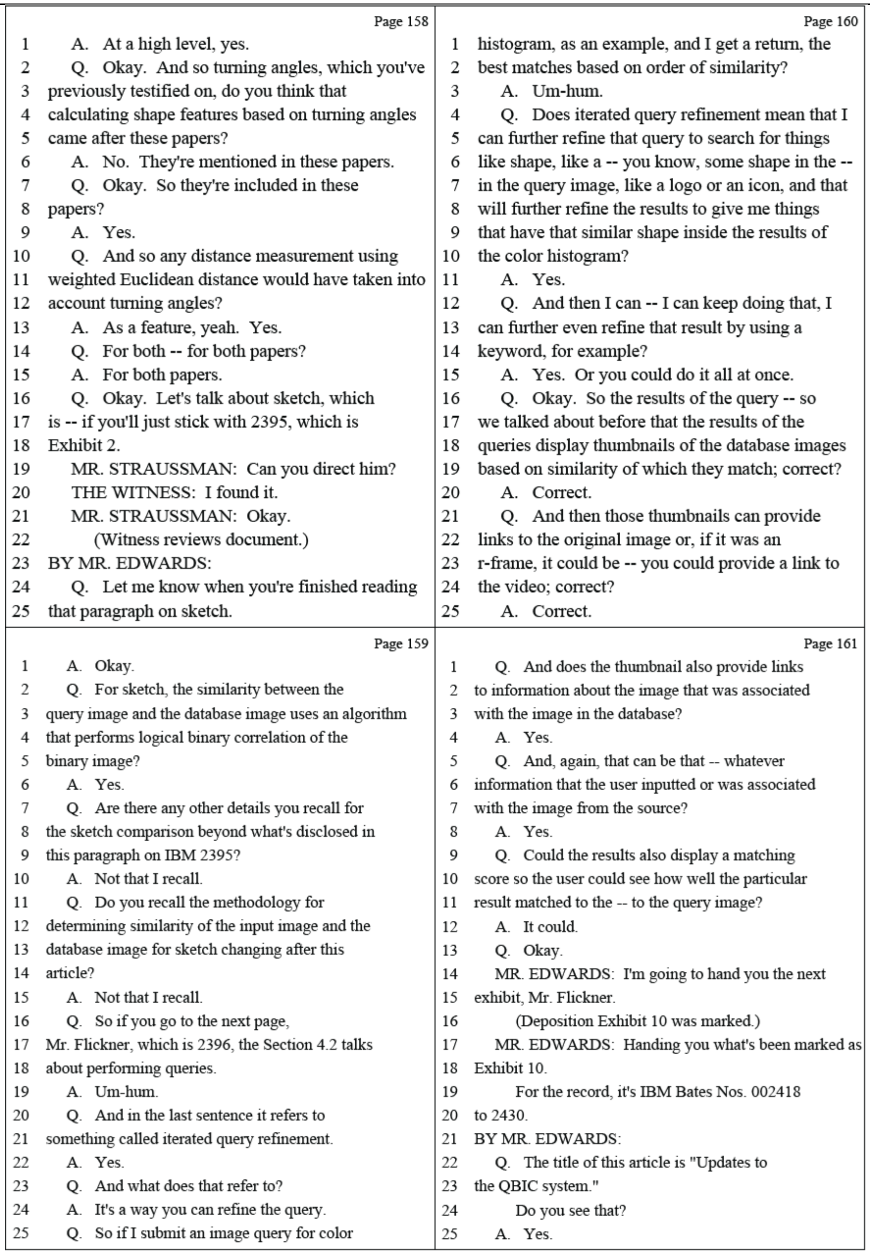

Page 158

A. At a high level, yes.
Q. Okay. And so turning angles, which you've previously testified on, do you think that calculating shape features based on turning angles came after these papers?
A. No. They're mentioned in these papers.
Q. Okay. So they're included in these papers?
A. Yes.
Q. And so any distance measurement using weighted Euclidean distance would have taken into account turning angles?
A. As a feature, yeah. Yes.
Q. For both -- for both papers?
A. For both papers.
Q. Okay. Let's talk about sketch, which is -- if you'll just stick with 2395, which is Exhibit 2.
MR. STRAUSSMAN: Can you direct him?
THE WITNESS: I found it.
MR. STRAUSSMAN: Okay.
(Witness reviews document.)
BY MR. EDWARDS:
Q. Let me know when you're finished reading that paragraph on sketch.

Page 159

A. Okay.
Q. For sketch, the similarity between the query image and the database image uses an algorithm that performs logical binary correlation of the binary image?
A. Yes.
Q. Are there any other details you recall for the sketch comparison beyond what's disclosed in this paragraph on IBM 2395?
A. Not that I recall.
Q. Do you recall the methodology for determining similarity of the input image and the database image for sketch changing after this article?
A. Not that I recall.
Q. So if you go to the next page, Mr. Flickner, which is 2396, the Section 4.2 talks about performing queries.
A. Um-hum.
Q. And in the last sentence it refers to something called iterated query refinement.
A. Yes.
Q. And what does that refer to?
A. It's a way you can refine the query.
Q. So if I submit an image query for color

Page 160

histogram, as an example, and I get a return, the best matches based on order of similarity?
A. Um-hum.
Q. Does iterated query refinement mean that I can further refine that query to search for things like shape, like a -- you know, some shape in the -- in the query image, like a logo or an icon, and that will further refine the results to give me things that have that similar shape inside the results of the color histogram?
A. Yes.
Q. And then I can -- I can keep doing that, I can further even refine that result by using a keyword, for example?
A. Yes. Or you could do it all at once.
Q. Okay. So the results of the query -- so we talked about before that the results of the queries display thumbnails of the database images based on similarity of which they match; correct?
A. Correct.
Q. And then those thumbnails can provide links to the original image or, if it was an r-frame, it could be -- you could provide a link to the video; correct?
A. Correct.

Page 161

Q. And does the thumbnail also provide links to information about the image that was associated with the image in the database?
A. Yes.
Q. And, again, that can be that -- whatever information that the user inputted or was associated with the image from the source?
A. Yes.
Q. Could the results also display a matching score so the user could see how well the particular result matched to the -- to the query image?
A. It could.
Q. Okay.
MR. EDWARDS: I'm going to hand you the next exhibit, Mr. Flickner.
(Deposition Exhibit 10 was marked.)
MR. EDWARDS: Handing you what's been marked as Exhibit 10.
For the record, it's IBM Bates Nos. 002418 to 2430.
BY MR. EDWARDS:
Q. The title of this article is "Updates to the QBIC system."
Do you see that?
A. Yes.

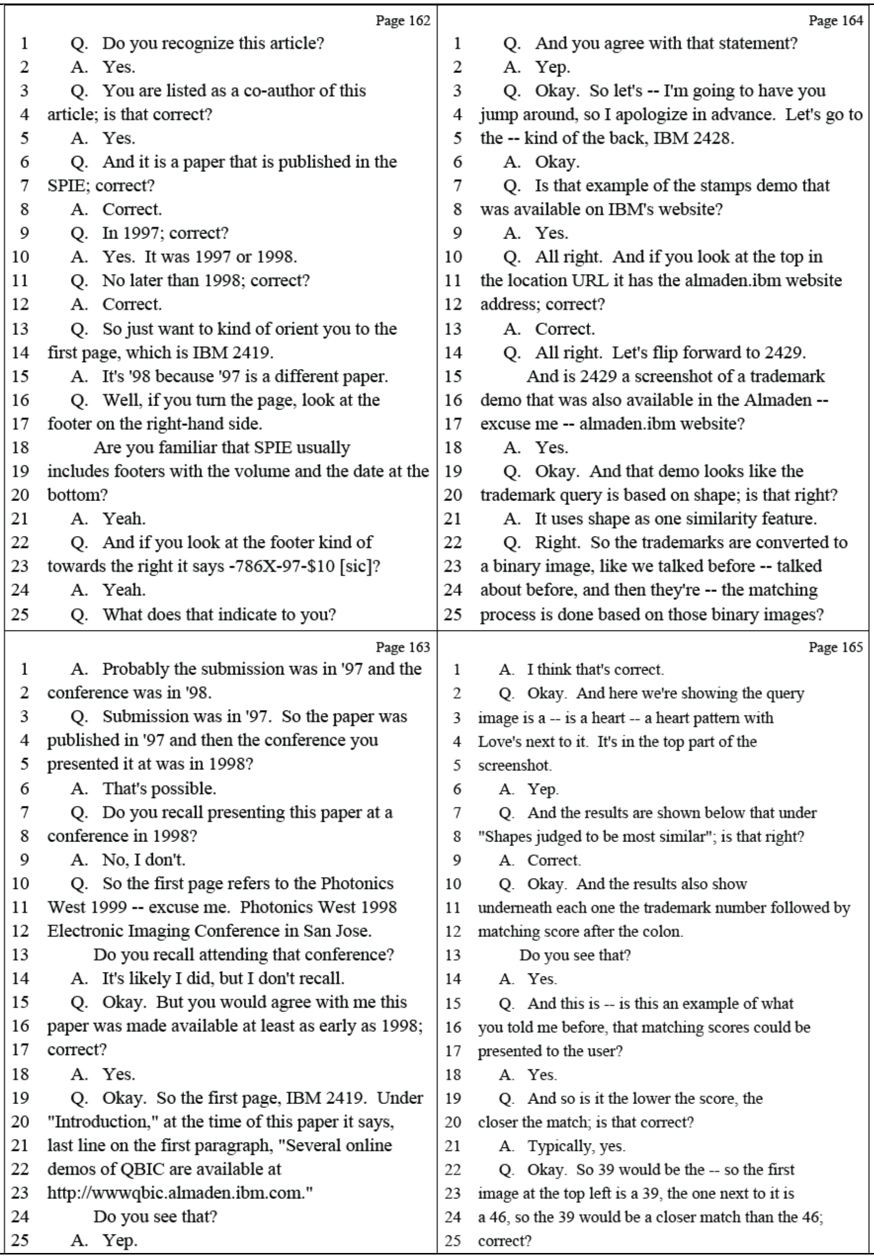

Page 162

Q. Do you recognize this article?
A. Yes.
Q. You are listed as a co-author of this article; is that correct?
A. Yes.
Q. And it is a paper that is published in the SPIE; correct?
A. Correct.
Q. In 1997; correct?
A. Yes. It was 1997 or 1998.
Q. No later than 1998; correct?
A. Correct.
Q. So just want to kind of orient you to the first page, which is IBM 2419.
A. It's '98 because '97 is a different paper.
Q. Well, if you turn the page, look at the footer on the right-hand side.
Are you familiar that SPIE usually includes footers with the volume and the date at the bottom?
A. Yeah.
Q. And if you look at the footer kind of towards the right it says -786X-97-$10 [sic]?
A. Yeah.
Q. What does that indicate to you?

Page 163

A. Probably the submission was in '97 and the conference was in '98.
Q. Submission was in '97. So the paper was published in '97 and then the conference you presented it at was in 1998?
A. That's possible.
Q. Do you recall presenting this paper at a conference in 1998?
A. No, I don't.
Q. So the first page refers to the Photonics West 1999 -- excuse me. Photonics West 1998 Electronic Imaging Conference in San Jose.
Do you recall attending that conference?
A. It's likely I did, but I don't recall.
Q. Okay. But you would agree with me this paper was made available at least as early as 1998; correct?
A. Yes.
Q. Okay. So the first page, IBM 2419. Under "Introduction," at the time of this paper it says, last line on the first paragraph, "Several online demos of QBIC are available at http://wwwqbic.almaden.ibm.com."
Do you see that?
A. Yep.

Page 164

Q. And you agree with that statement?
A. Yep.
Q. Okay. So let's -- I'm going to have you jump around, so I apologize in advance. Let's go to the -- kind of the back, IBM 2428.
A. Okay.
Q. Is that example of the stamps demo that was available on IBM's website?
A. Yes.
Q. All right. And if you look at the top in the location URL it has the almaden.ibm website address; correct?
A. Correct.
Q. All right. Let's flip forward to 2429.
And is 2429 a screenshot of a trademark demo that was also available in the Almaden -- excuse me -- almaden.ibm website?
A. Yes.
Q. Okay. And that demo looks like the trademark query is based on shape; is that right?
A. It uses shape as one similarity feature.
Q. Right. So the trademarks are converted to a binary image, like we talked before -- talked about before, and then they're -- the matching process is done based on those binary images?

Page 165

A. I think that's correct.
Q. Okay. And here we're showing the query image is a -- is a heart -- a heart pattern with Love's next to it. It's in the top part of the screenshot.
A. Yep.
Q. And the results are shown below that under "Shapes judged to be most similar"; is that right?
A. Correct.
Q. Okay. And the results also show underneath each one the trademark number followed by matching score after the colon.
Do you see that?
A. Yes.
Q. And this is -- is this an example of what you told me before, that matching scores could be presented to the user?
A. Yes.
Q. And so is it the lower the score, the closer the match; is that correct?
A. Typically, yes.
Q. Okay. So 39 would be the -- so the first image at the top left is a 39, the one next to it is a 46, so the 39 would be a closer match than the 46; correct?

Element	'036 Claim Recitation	Exemplary Citations in Reference
		Id. at 174–182: 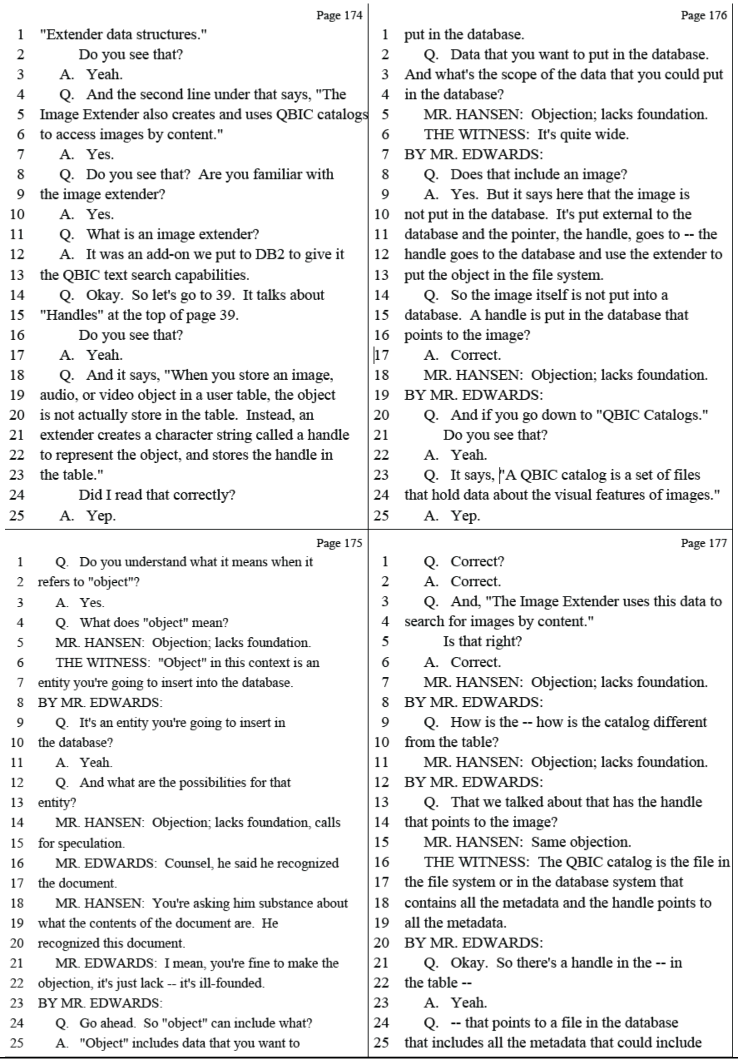 Page 174 "Extender data structures." Do you see that? A. Yeah. Q. And the second line under that says, "The Image Extender also creates and uses QBIC catalogs to access images by content." A. Yes. Q. Do you see that? Are you familiar with the image extender? A. Yes. Q. What is an image extender? A. It was an add-on we put to DB2 to give it the QBIC text search capabilities. Q. Okay. So let's go to 39. It talks about "Handles" at the top of page 39. Do you see that? A. Yeah. Q. And it says, "When you store an image, audio, or video object in a user table, the object is not actually store in the table. Instead, an extender creates a character string called a handle to represent the object, and stores the handle in the table." Did I read that correctly? A. Yep. Page 175 Q. Do you understand what it means when it refers to "object"? A. Yes. Q. What does "object" mean? MR. HANSEN: Objection; lacks foundation. THE WITNESS: "Object" in this context is an entity you're going to insert into the database. BY MR. EDWARDS: Q. It's an entity you're going to insert in the database? A. Yeah. Q. And what are the possibilities for that entity? MR. HANSEN: Objection; lacks foundation, calls for speculation. MR. EDWARDS: Counsel, he said he recognized the document. MR. HANSEN: You're asking him substance about what the contents of the document are. He recognized this document. MR. EDWARDS: I mean, you're fine to make the objection, it's just lack -- it's ill-founded. BY MR. EDWARDS: Q. Go ahead. So "object" can include what? A. "Object" includes data that you want to Page 176 put in the database. Q. Data that you want to put in the database. And what's the scope of the data that you could put in the database? MR. HANSEN: Objection; lacks foundation. THE WITNESS: It's quite wide. BY MR. EDWARDS: Q. Does that include an image? A. Yes. But it says here that the image is not put in the database. It's put external to the database and the pointer, the handle, goes to -- the handle goes to the database and use the extender to put the object in the file system. Q. So the image itself is not put into a database. A handle is put in the database that points to the image? A. Correct. MR. HANSEN: Objection; lacks foundation. BY MR. EDWARDS: Q. And if you go down to "QBIC Catalogs." Do you see that? A. Yeah. Q. It says, "A QBIC catalog is a set of files that hold data about the visual features of images." A. Yep. Page 177 Q. Correct? A. Correct. Q. And, "The Image Extender uses this data to search for images by content." Is that right? A. Correct. MR. HANSEN: Objection; lacks foundation. BY MR. EDWARDS: Q. How is the -- how is the catalog different from the table? MR. HANSEN: Objection; lacks foundation. BY MR. EDWARDS: Q. That we talked about that has the handle that points to the image? MR. HANSEN: Same objection. THE WITNESS: The QBIC catalog is the file in the file system or in the database system that contains all the metadata and the handle points to all the metadata. BY MR. EDWARDS: Q. Okay. So there's a handle in the -- in the table -- A. Yeah. Q. -- that points to a file in the database that includes all the metadata that could include

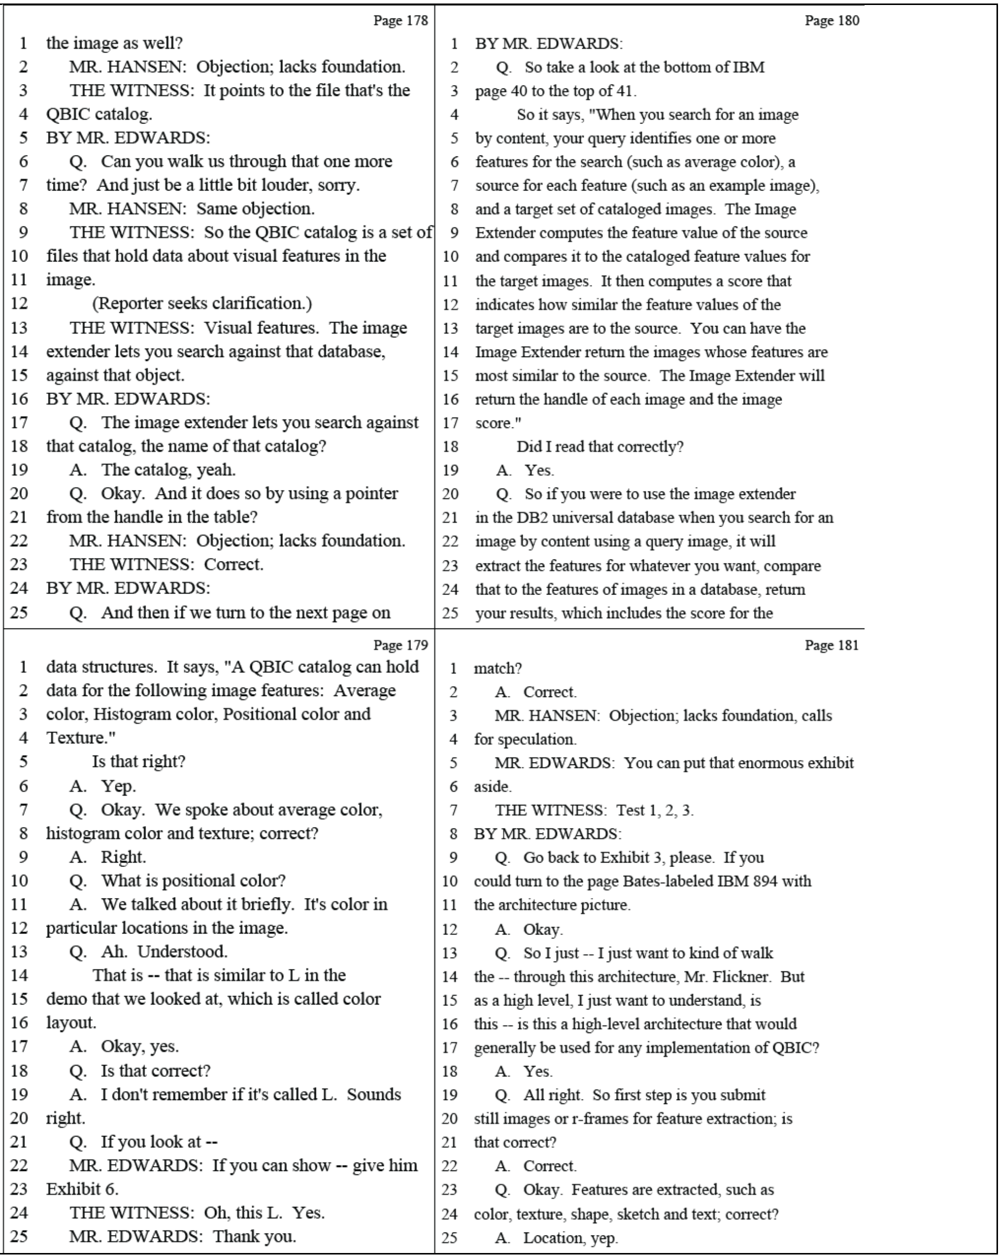

Page 178

the image as well?

MR. HANSEN: Objection; lacks foundation.

THE WITNESS: It points to the file that's the QBIC catalog.

BY MR. EDWARDS:

Q. Can you walk us through that one more time? And just be a little bit louder, sorry.

MR. HANSEN: Same objection.

THE WITNESS: So the QBIC catalog is a set of files that hold data about visual features in the image.

(Reporter seeks clarification.)

THE WITNESS: Visual features. The image extender lets you search against that database, against that object.

BY MR. EDWARDS:

Q. The image extender lets you search against that catalog, the name of that catalog?

A. The catalog, yeah.

Q. Okay. And it does so by using a pointer from the handle in the table?

MR. HANSEN: Objection; lacks foundation.

THE WITNESS: Correct.

BY MR. EDWARDS:

Q. And then if we turn to the next page on

Page 179

data structures. It says, "A QBIC catalog can hold data for the following image features: Average color, Histogram color, Positional color and Texture."

Is that right?

A. Yep.

Q. Okay. We spoke about average color, histogram color and texture; correct?

A. Right.

Q. What is positional color?

A. We talked about it briefly. It's color in particular locations in the image.

Q. Ah. Understood.

That is -- that is similar to L in the demo that we looked at, which is called color layout.

A. Okay, yes.

Q. Is that correct?

A. I don't remember if it's called L. Sounds right.

Q. If you look at --

MR. EDWARDS: If you can show -- give him Exhibit 6.

THE WITNESS: Oh, this L. Yes.

MR. EDWARDS: Thank you.

Page 180

BY MR. EDWARDS:

Q. So take a look at the bottom of IBM page 40 to the top of 41.

So it says, "When you search for an image by content, your query identifies one or more features for the search (such as average color), a source for each feature (such as an example image), and a target set of cataloged images. The Image Extender computes the feature value of the source and compares it to the cataloged feature values for the target images. It then computes a score that indicates how similar the feature values of the target images are to the source. You can have the Image Extender return the images whose features are most similar to the source. The Image Extender will return the handle of each image and the image score."

Did I read that correctly?

A. Yes.

Q. So if you were to use the image extender in the DB2 universal database when you search for an image by content using a query image, it will extract the features for whatever you want, compare that to the features of images in a database, return your results, which includes the score for the

Page 181

match?

A. Correct.

MR. HANSEN: Objection; lacks foundation, calls for speculation.

MR. EDWARDS: You can put that enormous exhibit aside.

THE WITNESS: Test 1, 2, 3.

BY MR. EDWARDS:

Q. Go back to Exhibit 3, please. If you could turn to the page Bates-labeled IBM 894 with the architecture picture.

A. Okay.

Q. So I just -- I just want to kind of walk the -- through this architecture, Mr. Flickner. But as a high level, I just want to understand, is this -- is this a high-level architecture that would generally be used for any implementation of QBIC?

A. Yes.

Q. All right. So first step is you submit still images or r-frames for feature extraction; is that correct?

A. Correct.

Q. Okay. Features are extracted, such as color, texture, shape, sketch and text; correct?

A. Location, yep.

Element	'036 Claim Recitation	Exemplary Citations in Reference
		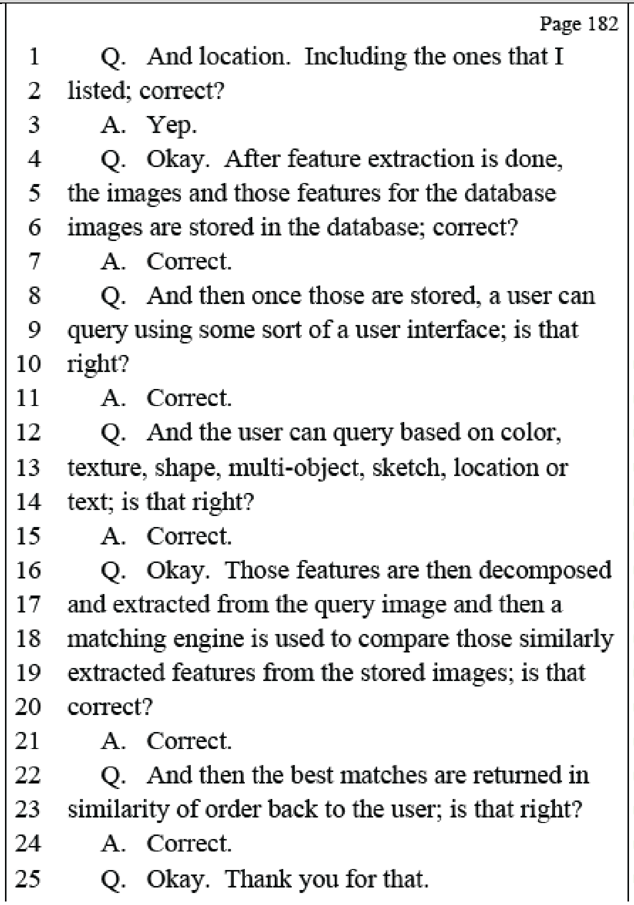 Page 182 Q. And location. Including the ones that I listed; correct? A. Yep. Q. Okay. After feature extraction is done, the images and those features for the database images are stored in the database; correct? A. Correct. Q. And then once those are stored, a user can query using some sort of a user interface; is that right? A. Correct. Q. And the user can query based on color, texture, shape, multi-object, sketch, location or text; is that right? A. Correct. Q. Okay. Those features are then decomposed and extracted from the query image and then a matching engine is used to compare those similarly extracted features from the stored images; is that correct? A. Correct. Q. And then the best matches are returned in similarity of order back to the user; is that right? A. Correct. Q. Okay. Thank you for that. *Id.* at 190–191: 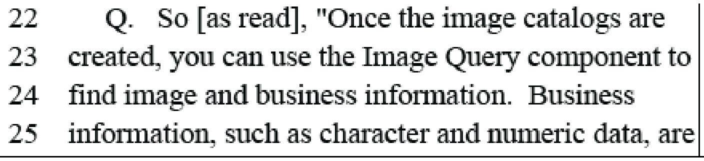 Q. So [as read], "Once the image catalogs are created, you can use the Image Query component to find image and business information. Business information, such as character and numeric data, are

		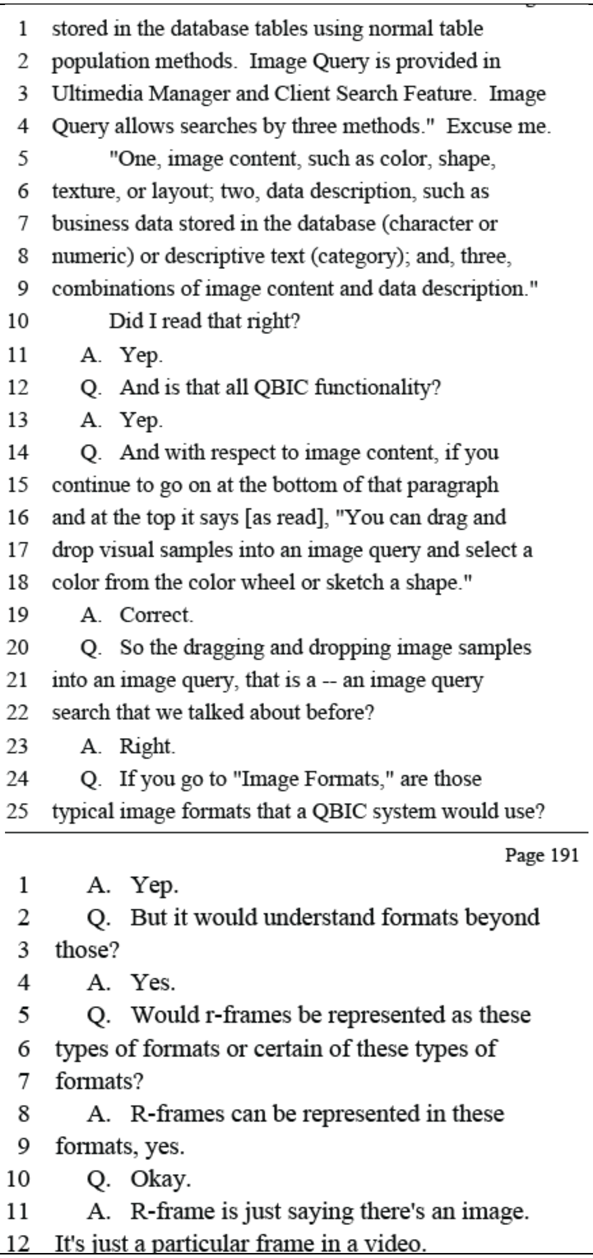stored in the database tables using normal table population methods. Image Query is provided in Ultimedia Manager and Client Search Feature. Image Query allows searches by three methods." Excuse me. "One, image content, such as color, shape, texture, or layout; two, data description, such as business data stored in the database (character or numeric) or descriptive text (category); and, three, combinations of image content and data description." Did I read that right? A. Yep. Q. And is that all QBIC functionality? A. Yep. Q. And with respect to image content, if you continue to go on at the bottom of that paragraph and at the top it says [as read], "You can drag and drop visual samples into an image query and select a color from the color wheel or sketch a shape." A. Correct. Q. So the dragging and dropping image samples into an image query, that is a -- an image query search that we talked about before? A. Right. Q. If you go to "Image Formats," are those typical image formats that a QBIC system would use? Page 191 A. Yep. Q. But it would understand formats beyond those? A. Yes. Q. Would r-frames be represented as these types of formats or certain of these types of formats? A. R-frames can be represented in these formats, yes. Q. Okay. A. R-frame is just saying there's an image. It's just a particular frame in a video.

Element	'036 Claim Recitation	Exemplary Citations in Reference
		Id. at 198 – 199: 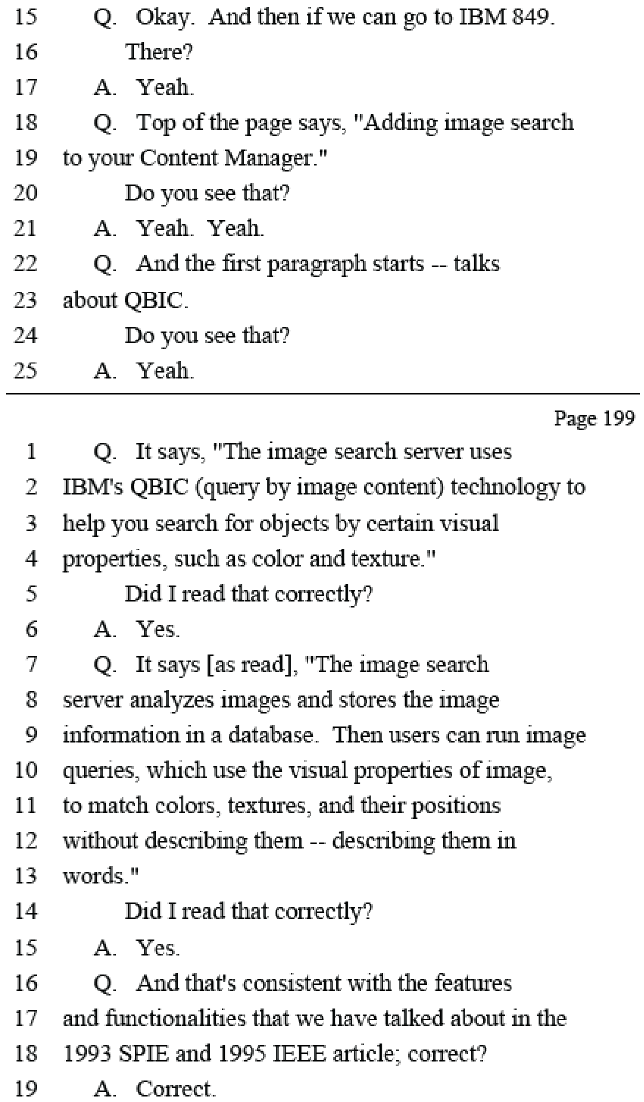 Q. Okay. And then if we can go to IBM 849. There? A. Yeah. Q. Top of the page says, "Adding image search to your Content Manager." Do you see that? A. Yeah. Yeah. Q. And the first paragraph starts -- talks about QBIC. Do you see that? A. Yeah. Page 199 Q. It says, "The image search server uses IBM's QBIC (query by image content) technology to help you search for objects by certain visual properties, such as color and texture." Did I read that correctly? A. Yes. Q. It says [as read], "The image search server analyzes images and stores the image information in a database. Then users can run image queries, which use the visual properties of image, to match colors, textures, and their positions without describing them -- describing them in words." Did I read that correctly? A. Yes. Q. And that's consistent with the features and functionalities that we have talked about in the 1993 SPIE and 1995 IEEE article; correct? A. Correct. *Id.* at 312 – 313:

Element	'036 Claim Recitation	Exemplary Citations in Reference
		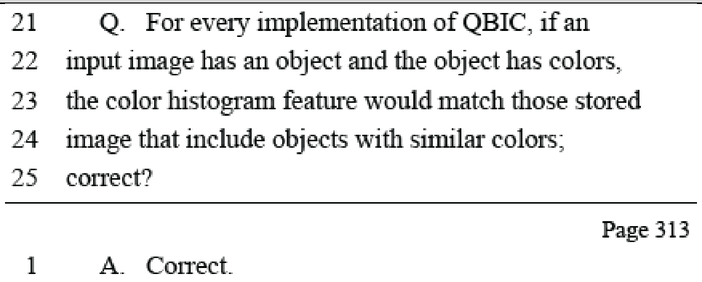 Q. For every implementation of QBIC, if an input image has an object and the object has colors, the color histogram feature would match those stored image that include objects with similar colors; correct? Page 313 A. Correct. *Id.* at 317–322: 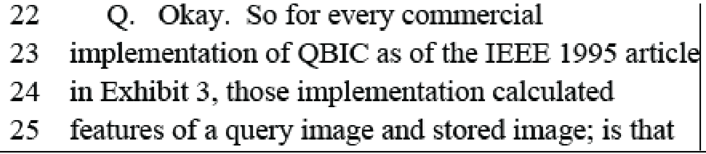 Q. Okay. So for every commercial implementation of QBIC as of the IEEE 1995 article in Exhibit 3, those implementation calculated features of a query image and stored image; is that

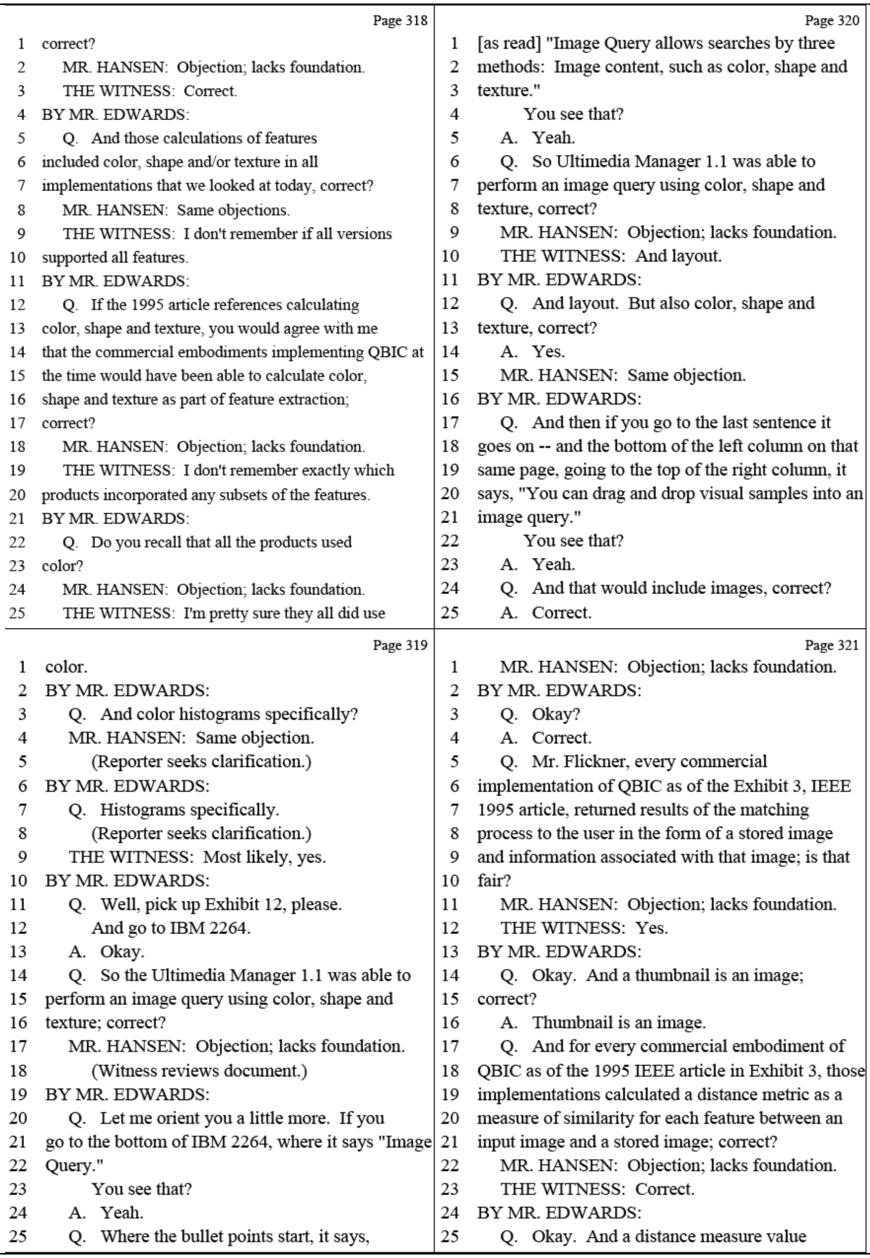

Page 318

correct?
MR. HANSEN: Objection; lacks foundation.
THE WITNESS: Correct.
BY MR. EDWARDS:
Q. And those calculations of features included color, shape and/or texture in all implementations that we looked at today, correct?
MR. HANSEN: Same objections.
THE WITNESS: I don't remember if all versions supported all features.
BY MR. EDWARDS:
Q. If the 1995 article references calculating color, shape and texture, you would agree with me that the commercial embodiments implementing QBIC at the time would have been able to calculate color, shape and texture as part of feature extraction; correct?
MR. HANSEN: Objection; lacks foundation.
THE WITNESS: I don't remember exactly which products incorporated any subsets of the features.
BY MR. EDWARDS:
Q. Do you recall that all the products used color?
MR. HANSEN: Objection; lacks foundation.
THE WITNESS: I'm pretty sure they all did use

Page 319

color.
BY MR. EDWARDS:
Q. And color histograms specifically?
MR. HANSEN: Same objection.
(Reporter seeks clarification.)
BY MR. EDWARDS:
Q. Histograms specifically.
(Reporter seeks clarification.)
THE WITNESS: Most likely, yes.
BY MR. EDWARDS:
Q. Well, pick up Exhibit 12, please.
And go to IBM 2264.
A. Okay.
Q. So the Ultimedia Manager 1.1 was able to perform an image query using color, shape and texture; correct?
MR. HANSEN: Objection; lacks foundation.
(Witness reviews document.)
BY MR. EDWARDS:
Q. Let me orient you a little more. If you go to the bottom of IBM 2264, where it says "Image Query."
You see that?
A. Yeah.
Q. Where the bullet points start, it says,

Page 320

[as read] "Image Query allows searches by three methods: Image content, such as color, shape and texture."
You see that?
A. Yeah.
Q. So Ultimedia Manager 1.1 was able to perform an image query using color, shape and texture, correct?
MR. HANSEN: Objection; lacks foundation.
THE WITNESS: And layout.
BY MR. EDWARDS:
Q. And layout. But also color, shape and texture, correct?
A. Yes.
MR. HANSEN: Same objection.
BY MR. EDWARDS:
Q. And then if you go to the last sentence it goes on -- and the bottom of the left column on that same page, going to the top of the right column, it says, "You can drag and drop visual samples into an image query."
You see that?
A. Yeah.
Q. And that would include images, correct?
A. Correct.

Page 321

MR. HANSEN: Objection; lacks foundation.
BY MR. EDWARDS:
Q. Okay?
A. Correct.
Q. Mr. Flickner, every commercial implementation of QBIC as of the Exhibit 3, IEEE 1995 article, returned results of the matching process to the user in the form of a stored image and information associated with that image; is that fair?
MR. HANSEN: Objection; lacks foundation.
THE WITNESS: Yes.
BY MR. EDWARDS:
Q. Okay. And a thumbnail is an image; correct?
A. Thumbnail is an image.
Q. And for every commercial embodiment of QBIC as of the 1995 IEEE article in Exhibit 3, those implementations calculated a distance metric as a measure of similarity for each feature between an input image and a stored image; correct?
MR. HANSEN: Objection; lacks foundation.
THE WITNESS: Correct.
BY MR. EDWARDS:
Q. Okay. And a distance measure value

Element	'036 Claim Recitation	Exemplary Citations in Reference
		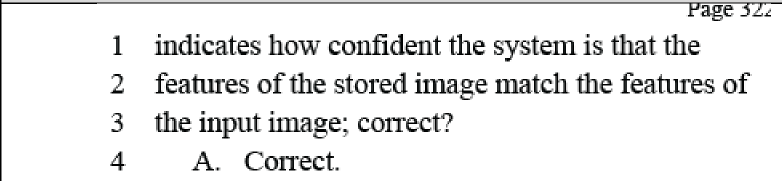 Page 322 indicates how confident the system is that the features of the stored image match the features of the input image; correct? A. Correct. *Id.* at 305: 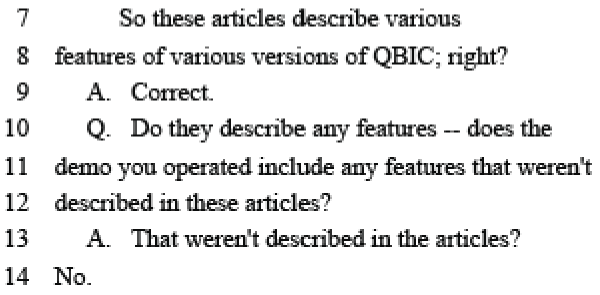 So these articles describe various features of various versions of QBIC; right? A. Correct. Q. Do they describe any features -- does the demo you operated include any features that weren't described in these articles? A. That weren't described in the articles? No. *Id.* at 35:

Element	'036 Claim Recitation	Exemplary Citations in Reference
		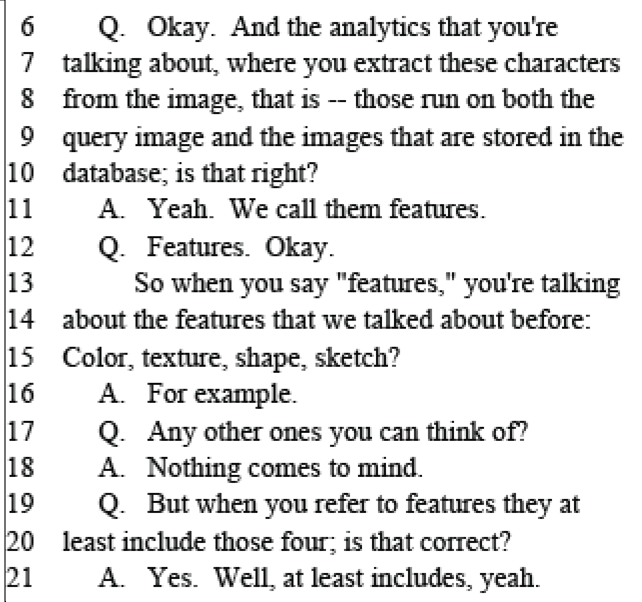 Q. Okay. And the analytics that you're talking about, where you extract these characters from the image, that is -- those run on both the query image and the images that are stored in the database; is that right? A. Yeah. We call them features. Q. Features. Okay. So when you say "features," you're talking about the features that we talked about before: Color, texture, shape, sketch? A. For example. Q. Any other ones you can think of? A. Nothing comes to mind. Q. But when you refer to features they at least include those four; is that correct? A. Yes. Well, at least includes, yeah. *Id.* at 109-110:

Element	'036 Claim Recitation	Exemplary Citations in Reference
		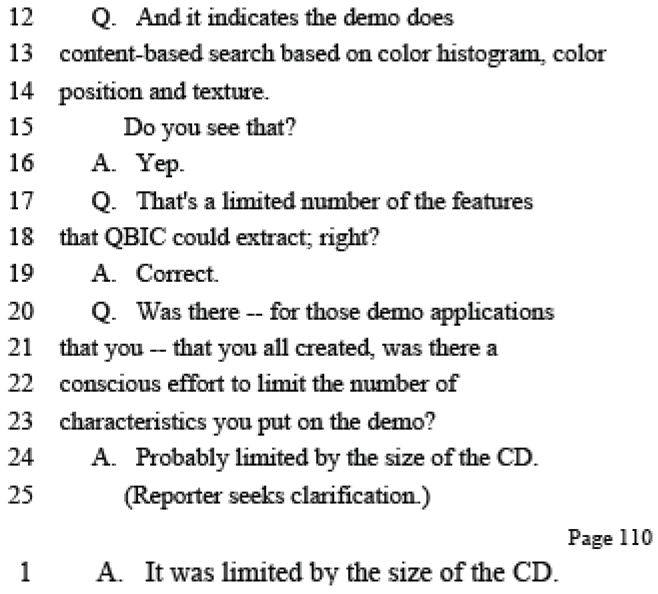 Q. And it indicates the demo does content-based search based on color histogram, color position and texture. Do you see that? A. Yep. Q. That's a limited number of the features that QBIC could extract; right? A. Correct. Q. Was there -- for those demo applications that you -- that you all created, was there a conscious effort to limit the number of characteristics you put on the demo? A. Probably limited by the size of the CD. (Reporter seeks clarification.) Page 110 A. It was limited by the size of the CD. *Id.* at 311: 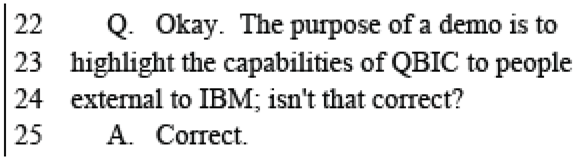 Q. Okay. The purpose of a demo is to highlight the capabilities of QBIC to people external to IBM; isn't that correct? A. Correct. *Id.* at 313-317:

Element	'036 Claim Recitation	Exemplary Citations in Reference
		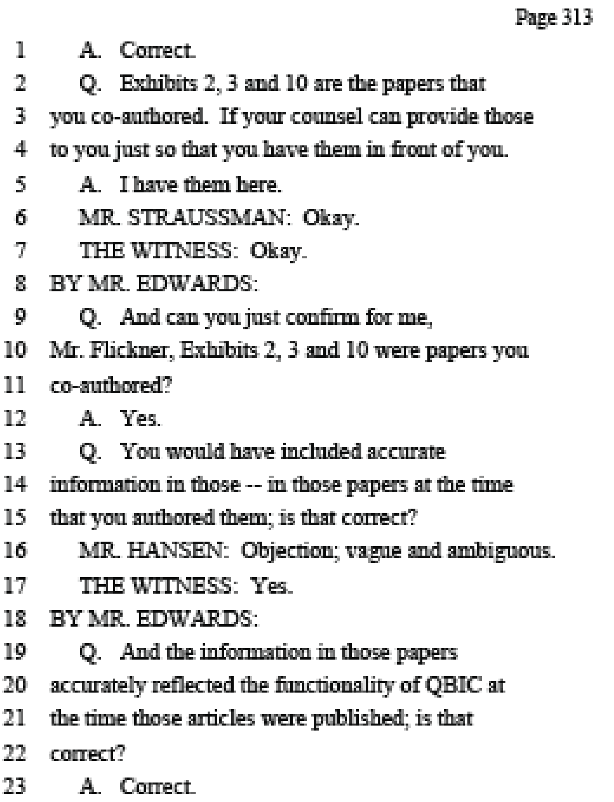 Page 313 A. Correct. Q. Exhibits 2, 3 and 10 are the papers that you co-authored. If your counsel can provide those to you just so that you have them in front of you. A. I have them here. MR. STRAUSSMAN: Okay. THE WITNESS: Okay. BY MR. EDWARDS: Q. And can you just confirm for me, Mr. Flickner, Exhibits 2, 3 and 10 were papers you co-authored? A. Yes. Q. You would have included accurate information in those -- in those papers at the time that you authored them; is that correct? MR. HANSEN: Objection; vague and ambiguous. THE WITNESS: Yes. BY MR. EDWARDS: Q. And the information in those papers accurately reflected the functionality of QBIC at the time those articles were published; is that correct? A. Correct.

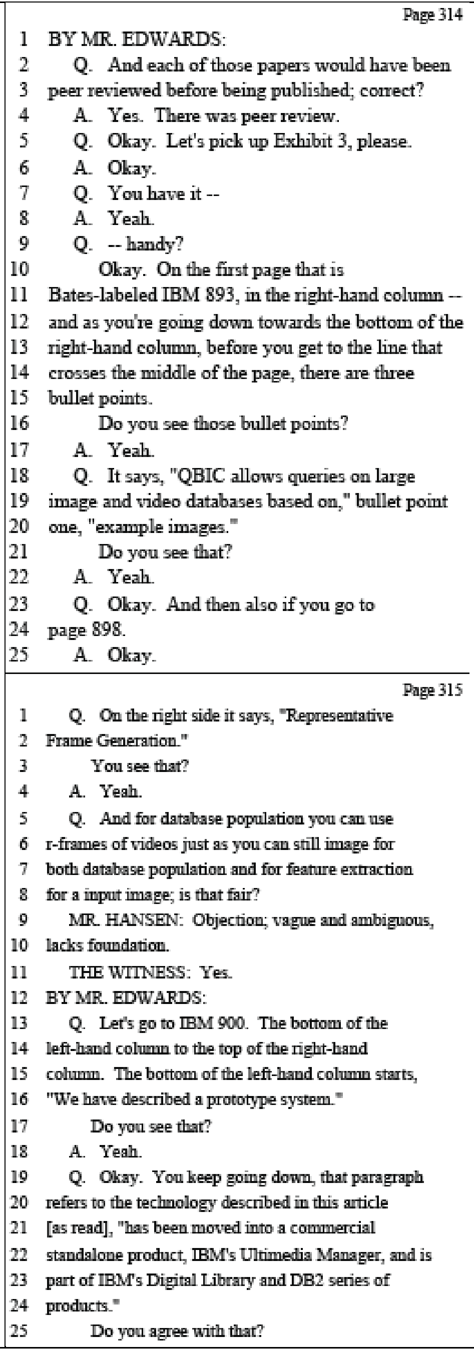

Page 314

BY MR. EDWARDS:
Q. And each of those papers would have been peer reviewed before being published; correct?
A. Yes. There was peer review.
Q. Okay. Let's pick up Exhibit 3, please.
A. Okay.
Q. You have it --
A. Yeah.
Q. -- handy?
Okay. On the first page that is Bates-labeled IBM 893, in the right-hand column -- and as you're going down towards the bottom of the right-hand column, before you get to the line that crosses the middle of the page, there are three bullet points.
Do you see those bullet points?
A. Yeah.
Q. It says, "QBIC allows queries on large image and video databases based on," bullet point one, "example images."
Do you see that?
A. Yeah.
Q. Okay. And then also if you go to page 898.
A. Okay.

Page 315

Q. On the right side it says, "Representative Frame Generation."
You see that?
A. Yeah.
Q. And for database population you can use r-frames of videos just as you can still image for both database population and for feature extraction for a input image; is that fair?
MR. HANSEN: Objection; vague and ambiguous, lacks foundation.
THE WITNESS: Yes.
BY MR. EDWARDS:
Q. Let's go to IBM 900. The bottom of the left-hand column to the top of the right-hand column. The bottom of the left-hand column starts, "We have described a prototype system."
Do you see that?
A. Yeah.
Q. Okay. You keep going down, that paragraph refers to the technology described in this article [as read], "has been moved into a commercial standalone product, IBM's Ultimedia Manager, and is part of IBM's Digital Library and DB2 series of products."
Do you agree with that?

Element	'036 Claim Recitation	Exemplary Citations in Reference
		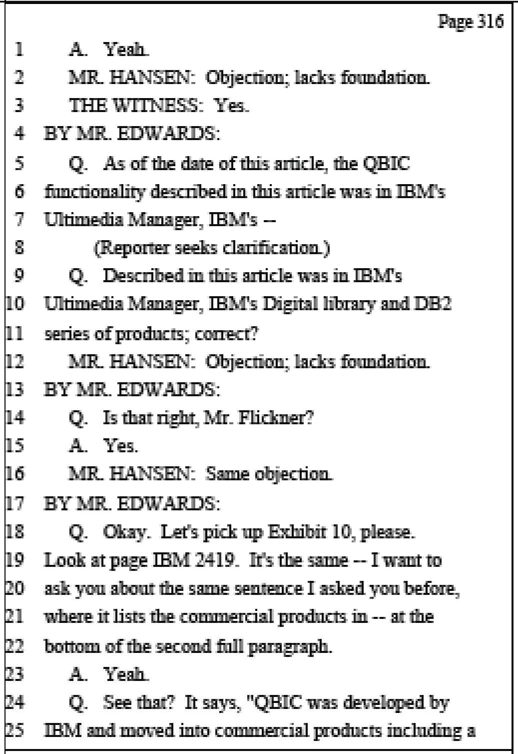 Page 316 A. Yeah. MR. HANSEN: Objection; lacks foundation. THE WITNESS: Yes. BY MR. EDWARDS: Q. As of the date of this article, the QBIC functionality described in this article was in IBM's Ultimedia Manager, IBM's -- (Reporter seeks clarification.) Q. Described in this article was in IBM's Ultimedia Manager, IBM's Digital library and DB2 series of products; correct? MR. HANSEN: Objection; lacks foundation. BY MR. EDWARDS: Q. Is that right, Mr. Flickner? A. Yes. MR. HANSEN: Same objection. BY MR. EDWARDS: Q. Okay. Let's pick up Exhibit 10, please. Look at page IBM 2419. It's the same -- I want to ask you about the same sentence I asked you before, where it lists the commercial products in -- at the bottom of the second full paragraph. A. Yeah. Q. See that? It says, "QBIC was developed by IBM and moved into commercial products including a 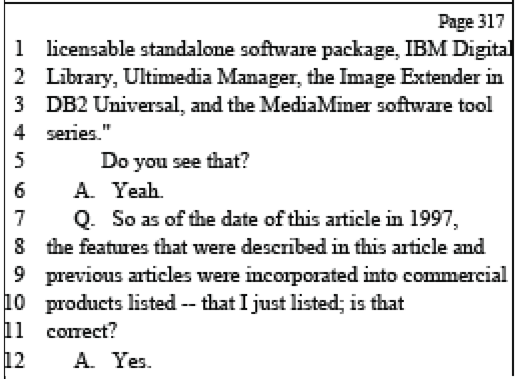 Page 317 licensable standalone software package, IBM Digital Library, Ultimedia Manager, the Image Extender in DB2 Universal, and the MediaMiner software tool series." Do you see that? A. Yeah. Q. So as of the date of this article in 1997, the features that were described in this article and previous articles were incorporated into commercial products listed -- that I just listed; is that correct? A. Yes.

Element	'036 Claim Recitation	Exemplary Citations in Reference
		IBM000001 at 39 – 41 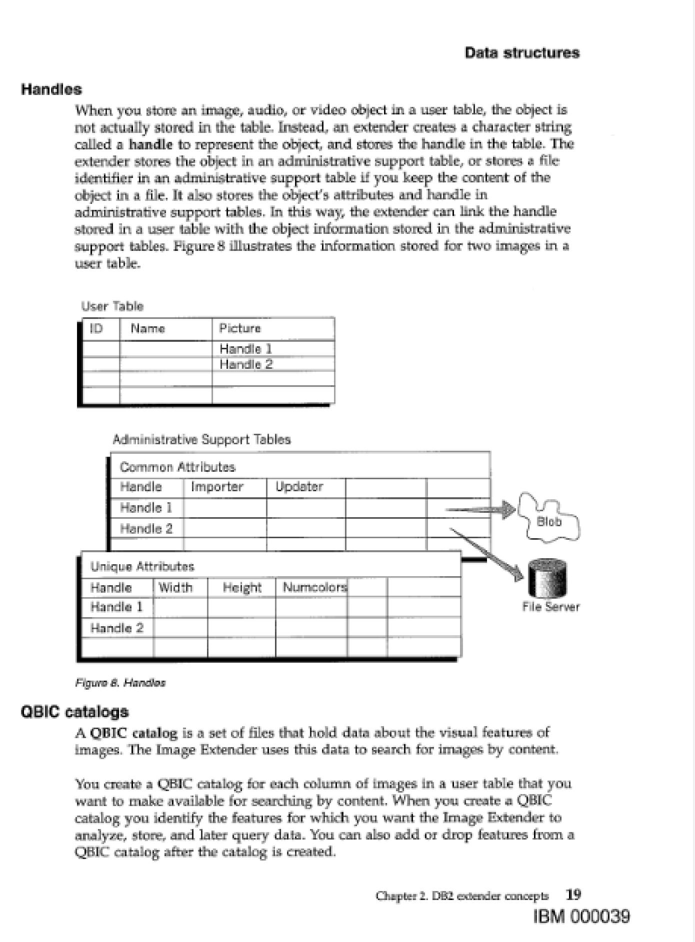

Data structures

Handles

When you store an image, audio, or video object in a user table, the object is not actually stored in the table. Instead, an extender creates a character string called a **handle** to represent the object, and stores the handle in the table. The extender stores the object in an administrative support table, or stores a file identifier in an administrative support table if you keep the content of the object in a file. It also stores the object's attributes and handle in administrative support tables. In this way, the extender can link the handle stored in a user table with the object information stored in the administrative support tables. Figure 8 illustrates the information stored for two images in a user table.

Figure 8. Handles

QBIC catalogs

A **QBIC catalog** is a set of files that hold data about the visual features of images. The Image Extender uses this data to search for images by content.

You create a QBIC catalog for each column of images in a user table that you want to make available for searching by content. When you create a QBIC catalog you identify the features for which you want the Image Extender to analyze, store, and later query data. You can also add or drop features from a QBIC catalog after the catalog is created.

Chapter 2. DB2 extender concepts **19**

IBM 000039

Element	'036 Claim Recitation	Exemplary Citations in Reference
		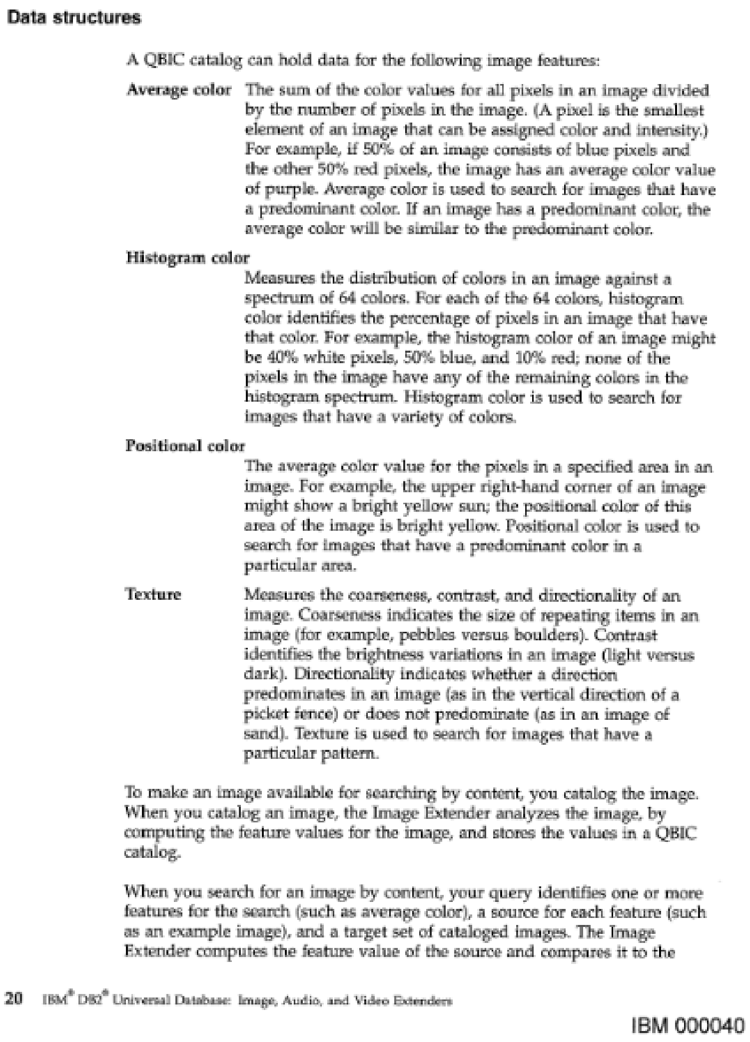 **Data structures** A QBIC catalog can hold data for the following image features: **Average color** The sum of the color values for all pixels in an image divided by the number of pixels in the image. (A pixel is the smallest element of an image that can be assigned color and intensity.) For example, if 50% of an image consists of blue pixels and the other 50% red pixels, the image has an average color value of purple. Average color is used to search for images that have a predominant color. If an image has a predominant color, the average color will be similar to the predominant color. **Histogram color** Measures the distribution of colors in an image against a spectrum of 64 colors. For each of the 64 colors, histogram color identifies the percentage of pixels in an image that have that color. For example, the histogram color of an image might be 40% white pixels, 50% blue, and 10% red; none of the pixels in the image have any of the remaining colors in the histogram spectrum. Histogram color is used to search for images that have a variety of colors. **Positional color** The average color value for the pixels in a specified area in an image. For example, the upper right-hand corner of an image might show a bright yellow sun; the positional color of this area of the image is bright yellow. Positional color is used to search for images that have a predominant color in a particular area. **Texture** Measures the coarseness, contrast, and directionality of an image. Coarseness indicates the size of repeating items in an image (for example, pebbles versus boulders). Contrast identifies the brightness variations in an image (light versus dark). Directionality indicates whether a direction predominates in an image (as in the vertical direction of a picket fence) or does not predominate (as in an image of sand). Texture is used to search for images that have a particular pattern. To make an image available for searching by content, you catalog the image. When you catalog an image, the Image Extender analyzes the image, by computing the feature values for the image, and stores the values in a QBIC catalog. When you search for an image by content, your query identifies one or more features for the search (such as average color), a source for each feature (such as an example image), and a target set of cataloged images. The Image Extender computes the feature value of the source and compares it to the 20 IBM® DB2® Universal Database: Image, Audio, and Video Extenders IBM 000040

Element	'036 Claim Recitation	Exemplary Citations in Reference
		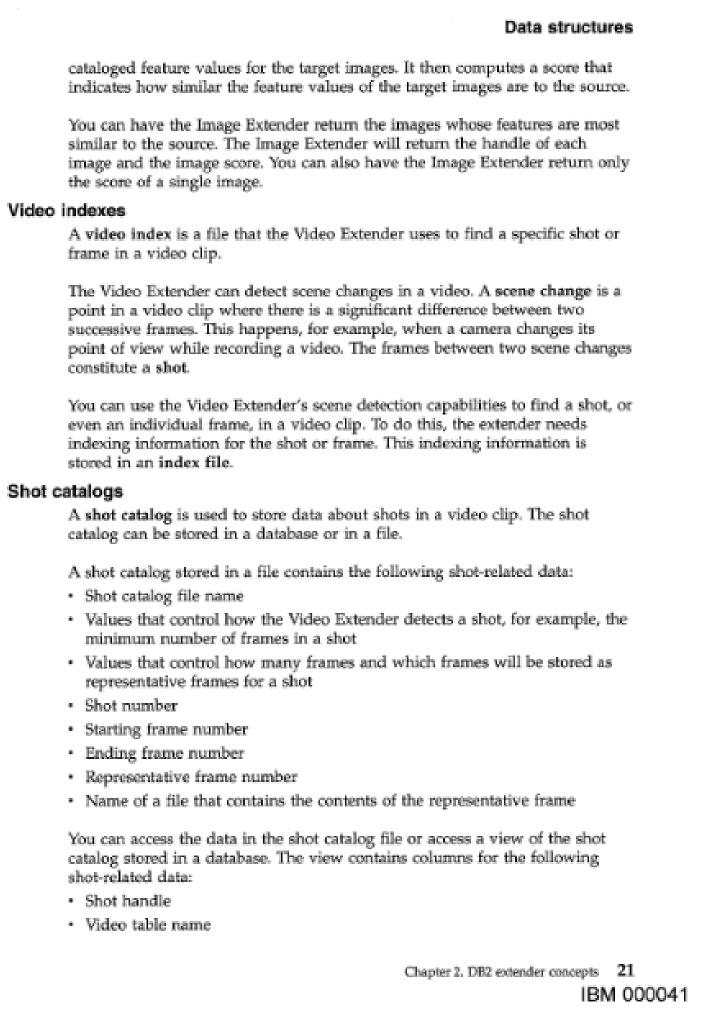 **Data structures** cataloged feature values for the target images. It then computes a score that indicates how similar the feature values of the target images are to the source. You can have the Image Extender return the images whose features are most similar to the source. The Image Extender will return the handle of each image and the image score. You can also have the Image Extender return only the score of a single image. **Video indexes** A **video index** is a file that the Video Extender uses to find a specific shot or frame in a video clip. The Video Extender can detect scene changes in a video. A **scene change** is a point in a video clip where there is a significant difference between two successive frames. This happens, for example, when a camera changes its point of view while recording a video. The frames between two scene changes constitute a **shot**. You can use the Video Extender's scene detection capabilities to find a shot, or even an individual frame, in a video clip. To do this, the extender needs indexing information for the shot or frame. This indexing information is stored in an **index file**. **Shot catalogs** A **shot catalog** is used to store data about shots in a video clip. The shot catalog can be stored in a database or in a file. A shot catalog stored in a file contains the following shot-related data: • Shot catalog file name • Values that control how the Video Extender detects a shot, for example, the minimum number of frames in a shot • Values that control how many frames and which frames will be stored as representative frames for a shot • Shot number • Starting frame number • Ending frame number • Representative frame number • Name of a file that contains the contents of the representative frame You can access the data in the shot catalog file or access a view of the shot catalog stored in a database. The view contains columns for the following shot-related data: • Shot handle • Video table name Chapter 2. DB2 extender concepts 21 IBM 000041 IBM0002263 at 226

Element	'036 Claim Recitation	Exemplary Citations in Reference
		 IBM000777 at 794

Description

Ultimedia Manager 1.1 opens new dimensions in visual queries, helping you manage your image database more efficiently. With its easy-to-use interface, even entry-level users can perform sophisticated searches. You can search by image content, data description, or a combination of both. Once the images are found, you can view the images in full resolution, use the images in another application, or read the associated textual comments.

You can use Ultimedia Manager 1.1 either as a stand-alone product or in a LAN environment. As a stand-alone product, Ultimedia Manager 1.1 employs full image classification and image query capabilities. Within the LAN environment, the Ultimedia Manager image catalogs can be sorted on a DB2 for OS/2 or DB2 for AIX database server. The image catalogs can be accessed by Windows or OS/2 clients using the Client Search Feature or the Visualizer Ultimedia Query.

Ultimedia Manager runs as an OS/2 application to a DB2 for OS/2 or DB2 for AIX database server in a LAN, or as a stand-alone application with a local DB2 for OS/2 database. The Client Search Feature runs on DOS or Windows, and accesses files in dBASE flat file format. A database manager is not required to access catalogs in dBASE file format.

Image Classification

With the Image Classifier, you can create visual indexes based on color, shape, category, and layout elements, which are areas of the image containing the patterns you want to use for searches. You can store these elements in catalogs in DB2 for OS/2 databases. The catalogs can be used locally or, through the use of utility programs, be transferred to DB2 for OS/2 or AIX servers. In addition, you can convert the catalog data into dBASE flat file format for access by Windows client applications.

The Image Classifier and Analyzer allows automatic or interactive classification of images for later image queries, flexibility in storage of catalog information, and access to multiple users. You can organize large image catalogs so that later searches for images can be performed easily. You can classify automatically using a batch program, or interactively by specifying areas of interest in an image and indicating the visual cue (color, shape, or texture) to be saved for later queries.

Image Query

Once the image catalogs are created, you can use the Image Query component to find image and business information. Business information, such as character and numeric data, are stored in the database tables using normal table population methods. Image Query is provided in Ultimedia Manager and Client Search Feature.

Image Query allows searches by three methods:

- Image content, such as color, shape, texture, or layout
- Data description, such as business data stored in the database (character or numeric) or descriptive text (category)
- Combinations of image content and data description.

Image Content: To search by image content, specify the attributes (color, shape, texture, or layout). You can drag and drop visual samples into an image query, select a color from the color wheel, or sketch a shape.

Data Description: To search by data description, enter character and numeric data values in the database tables using the simple fill-in-the-blanks query panel. You can enter data such as name, number, date, or a range of dates. You can integrate character or numeric data as part of the image query, enabling a more effective search for better results. The simple fill-in-the-blanks method for specifying desired conditions is so easy even non-programmers can use database table information as part of their searches.

Combinations: Image queries using either image content or data description, or combinations of both, allow you to search through thousands of images more quickly and effectively than the traditional text and keyword search methods.

Image Formats

Ultimedia Manager 1.1 understands these image formats:

- Audio Visual Connection (AVC) images
- OS/2 bitmap
- Windows 3.x bitmap
- PCX
- Intel and Motorola TIFF
- TARGA TGA
- TIFF (FAX)
- GIF
- Kodak Photo-CD

Product Positioning

The advanced image search technology in Ultimedia Manager, query by image content, was developed in conjunction with the IBM Almaden Research Center. It provides function not available with conventional, text-only image management systems.

Ultimedia Manager fulfills image search and retrieval needs of business users and creative professionals through use of online image libraries or collections. It is intended for users who have large image libraries and need fast, flexible access to their image data, users who may also want to add character or numeric data to their image queries, and users who feel limited by text-based systems.

Ultimedia Manager 1.1 works with the DB2 Family of products and the Visualizer Ultimedia Query product. Catalogs created by Ultimedia Manager can be stored in DB2 for OS/2 database tables. Image catalogs also can be stored in DB2 for OS/2 or DB2 for AIX database servers in client/server environments.

OS/2 and Windows clients in a LAN environment using the Client Search Feature and the Visualizer Ultimedia Query can access catalogs created with Ultimedia Manager 1.1. Visualizer Ultimedia Query provides users with the capability to:

- Add multimedia columns (such as image, video, audio, and formatted text) to DB2 for OS/2 database tables
- Query and display images
- View or listen to video and audio information
- Display formatted text

Users of Ultimedia Manager 1.1 and Client Search Feature are focused on image management and query applications. Visualizer Ultimedia Query users are

Element	'036 Claim Recitation	Exemplary Citations in Reference
		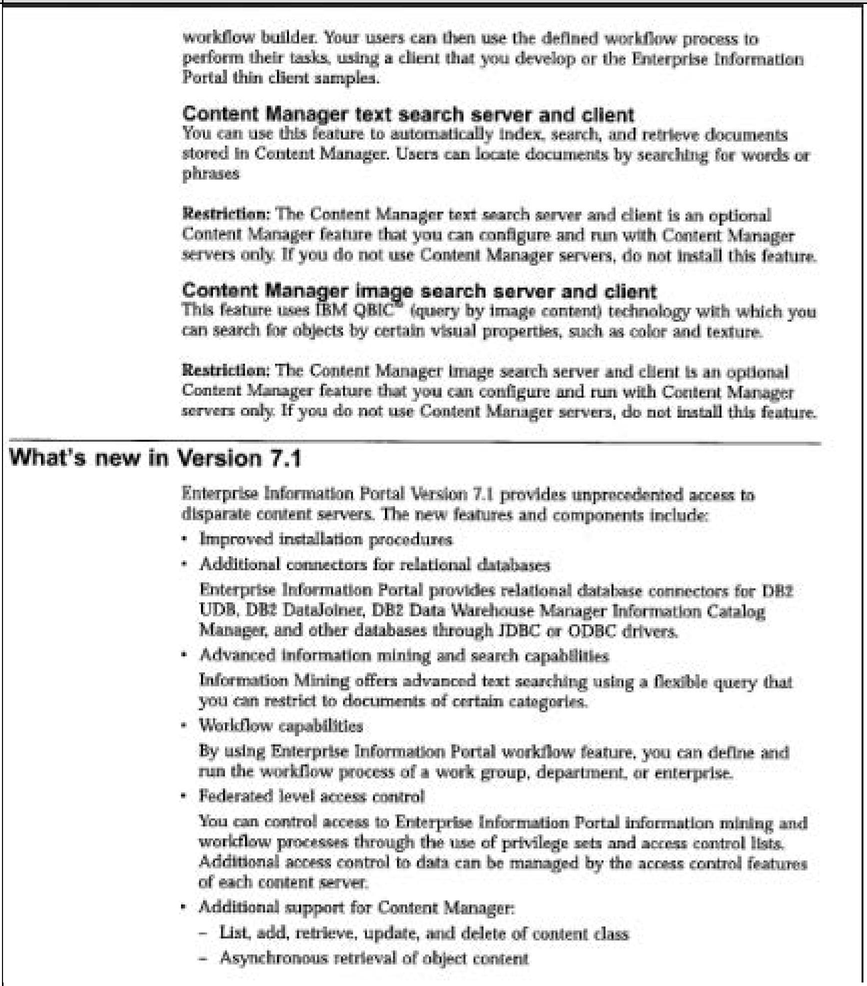 workflow builder. Your users can then use the defined workflow process to perform their tasks, using a client that you develop or the Enterprise Information Portal thin client samples. **Content Manager text search server and client** You can use this feature to automatically index, search, and retrieve documents stored in Content Manager. Users can locate documents by searching for words or phrases **Restriction:** The Content Manager text search server and client is an optional Content Manager feature that you can configure and run with Content Manager servers only. If you do not use Content Manager servers, do not install this feature. **Content Manager image search server and client** This feature uses IBM QBIC™ (query by image content) technology with which you can search for objects by certain visual properties, such as color and texture. **Restriction:** The Content Manager image search server and client is an optional Content Manager feature that you can configure and run with Content Manager servers only. If you do not use Content Manager servers, do not install this feature. **What's new in Version 7.1** Enterprise Information Portal Version 7.1 provides unprecedented access to disparate content servers. The new features and components include: • Improved installation procedures • Additional connectors for relational databases Enterprise Information Portal provides relational database connectors for DB2 UDB, DB2 DataJoiner, DB2 Data Warehouse Manager Information Catalog Manager, and other databases through JDBC or ODBC drivers. • Advanced information mining and search capabilities Information Mining offers advanced text searching using a flexible query that you can restrict to documents of certain categories. • Workflow capabilities By using Enterprise Information Portal workflow feature, you can define and run the workflow process of a work group, department, or enterprise. • Federated level access control You can control access to Enterprise Information Portal information mining and workflow processes through the use of privilege sets and access control lists. Additional access control to data can be managed by the access control features of each content server. • Additional support for Content Manager: – List, add, retrieve, update, and delete of content class – Asynchronous retrieval of object content *See also* Flickner Depo at 22-25, 36, 54-55, 89-90, 78-79, 159, 201, 209-210

Element	'036 Claim Recitation	Exemplary Citations in Reference
		To the extent it is found that QBIC does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that QBIC does not anticipate this claim, QBIC renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine QBIC with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1C]	a content service coupled with the identification platform, and that:	QBIC discloses a content service coupled with the identification platform. The QBIC Project: Querying Images by Content Using Color, Texture, and Shape (IBM 0002390–404) at 8 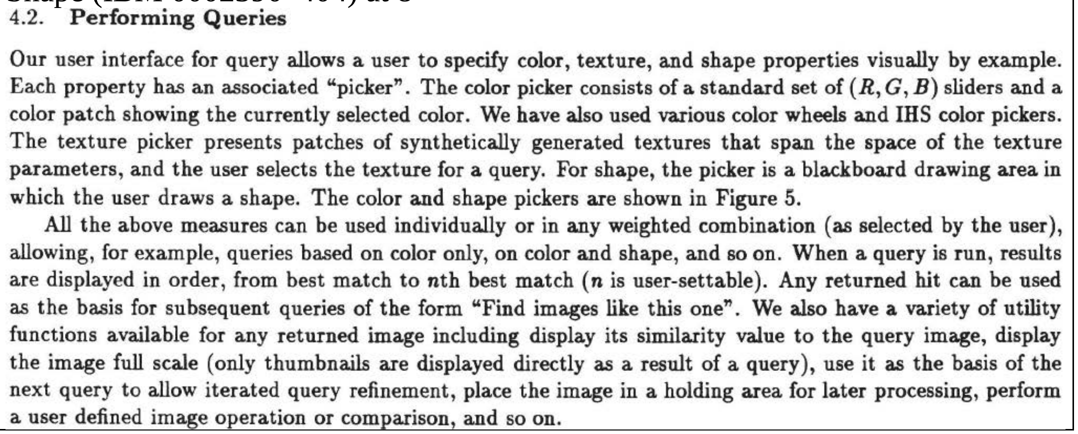 4.2. **Performing Queries** Our user interface for query allows a user to specify color, texture, and shape properties visually by example. Each property has an associated "picker". The color picker consists of a standard set of (R, G, B) sliders and a color patch showing the currently selected color. We have also used various color wheels and IHS color pickers. The texture picker presents patches of synthetically generated textures that span the space of the texture parameters, and the user selects the texture for a query. For shape, the picker is a blackboard drawing area in which the user draws a shape. The color and shape pickers are shown in Figure 5. All the above measures can be used individually or in any weighted combination (as selected by the user), allowing, for example, queries based on color only, on color and shape, and so on. When a query is run, results are displayed in order, from best match to nth best match (n is user-settable). Any returned hit can be used as the basis for subsequent queries of the form "Find images like this one". We also have a variety of utility functions available for any returned image including display its similarity value to the query image, display the image full scale (only thumbnails are displayed directly as a result of a query), use it as the basis of the next query to allow iterated query refinement, place the image in a holding area for later processing, perform a user defined image operation or comparison, and so on.

Element	'036 Claim Recitation	Exemplary Citations in Reference
		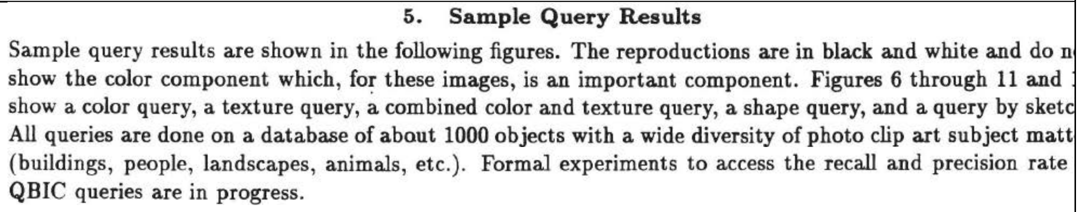 **5. Sample Query Results** Sample query results are shown in the following figures. The reproductions are in black and white and do n show the color component which, for these images, is an important component. Figures 6 through 11 and show a color query, a texture query, a combined color and texture query, a shape query, and a query by sketc All queries are done on a database of about 1000 objects with a wide diversity of photo clip art subject matt (buildings, people, landscapes, animals, etc.). Formal experiments to access the recall and precision rate QBIC queries are in progress. Query by Image and Video Content: The QBIC System (IBM 0000893–902): "To achieve this functionality, QBIC has two main components: database population (the process of creating an image database) and database query. During the population, images and videos are processed to extract features describing their content-colors, textures, shapes, and camera and object motion-and the features are stored in a database. During the query, the user composes a query graphically. Features are generated from the graphical query and then input to a matching engine that finds images or videos from the database with similar features. Figure 2 shows the system architecture." *Id*. at p. 3. "For both population and query, the QBIC data model has • still images or scenes (full images) that contain objects (subsets of an image), and • video shots that consist of sets of contiguous frames and contain motion objects. For still images, the QBIC data model distinguishes between 'scenes' (or images) and 'objects.' A scene is an image or single representative frame of video. An object is a part of a scene-for example, the fox in Figure 3-or a moving entity in a video. For still image database population, features are extracted from images and objects and stored in a database as shown in the top left part of Figure 2. Videos are broken into clips called shots. Representative frames, or r-frames, are generated for each extracted shot. R-frames are treated as still images, and features are extracted and stored in the database. Further processing of shots generates motion objects-for example, a car moving across the screen. Queries are allowed on objects ('Find images with a red, round object'), scenes ('Find images that have approximately 30-percent red and IS-percent blue colors'), shots ('Find all shots

Element	'036 Claim Recitation	Exemplary Citations in Reference
		panning from left to right'), or any combination ('Find images that have 30 percent red and contain a blue textured object')." *Id.* at p. 3. "In still image database population, the images are reduced to a standard-sized icon called a thumbnail and annotated with any available text information. Object identification is an optional but key part of this step. It lets users manually, semiautomatically, or fully automatically identify interesting regions-which we call objects-in the images. . . . For video data, database population has three major components: • shot detection, • representative frame creation for each shot, and • derivation of a layered representation of coherently moving structures/ objects. . . . Once the shot boundaries have been detected, each shot is represented using an r-frame. R-frames are used for several purposes. First, during database population, r-frames are treated as still images in which objects can be identified by using the previously described methods." *Id.* at pp. 3–6. Image, Audio, and Video Extenders Administration and Programming (IBM 000001–617): "To make an image available for searching by content, you catalog the image. When you catalog an image, the Image Extender analyzes the image, by computing the feature values for the image, and stores the values in a QBIC catalog. When you search for an image by content, your query identifies one or more features for the search (such as average color), a source for each feature (such as an example image), and a target set of cataloged images. The Image Extender computes the feature value of the source and compares it to the cataloged feature values for the target images. It then computes a score that indicates how similar the feature values of the target images are to the source. You can have the Image Extender return the images whose features are most similar to the source. The Image Extender will return the handle of each image and the image score. You can also have the Image Extender return only the score of a single image." *Id.* at pp. 20–21.

Element	'036 Claim Recitation	Exemplary Citations in Reference
		QBIC Demo (IBM000747): 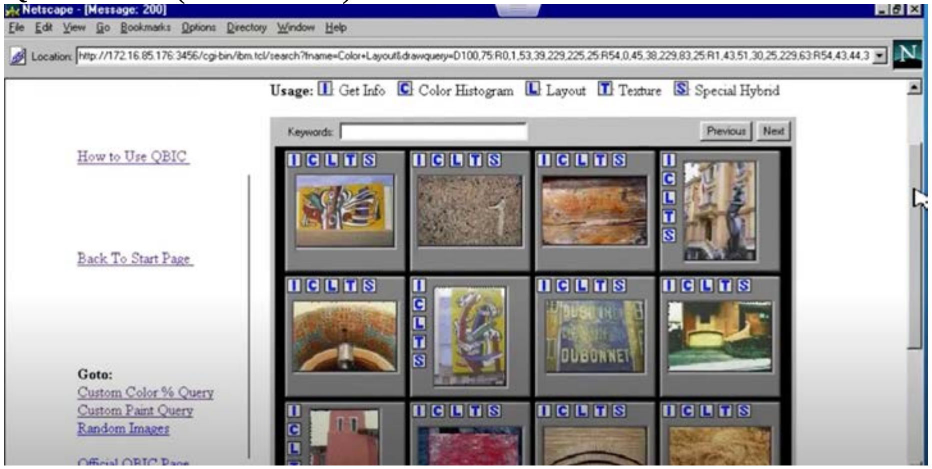 Flickner 8/29/23 Depo at 53 – 56: 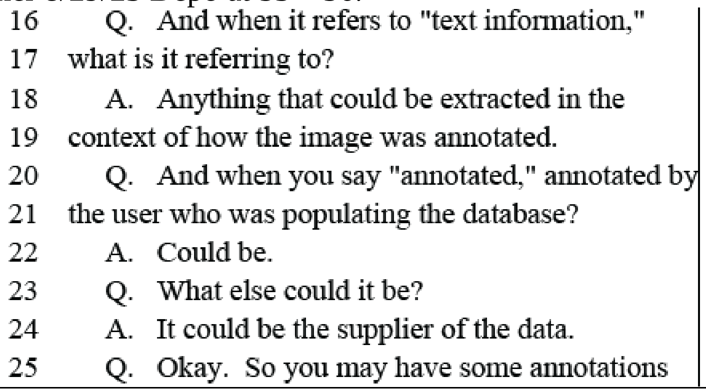 Q. And when it refers to "text information," what is it referring to? A. Anything that could be extracted in the context of how the image was annotated. Q. And when you say "annotated," annotated by the user who was populating the database? A. Could be. Q. What else could it be? A. It could be the supplier of the data. Q. Okay. So you may have some annotations

Element	'036 Claim Recitation	Exemplary Citations in Reference
		See transcript excerpt below

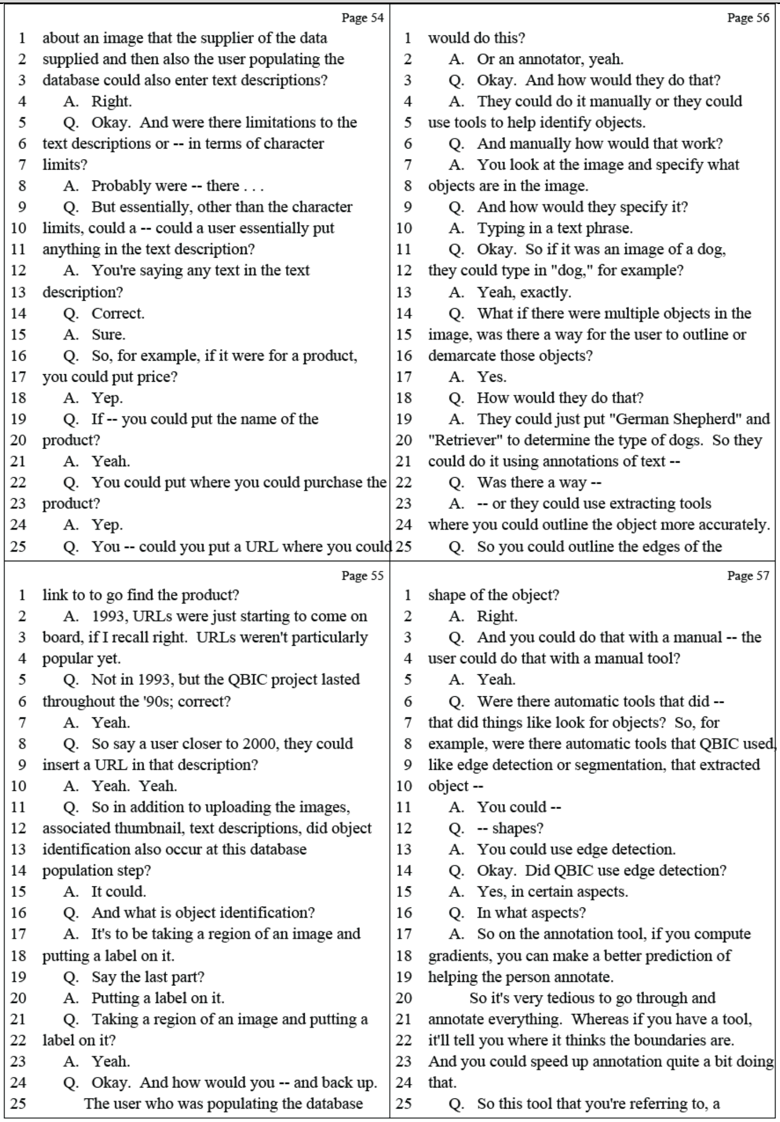

Page 54

about an image that the supplier of the data supplied and then also the user populating the database could also enter text descriptions?
A. Right.
Q. Okay. And were there limitations to the text descriptions or -- in terms of character limits?
A. Probably were -- there . . .
Q. But essentially, other than the character limits, could a -- could a user essentially put anything in the text description?
A. You're saying any text in the text description?
Q. Correct.
A. Sure.
Q. So, for example, if it were for a product, you could put price?
A. Yep.
Q. If -- you could put the name of the product?
A. Yeah.
Q. You could put where you could purchase the product?
A. Yep.
Q. You -- could you put a URL where you could

Page 55

link to to go find the product?
A. 1993, URLs were just starting to come on board, if I recall right. URLs weren't particularly popular yet.
Q. Not in 1993, but the QBIC project lasted throughout the '90s; correct?
A. Yeah.
Q. So say a user closer to 2000, they could insert a URL in that description?
A. Yeah. Yeah.
Q. So in addition to uploading the images, associated thumbnail, text descriptions, did object identification also occur at this database population step?
A. It could.
Q. And what is object identification?
A. It's to be taking a region of an image and putting a label on it.
Q. Say the last part?
A. Putting a label on it.
Q. Taking a region of an image and putting a label on it?
A. Yeah.
Q. Okay. And how would you -- and back up.
The user who was populating the database

Page 56

would do this?
A. Or an annotator, yeah.
Q. Okay. And how would they do that?
A. They could do it manually or they could use tools to help identify objects.
Q. And manually how would that work?
A. You look at the image and specify what objects are in the image.
Q. And how would they specify it?
A. Typing in a text phrase.
Q. Okay. So if it was an image of a dog, they could type in "dog," for example?
A. Yeah, exactly.
Q. What if there were multiple objects in the image, was there a way for the user to outline or demarcate those objects?
A. Yes.
Q. How would they do that?
A. They could just put "German Shepherd" and "Retriever" to determine the type of dogs. So they could do it using annotations of text --
Q. Was there a way --
A. -- or they could use extracting tools where you could outline the object more accurately.
Q. So you could outline the edges of the

Page 57

shape of the object?
A. Right.
Q. And you could do that with a manual -- the user could do that with a manual tool?
A. Yeah.
Q. Were there automatic tools that did -- that did things like look for objects? So, for example, were there automatic tools that QBIC used, like edge detection or segmentation, that extracted object --
A. You could --
Q. -- shapes?
A. You could use edge detection.
Q. Okay. Did QBIC use edge detection?
A. Yes, in certain aspects.
Q. In what aspects?
A. So on the annotation tool, if you compute gradients, you can make a better prediction of helping the person annotate.
So it's very tedious to go through and annotate everything. Whereas if you have a tool, it'll tell you where it thinks the boundaries are. And you could speed up annotation quite a bit doing that.
Q. So this tool that you're referring to, a

Element	'036 Claim Recitation	Exemplary Citations in Reference
		Id. at 60 – 62: 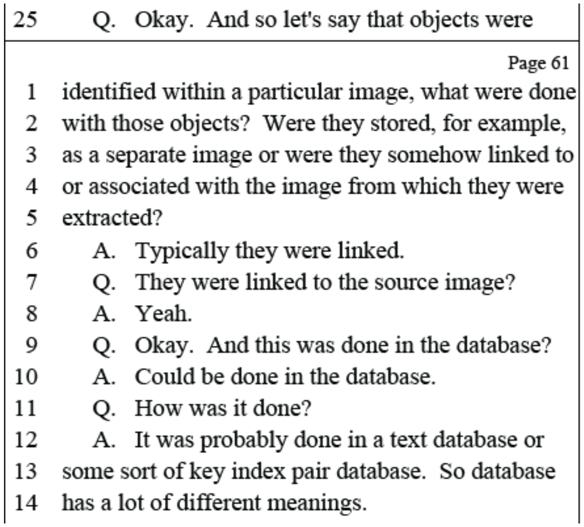 Q. Okay. And so let's say that objects were Page 61 identified within a particular image, what were done with those objects? Were they stored, for example, as a separate image or were they somehow linked to or associated with the image from which they were extracted? A. Typically they were linked. Q. They were linked to the source image? A. Yeah. Q. Okay. And this was done in the database? A. Could be done in the database. Q. How was it done? A. It was probably done in a text database or some sort of key index pair database. So database has a lot of different meanings. 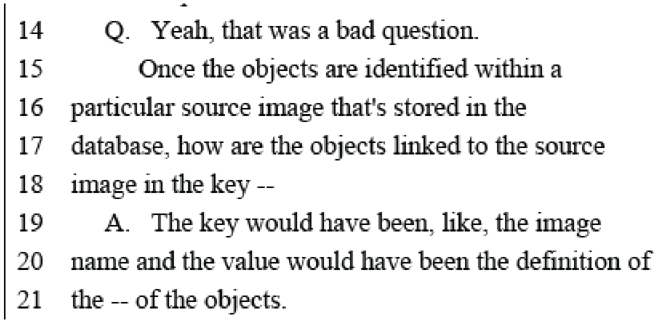 Q. Yeah, that was a bad question. Once the objects are identified within a particular source image that's stored in the database, how are the objects linked to the source image in the key -- A. The key would have been, like, the image name and the value would have been the definition of the -- of the objects. *Id.* at 103:

Element	**'036 Claim Recitation**	**Exemplary Citations in Reference**
		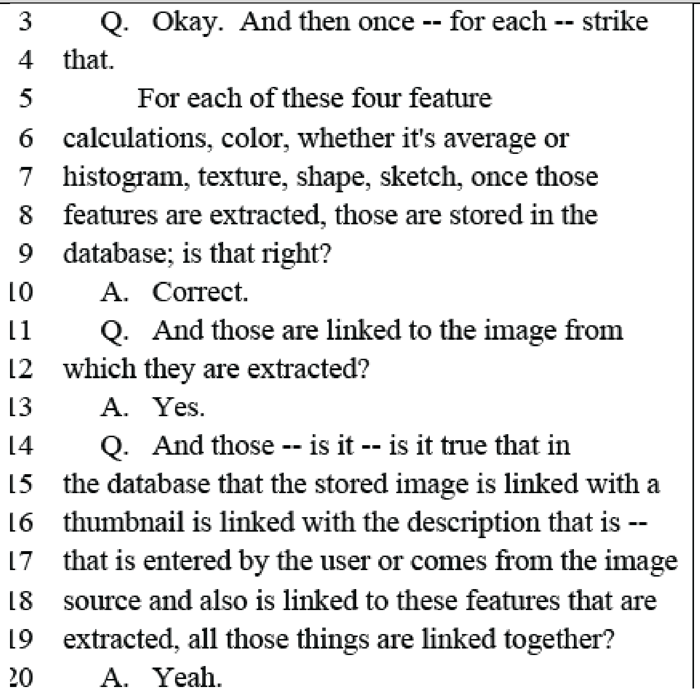 Q. Okay. And then once -- for each -- strike that. For each of these four feature calculations, color, whether it's average or histogram, texture, shape, sketch, once those features are extracted, those are stored in the database; is that right? A. Correct. Q. And those are linked to the image from which they are extracted? A. Yes. Q. And those -- is it -- is it true that in the database that the stored image is linked with a thumbnail is linked with the description that is -- that is entered by the user or comes from the image source and also is linked to these features that are extracted, all those things are linked together? A. Yeah. *Id.* at 129 – 131:

Element	'036 Claim Recitation	Exemplary Citations in Reference
		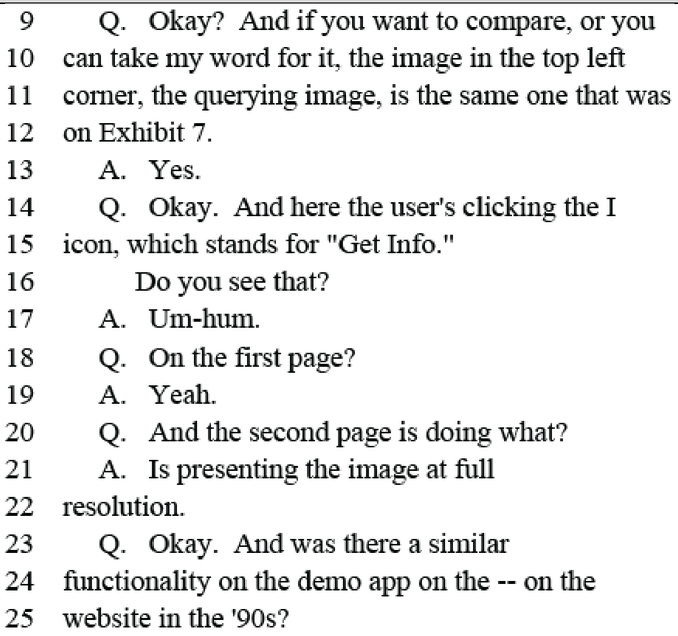 Q. Okay? And if you want to compare, or you can take my word for it, the image in the top left corner, the querying image, is the same one that was on Exhibit 7. A. Yes. Q. Okay. And here the user's clicking the I icon, which stands for "Get Info." Do you see that? A. Um-hum. Q. On the first page? A. Yeah. Q. And the second page is doing what? A. Is presenting the image at full resolution. Q. Okay. And was there a similar functionality on the demo app on the -- on the website in the '90s?

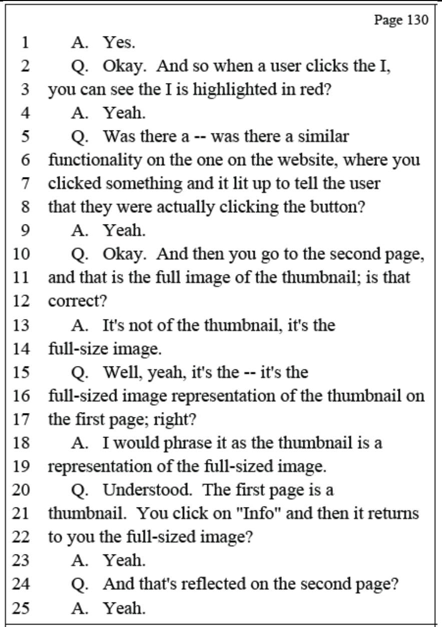

Page 130

A. Yes.
Q. Okay. And so when a user clicks the I, you can see the I is highlighted in red?
A. Yeah.
Q. Was there a -- was there a similar functionality on the one on the website, where you clicked something and it lit up to tell the user that they were actually clicking the button?
A. Yeah.
Q. Okay. And then you go to the second page, and that is the full image of the thumbnail; is that correct?
A. It's not of the thumbnail, it's the full-size image.
Q. Well, yeah, it's the -- it's the full-sized image representation of the thumbnail on the first page; right?
A. I would phrase it as the thumbnail is a representation of the full-sized image.
Q. Understood. The first page is a thumbnail. You click on "Info" and then it returns to you the full-sized image?
A. Yeah.
Q. And that's reflected on the second page?
A. Yeah.

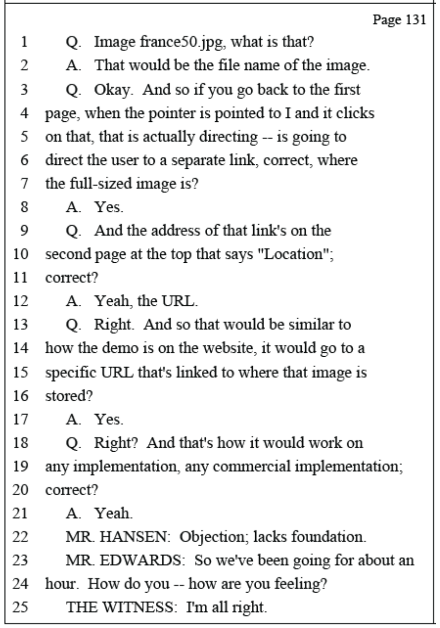

Page 131

Q. Image france50.jpg, what is that?
A. That would be the file name of the image.
Q. Okay. And so if you go back to the first page, when the pointer is pointed to I and it clicks on that, that is actually directing -- is going to direct the user to a separate link, correct, where the full-sized image is?
A. Yes.
Q. And the address of that link's on the second page at the top that says "Location"; correct?
A. Yeah, the URL.
Q. Right. And so that would be similar to how the demo is on the website, it would go to a specific URL that's linked to where that image is stored?
A. Yes.
Q. Right? And that's how it would work on any implementation, any commercial implementation; correct?
A. Yeah.
MR. HANSEN: Objection; lacks foundation.
MR. EDWARDS: So we've been going for about an hour. How do you -- how are you feeling?
THE WITNESS: I'm all right.

Element	’036 Claim Recitation	Exemplary Citations in Reference
		Id. at 161: 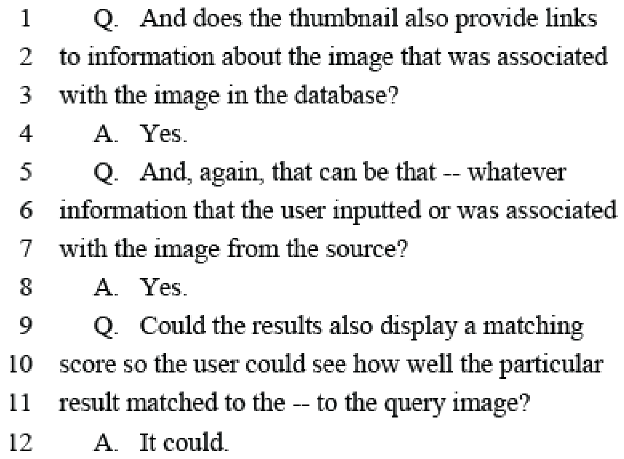 Q. And does the thumbnail also provide links to information about the image that was associated with the image in the database? A. Yes. Q. And, again, that can be that -- whatever information that the user inputted or was associated with the image from the source? A. Yes. Q. Could the results also display a matching score so the user could see how well the particular result matched to the -- to the query image? A. It could. *Id.* at 321: 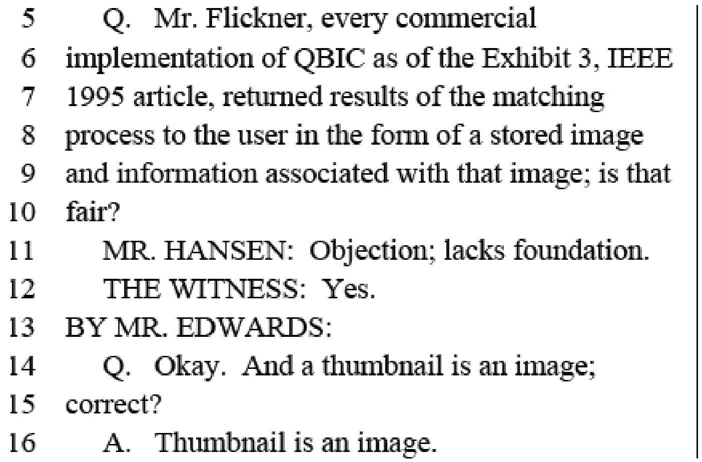 Q. Mr. Flickner, every commercial implementation of QBIC as of the Exhibit 3, IEEE 1995 article, returned results of the matching process to the user in the form of a stored image and information associated with that image; is that fair? MR. HANSEN: Objection; lacks foundation. THE WITNESS: Yes. BY MR. EDWARDS: Q. Okay. And a thumbnail is an image; correct? A. Thumbnail is an image. *Id.* at 305:

Element	’036 Claim Recitation	Exemplary Citations in Reference
		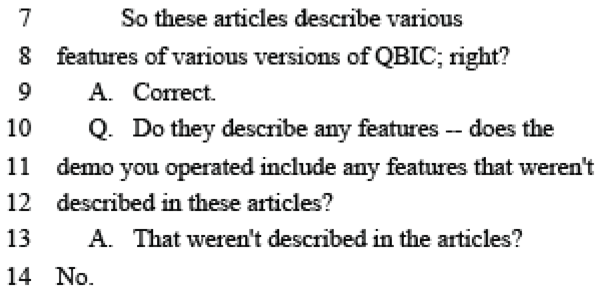 So these articles describe various features of various versions of QBIC; right? A. Correct. Q. Do they describe any features -- does the demo you operated include any features that weren't described in these articles? A. That weren't described in the articles? No. *Id.* at 35: 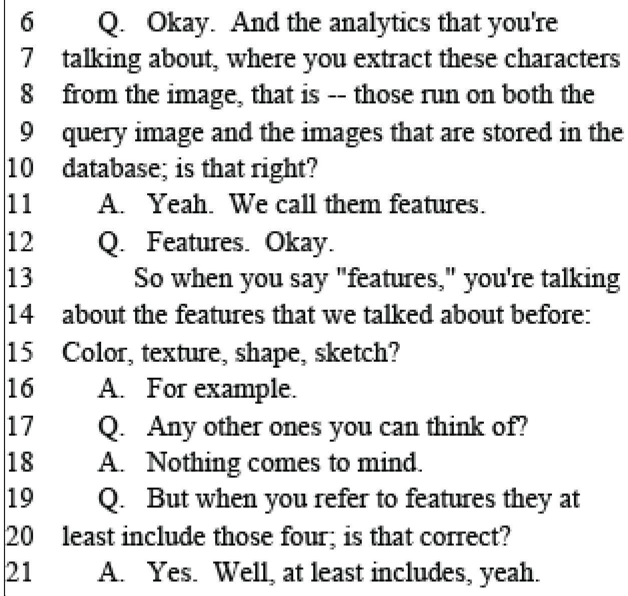 Q. Okay. And the analytics that you're talking about, where you extract these characters from the image, that is -- those run on both the query image and the images that are stored in the database; is that right? A. Yeah. We call them features. Q. Features. Okay. So when you say "features," you're talking about the features that we talked about before: Color, texture, shape, sketch? A. For example. Q. Any other ones you can think of? A. Nothing comes to mind. Q. But when you refer to features they at least include those four; is that correct? A. Yes. Well, at least includes, yeah. *Id.* at 109-110:

Element	'036 Claim Recitation	Exemplary Citations in Reference
		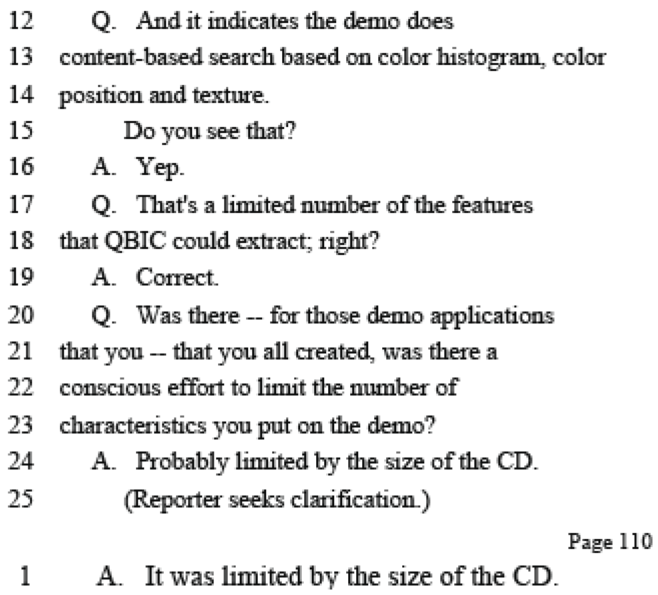 Q. And it indicates the demo does content-based search based on color histogram, color position and texture. Do you see that? A. Yep. Q. That's a limited number of the features that QBIC could extract; right? A. Correct. Q. Was there -- for those demo applications that you -- that you all created, was there a conscious effort to limit the number of characteristics you put on the demo? A. Probably limited by the size of the CD. (Reporter seeks clarification.) Page 110 A. It was limited by the size of the CD. *Id.* at 311: 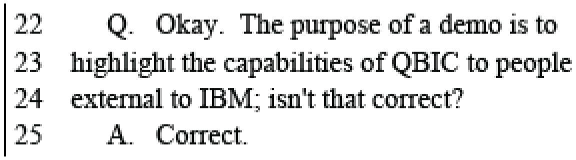 Q. Okay. The purpose of a demo is to highlight the capabilities of QBIC to people external to IBM; isn't that correct? A. Correct. *Id.* at 313-317:

Element	'036 Claim Recitation	Exemplary Citations in Reference
		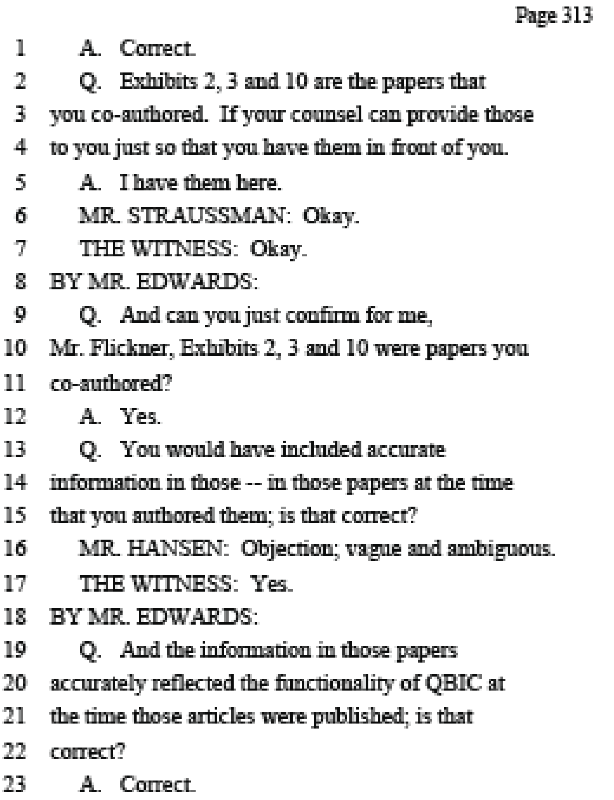 Page 313 A. Correct. Q. Exhibits 2, 3 and 10 are the papers that you co-authored. If your counsel can provide those to you just so that you have them in front of you. A. I have them here. MR. STRAUSSMAN: Okay. THE WITNESS: Okay. BY MR. EDWARDS: Q. And can you just confirm for me, Mr. Flickner, Exhibits 2, 3 and 10 were papers you co-authored? A. Yes. Q. You would have included accurate information in those -- in those papers at the time that you authored them; is that correct? MR. HANSEN: Objection; vague and ambiguous. THE WITNESS: Yes. BY MR. EDWARDS: Q. And the information in those papers accurately reflected the functionality of QBIC at the time those articles were published; is that correct? A. Correct.

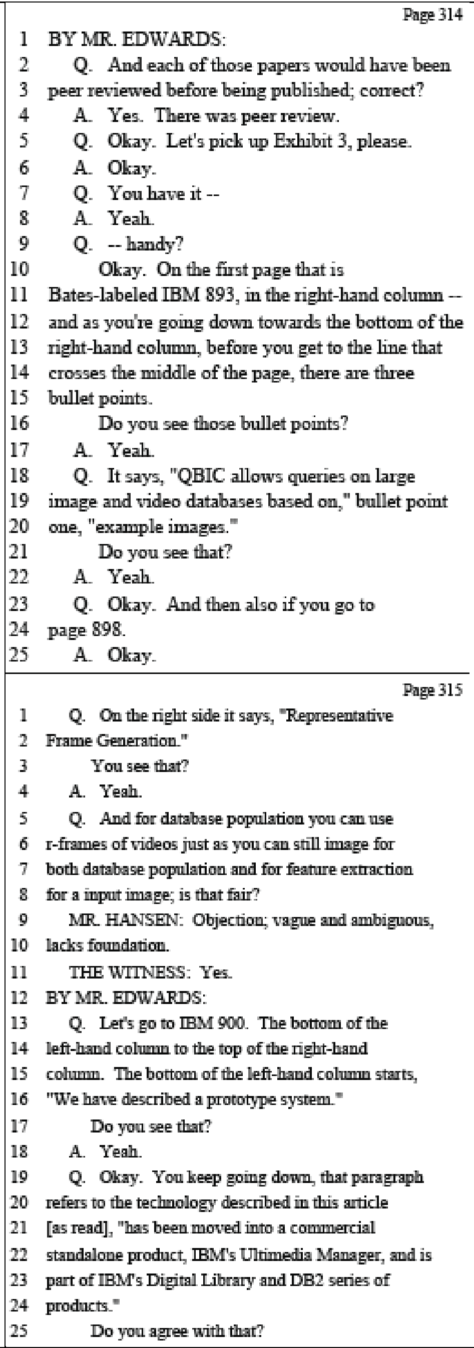

Page 314

BY MR. EDWARDS:
Q. And each of those papers would have been peer reviewed before being published; correct?
A. Yes. There was peer review.
Q. Okay. Let's pick up Exhibit 3, please.
A. Okay.
Q. You have it --
A. Yeah.
Q. -- handy?
Okay. On the first page that is Bates-labeled IBM 893, in the right-hand column -- and as you're going down towards the bottom of the right-hand column, before you get to the line that crosses the middle of the page, there are three bullet points.
Do you see those bullet points?
A. Yeah.
Q. It says, "QBIC allows queries on large image and video databases based on," bullet point one, "example images."
Do you see that?
A. Yeah.
Q. Okay. And then also if you go to page 898.
A. Okay.

Page 315

Q. On the right side it says, "Representative Frame Generation."
You see that?
A. Yeah.
Q. And for database population you can use r-frames of videos just as you can still image for both database population and for feature extraction for a input image; is that fair?
MR. HANSEN: Objection; vague and ambiguous, lacks foundation.
THE WITNESS: Yes.
BY MR. EDWARDS:
Q. Let's go to IBM 900. The bottom of the left-hand column to the top of the right-hand column. The bottom of the left-hand column starts, "We have described a prototype system."
Do you see that?
A. Yeah.
Q. Okay. You keep going down, that paragraph refers to the technology described in this article [as read], "has been moved into a commercial standalone product, IBM's Ultimedia Manager, and is part of IBM's Digital Library and DB2 series of products."
Do you agree with that?

Element	'036 Claim Recitation	Exemplary Citations in Reference
		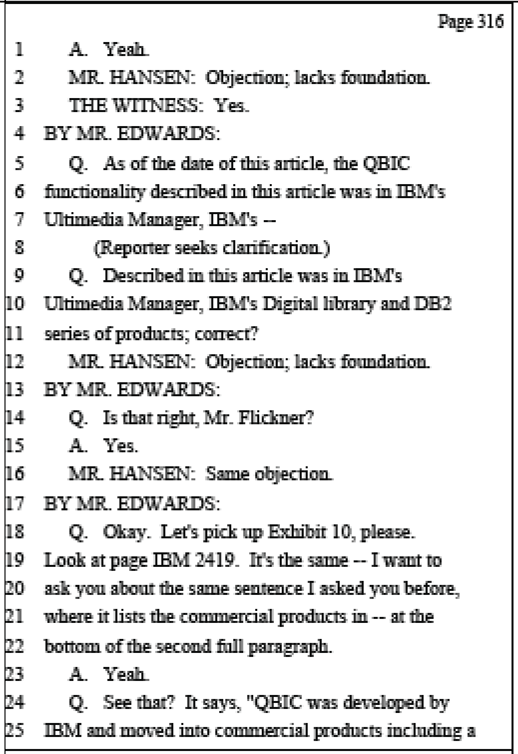 Page 316 A. Yeah. MR. HANSEN: Objection; lacks foundation. THE WITNESS: Yes. BY MR. EDWARDS: Q. As of the date of this article, the QBIC functionality described in this article was in IBM's Ultimedia Manager, IBM's -- (Reporter seeks clarification.) Q. Described in this article was in IBM's Ultimedia Manager, IBM's Digital library and DB2 series of products; correct? MR. HANSEN: Objection; lacks foundation. BY MR. EDWARDS: Q. Is that right, Mr. Flickner? A. Yes. MR. HANSEN: Same objection. BY MR. EDWARDS: Q. Okay. Let's pick up Exhibit 10, please. Look at page IBM 2419. It's the same -- I want to ask you about the same sentence I asked you before, where it lists the commercial products in -- at the bottom of the second full paragraph. A. Yeah. Q. See that? It says, "QBIC was developed by IBM and moved into commercial products including a 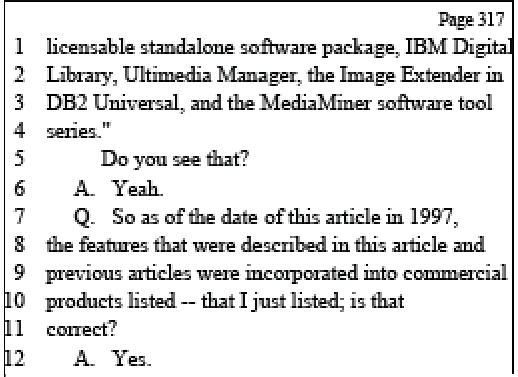 Page 317 licensable standalone software package, IBM Digital Library, Ultimedia Manager, the Image Extender in DB2 Universal, and the MediaMiner software tool series." Do you see that? A. Yeah. Q. So as of the date of this article in 1997, the features that were described in this article and previous articles were incorporated into commercial products listed -- that I just listed; is that correct? A. Yes.

Element	'036 Claim Recitation	Exemplary Citations in Reference
		See also Flickner Depo at 40, 50, 54-55, 72, 76, 89-90, 91, 95, 113; Flickner 8/29/23 Depo at 62–63, 74, 166–174, 180–181, 189, 193, 202, 316–317 To the extent it is found that QBIC does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that QBIC does not anticipate this claim, QBIC renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine QBIC with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1C.i]	obtains content information related to the known target; and	QBIC discloses a content service coupled with the identification platform] that obtains content information related to the known target. The QBIC Project: Querying Images by Content Using Color, Texture, and Shape (IBM 0002390–404) at 8

Element	'036 Claim Recitation	Exemplary Citations in Reference
		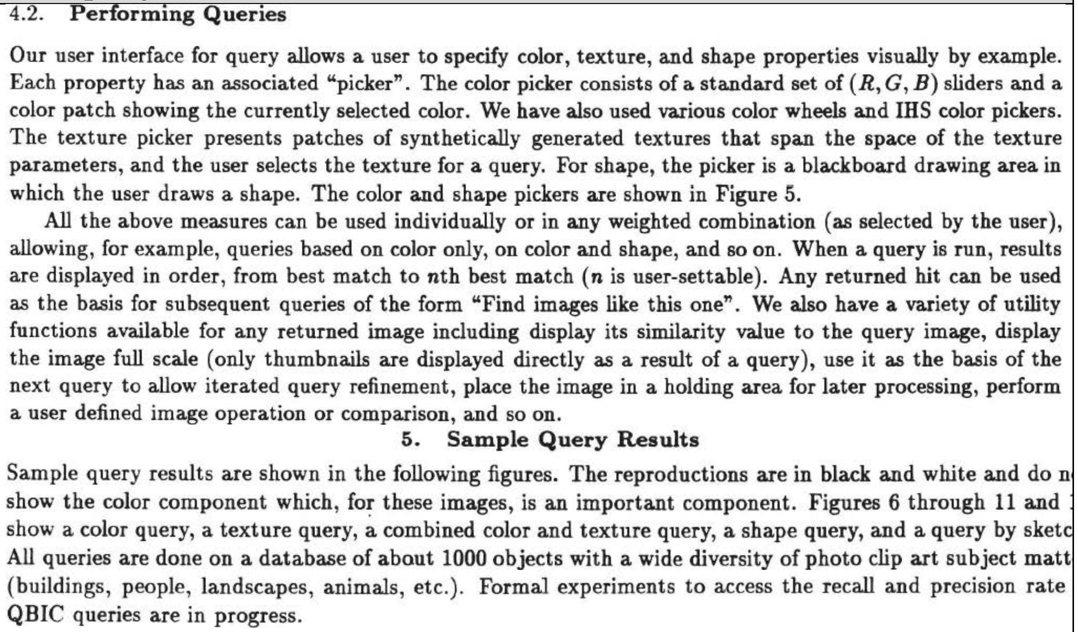 4.2. **Performing Queries** Our user interface for query allows a user to specify color, texture, and shape properties visually by example. Each property has an associated "picker". The color picker consists of a standard set of (R, G, B) sliders and a color patch showing the currently selected color. We have also used various color wheels and IHS color pickers. The texture picker presents patches of synthetically generated textures that span the space of the texture parameters, and the user selects the texture for a query. For shape, the picker is a blackboard drawing area in which the user draws a shape. The color and shape pickers are shown in Figure 5. All the above measures can be used individually or in any weighted combination (as selected by the user), allowing, for example, queries based on color only, on color and shape, and so on. When a query is run, results are displayed in order, from best match to nth best match (n is user-settable). Any returned hit can be used as the basis for subsequent queries of the form "Find images like this one". We also have a variety of utility functions available for any returned image including display its similarity value to the query image, display the image full scale (only thumbnails are displayed directly as a result of a query), use it as the basis of the next query to allow iterated query refinement, place the image in a holding area for later processing, perform a user defined image operation or comparison, and so on. 5. **Sample Query Results** Sample query results are shown in the following figures. The reproductions are in black and white and do n show the color component which, for these images, is an important component. Figures 6 through 11 and show a color query, a texture query, a combined color and texture query, a shape query, and a query by sketc All queries are done on a database of about 1000 objects with a wide diversity of photo clip art subject matt (buildings, people, landscapes, animals, etc.). Formal experiments to access the recall and precision rate QBIC queries are in progress. The QBIC Project: Querying Images by Content Using Color, Texture, and Shape (IBM 0002390–404) at Figs. 7–9

Element	'036 Claim Recitation	Exemplary Citations in Reference
		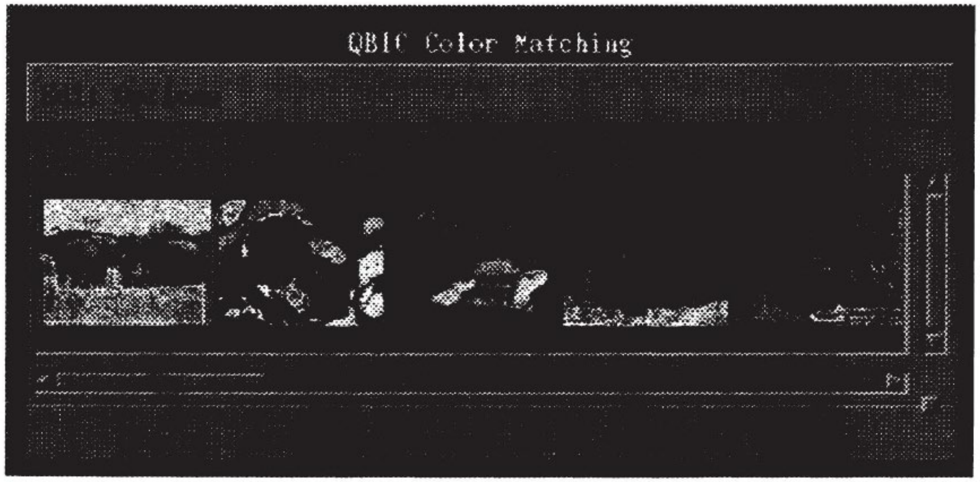 Figure 7: Result of query by example using color only.

Element	'036 Claim Recitation	Exemplary Citations in Reference
		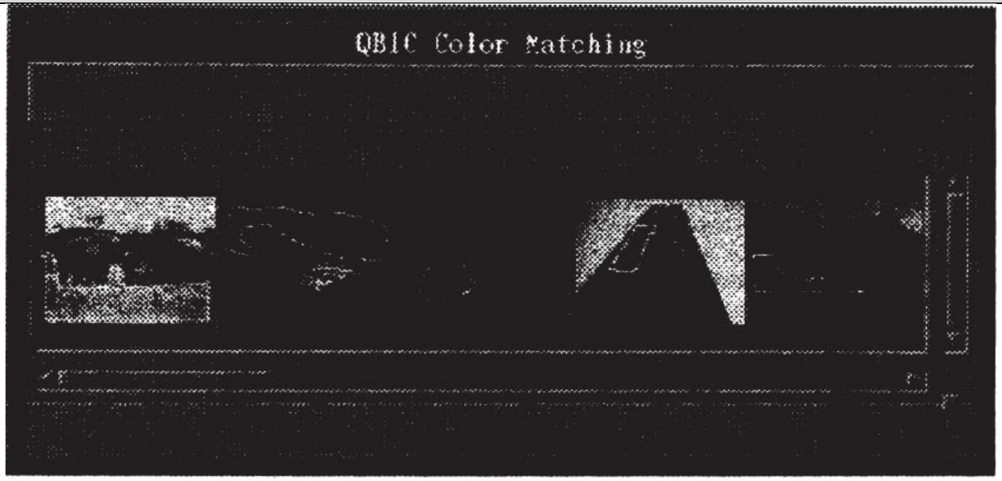 Figure 8: Result of query by example using texture only.

Element	'036 Claim Recitation	Exemplary Citations in Reference
		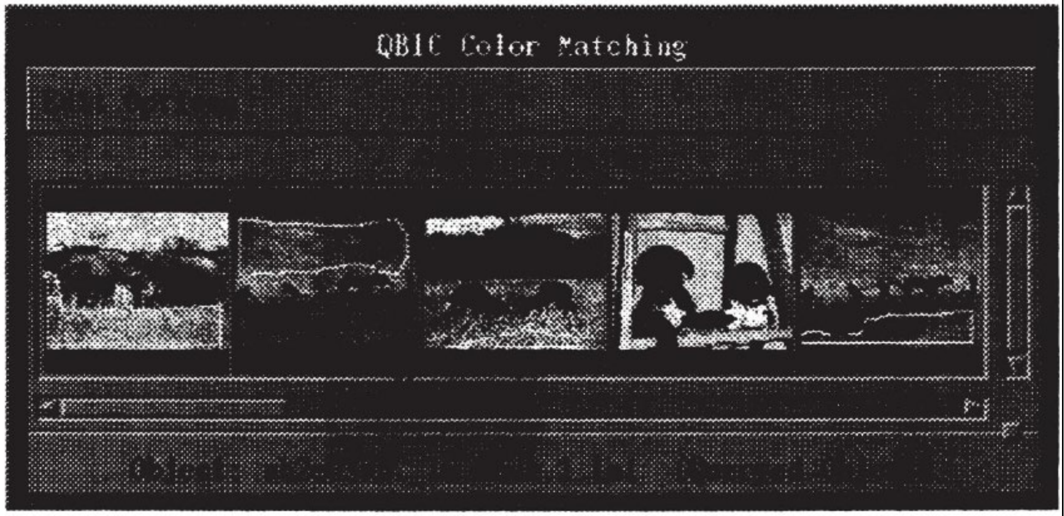 Figure 9: Result of query by example using color and texture. Query by Image and Video Content: The QBIC System (IBM 0000893–902): "To achieve this functionality, QBIC has two main components: database population (the process of creating an image database) and database query. During the population, images and videos are processed to extract features describing their content-colors, textures, shapes, and camera and object motion-and the features are stored in a database. During the query, the user composes a query graphically. Features are generated from the graphical query and then input to a matching engine that finds images or videos from the database with similar features. Figure 2 shows the system architecture." *Id*. at p. 3. "For both population and query, the QBIC data model has • still images or scenes (full images) that contain objects (subsets of an image), and • video shots that consist of sets of contiguous frames and contain motion objects. For still images, the QBIC data model distinguishes between 'scenes' (or images) and 'objects.' A scene is an image or single representative frame of video. An

Element	'036 Claim Recitation	Exemplary Citations in Reference
		object is a part of a scene-for example, the fox in Figure 3-or a moving entity in a video. For still image database population, features are extracted from images and objects and stored in a database as shown in the top left part of Figure 2. Videos are broken into clips called shots. Representative frames, or r-frames, are generated for each extracted shot. R-frames are treated as still images, and features are extracted and stored in the database. Further processing of shots generates motion objects-for example, a car moving across the screen. Queries are allowed on objects ('Find images with a red, round object'), scenes ('Find images that have approximately 30-percent red and IS-percent blue colors'), shots ('Find all shots panning from left to right'), or any combination ('Find images that have 30 percent red and contain a blue textured object')." *Id.* at p. 3. "In still image database population, the images are reduced to a standard-sized icon called a thumbnail and annotated with any available text information. Object identification is an optional but key part of this step. It lets users manually, semiautomatically, or fully automatically identify interesting regions-which we call objects-in the images. . . . For video data, database population has three major components: • shot detection, • representative frame creation for each shot, and • derivation of a layered representation of coherently moving structures/ objects. . . . Once the shot boundaries have been detected, each shot is represented using an r-frame. R-frames are used for several purposes. First, during database population, r-frames are treated as still images in which objects can be identified by using the previously described methods." *Id.* at pp. 3–6. Image, Audio, and Video Extenders Administration and Programming (IBM 000001–617): "To make an image available for searching by content, you catalog the image. When you catalog an image, the Image Extender analyzes the image, by computing the feature values for the image, and stores the values in a QBIC catalog. When you search for an image by content, your query identifies one or more

Element	'036 Claim Recitation	Exemplary Citations in Reference
		features for the search (such as average color), a source for each feature (such as an example image), and a target set of cataloged images. The Image Extender computes the feature value of the source and compares it to the cataloged feature values for the target images. It then computes a score that indicates how similar the feature values of the target images are to the source. You can have the Image Extender return the images whose features are most similar to the source. The Image Extender will return the handle of each image and the image score. You can also have the Image Extender return only the score of a single image." *Id.* at pp. 20–21. QBIC Demo (IBM000747): 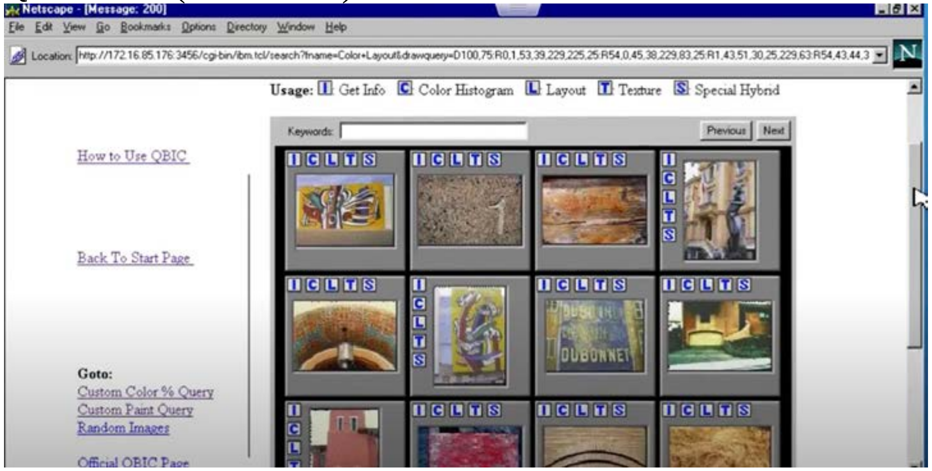

Element	'036 Claim Recitation	Exemplary Citations in Reference
		QBIC Demo (IBM000747):  Location: http://172.16.85.176:3456/cgi-bin/ibm.tcl/info?france50%2Ejpg Image france50.jpg Flickner Depo at 88: 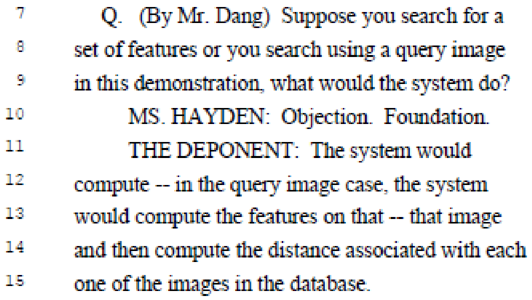 Q. (By Mr. Dang) Suppose you search for a set of features or you search using a query image in this demonstration, what would the system do? MS. HAYDEN: Objection. Foundation. THE DEPONENT: The system would compute -- in the query image case, the system would compute the features on that -- that image and then compute the distance associated with each one of the images in the database. Flickner 8/29/23 Depo at 53 – 56:

Element	'036 Claim Recitation	Exemplary Citations in Reference
		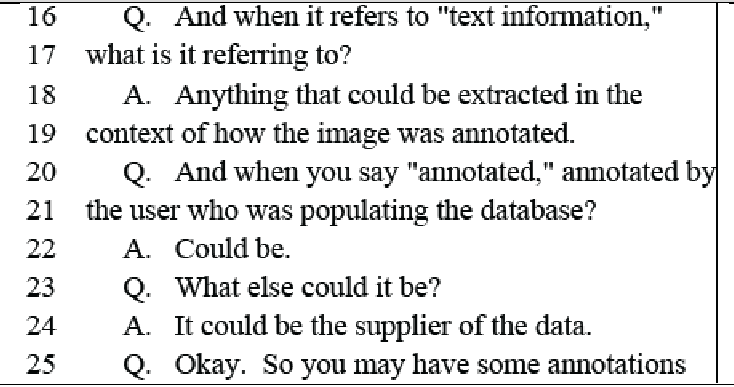 Q. And when it refers to "text information," what is it referring to? A. Anything that could be extracted in the context of how the image was annotated. Q. And when you say "annotated," annotated by the user who was populating the database? A. Could be. Q. What else could it be? A. It could be the supplier of the data. Q. Okay. So you may have some annotations

Element	'036 Claim Recitation	Exemplary Citations in Reference
		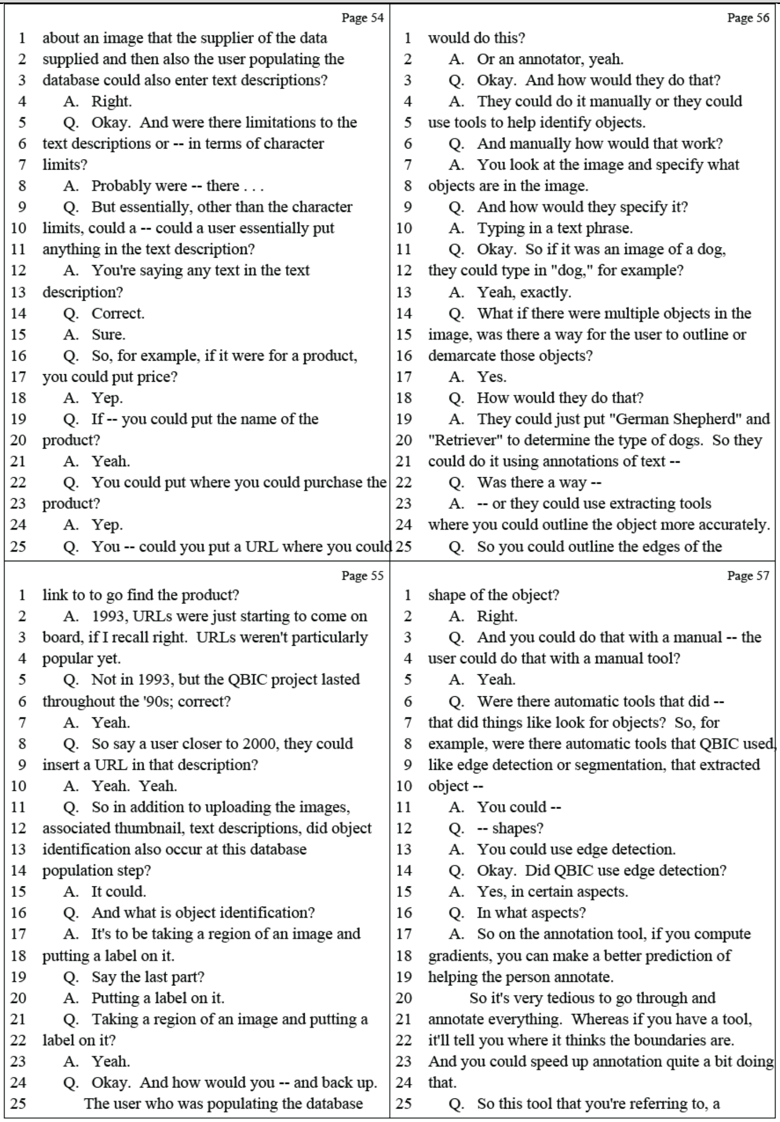 Page 54 about an image that the supplier of the data supplied and then also the user populating the database could also enter text descriptions? A. Right. Q. Okay. And were there limitations to the text descriptions or -- in terms of character limits? A. Probably were -- there . . . Q. But essentially, other than the character limits, could a -- could a user essentially put anything in the text description? A. You're saying any text in the text description? Q. Correct. A. Sure. Q. So, for example, if it were for a product, you could put price? A. Yep. Q. If -- you could put the name of the product? A. Yeah. Q. You could put where you could purchase the product? A. Yep. Q. You -- could you put a URL where you could Page 55 link to to go find the product? A. 1993, URLs were just starting to come on board, if I recall right. URLs weren't particularly popular yet. Q. Not in 1993, but the QBIC project lasted throughout the '90s; correct? A. Yeah. Q. So say a user closer to 2000, they could insert a URL in that description? A. Yeah. Yeah. Q. So in addition to uploading the images, associated thumbnail, text descriptions, did object identification also occur at this database population step? A. It could. Q. And what is object identification? A. It's to be taking a region of an image and putting a label on it. Q. Say the last part? A. Putting a label on it. Q. Taking a region of an image and putting a label on it? A. Yeah. Q. Okay. And how would you -- and back up. The user who was populating the database Page 56 would do this? A. Or an annotator, yeah. Q. Okay. And how would they do that? A. They could do it manually or they could use tools to help identify objects. Q. And manually how would that work? A. You look at the image and specify what objects are in the image. Q. And how would they specify it? A. Typing in a text phrase. Q. Okay. So if it was an image of a dog, they could type in "dog," for example? A. Yeah, exactly. Q. What if there were multiple objects in the image, was there a way for the user to outline or demarcate those objects? A. Yes. Q. How would they do that? A. They could just put "German Shepherd" and "Retriever" to determine the type of dogs. So they could do it using annotations of text -- Q. Was there a way -- A. -- or they could use extracting tools where you could outline the object more accurately. Q. So you could outline the edges of the Page 57 shape of the object? A. Right. Q. And you could do that with a manual -- the user could do that with a manual tool? A. Yeah. Q. Were there automatic tools that did -- that did things like look for objects? So, for example, were there automatic tools that QBIC used, like edge detection or segmentation, that extracted object -- A. You could -- Q. -- shapes? A. You could use edge detection. Q. Okay. Did QBIC use edge detection? A. Yes, in certain aspects. Q. In what aspects? A. So on the annotation tool, if you compute gradients, you can make a better prediction of helping the person annotate. So it's very tedious to go through and annotate everything. Whereas if you have a tool, it'll tell you where it thinks the boundaries are. And you could speed up annotation quite a bit doing that. Q. So this tool that you're referring to, a

Element	'036 Claim Recitation	Exemplary Citations in Reference
		Id. at 60 – 62: 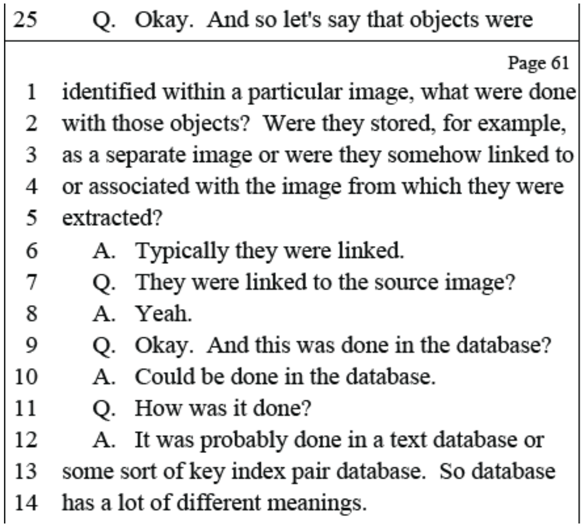 Q. Okay. And so let's say that objects were Page 61 identified within a particular image, what were done with those objects? Were they stored, for example, as a separate image or were they somehow linked to or associated with the image from which they were extracted? A. Typically they were linked. Q. They were linked to the source image? A. Yeah. Q. Okay. And this was done in the database? A. Could be done in the database. Q. How was it done? A. It was probably done in a text database or some sort of key index pair database. So database has a lot of different meanings. 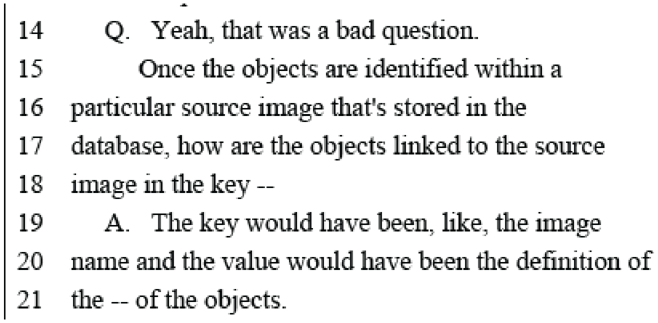 Q. Yeah, that was a bad question. Once the objects are identified within a particular source image that's stored in the database, how are the objects linked to the source image in the key -- A. The key would have been, like, the image name and the value would have been the definition of the -- of the objects. *Id.* at 103:

Element	'036 Claim Recitation	Exemplary Citations in Reference
		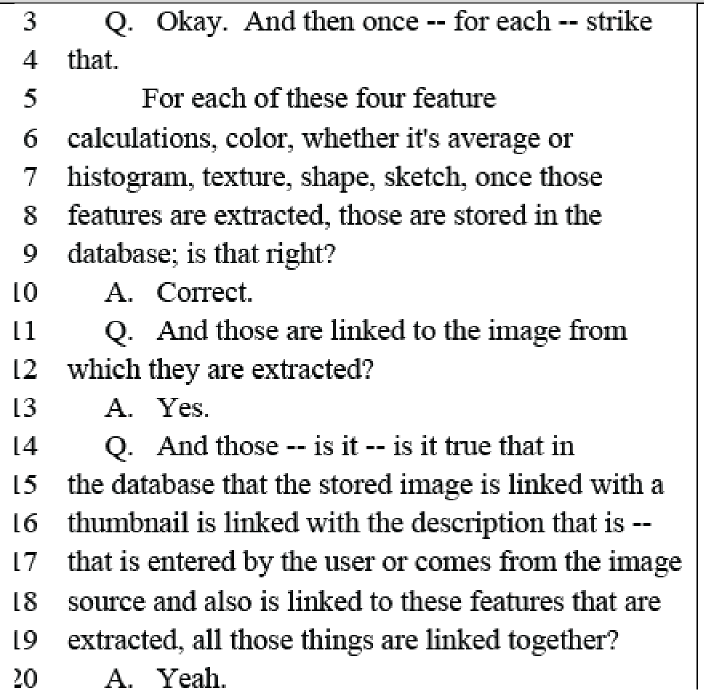 Q. Okay. And then once -- for each -- strike that. For each of these four feature calculations, color, whether it's average or histogram, texture, shape, sketch, once those features are extracted, those are stored in the database; is that right? A. Correct. Q. And those are linked to the image from which they are extracted? A. Yes. Q. And those -- is it -- is it true that in the database that the stored image is linked with a thumbnail is linked with the description that is -- that is entered by the user or comes from the image source and also is linked to these features that are extracted, all those things are linked together? A. Yeah. *Id.* at 129 – 131:

Element	'036 Claim Recitation	Exemplary Citations in Reference
		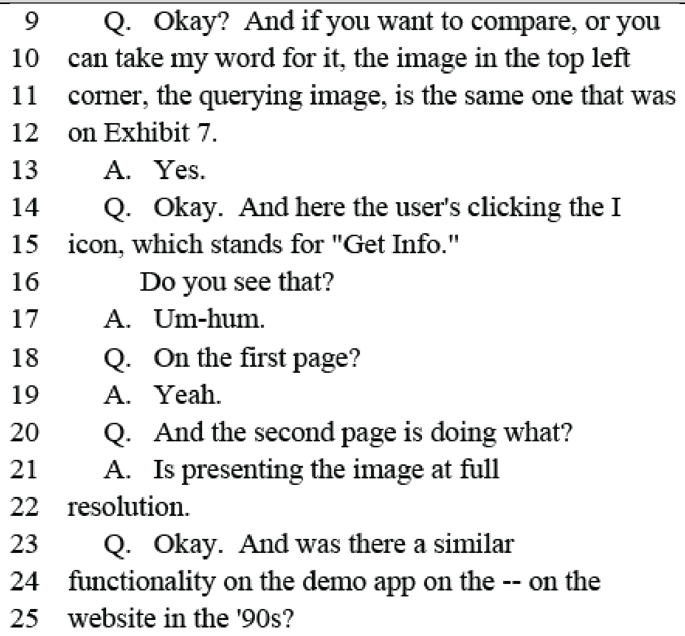 Q. Okay? And if you want to compare, or you can take my word for it, the image in the top left corner, the querying image, is the same one that was on Exhibit 7. A. Yes. Q. Okay. And here the user's clicking the I icon, which stands for "Get Info." Do you see that? A. Um-hum. Q. On the first page? A. Yeah. Q. And the second page is doing what? A. Is presenting the image at full resolution. Q. Okay. And was there a similar functionality on the demo app on the -- on the website in the '90s?

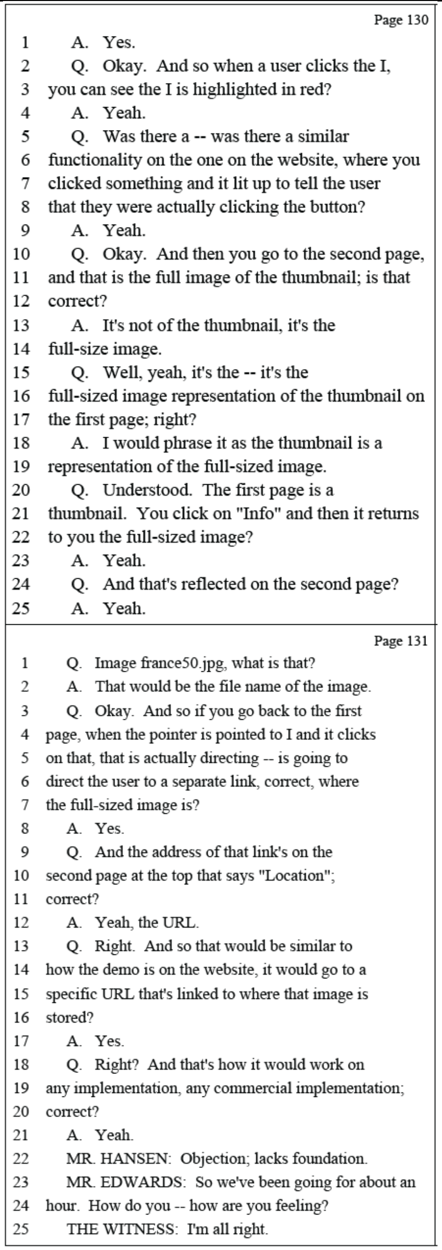

Page 130

A. Yes.
Q. Okay. And so when a user clicks the I, you can see the I is highlighted in red?
A. Yeah.
Q. Was there a -- was there a similar functionality on the one on the website, where you clicked something and it lit up to tell the user that they were actually clicking the button?
A. Yeah.
Q. Okay. And then you go to the second page, and that is the full image of the thumbnail; is that correct?
A. It's not of the thumbnail, it's the full-size image.
Q. Well, yeah, it's the -- it's the full-sized image representation of the thumbnail on the first page; right?
A. I would phrase it as the thumbnail is a representation of the full-sized image.
Q. Understood. The first page is a thumbnail. You click on "Info" and then it returns to you the full-sized image?
A. Yeah.
Q. And that's reflected on the second page?
A. Yeah.

Page 131

Q. Image france50.jpg, what is that?
A. That would be the file name of the image.
Q. Okay. And so if you go back to the first page, when the pointer is pointed to I and it clicks on that, that is actually directing -- is going to direct the user to a separate link, correct, where the full-sized image is?
A. Yes.
Q. And the address of that link's on the second page at the top that says "Location"; correct?
A. Yeah, the URL.
Q. Right. And so that would be similar to how the demo is on the website, it would go to a specific URL that's linked to where that image is stored?
A. Yes.
Q. Right? And that's how it would work on any implementation, any commercial implementation; correct?
A. Yeah.
MR. HANSEN: Objection; lacks foundation.
MR. EDWARDS: So we've been going for about an hour. How do you -- how are you feeling?
THE WITNESS: I'm all right.

Element	'036 Claim Recitation	Exemplary Citations in Reference
		Id. at 161: 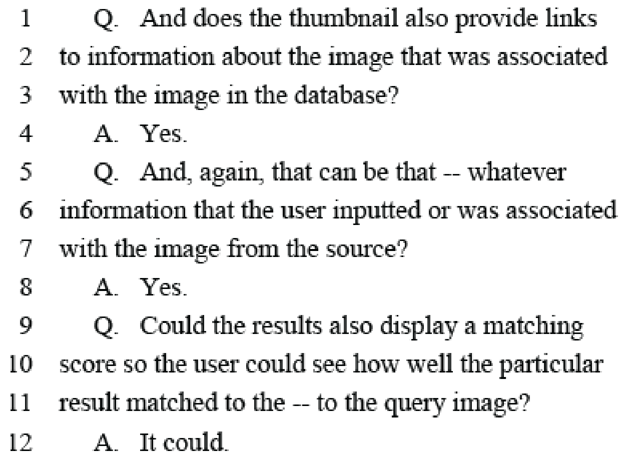 Q. And does the thumbnail also provide links to information about the image that was associated with the image in the database? A. Yes. Q. And, again, that can be that -- whatever information that the user inputted or was associated with the image from the source? A. Yes. Q. Could the results also display a matching score so the user could see how well the particular result matched to the -- to the query image? A. It could. *Id.* at 321: 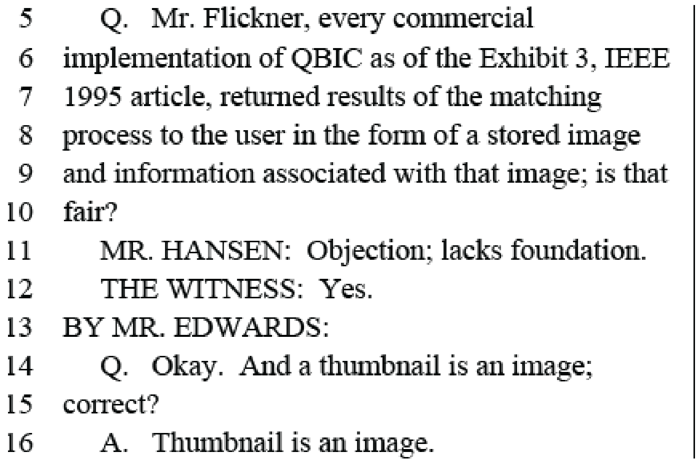 Q. Mr. Flickner, every commercial implementation of QBIC as of the Exhibit 3, IEEE 1995 article, returned results of the matching process to the user in the form of a stored image and information associated with that image; is that fair? MR. HANSEN: Objection; lacks foundation. THE WITNESS: Yes. BY MR. EDWARDS: Q. Okay. And a thumbnail is an image; correct? A. Thumbnail is an image. *Id.* at 305:

Element	'036 Claim Recitation	Exemplary Citations in Reference
		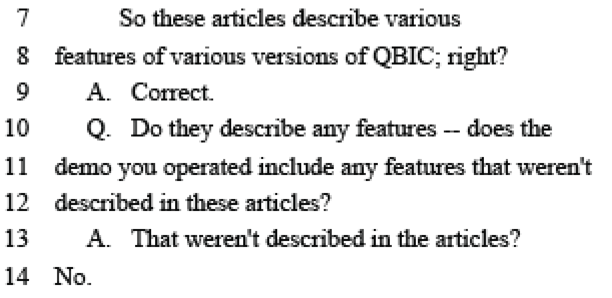 So these articles describe various features of various versions of QBIC; right? A. Correct. Q. Do they describe any features -- does the demo you operated include any features that weren't described in these articles? A. That weren't described in the articles? No. *Id.* at 35: 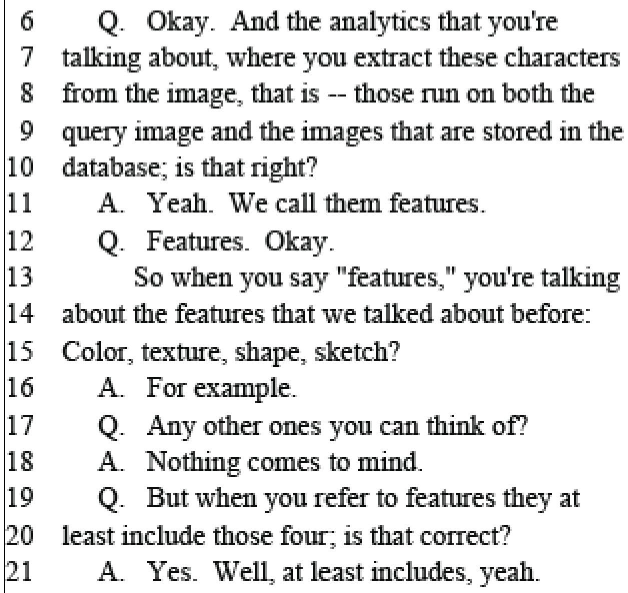 Q. Okay. And the analytics that you're talking about, where you extract these characters from the image, that is -- those run on both the query image and the images that are stored in the database; is that right? A. Yeah. We call them features. Q. Features. Okay. So when you say "features," you're talking about the features that we talked about before: Color, texture, shape, sketch? A. For example. Q. Any other ones you can think of? A. Nothing comes to mind. Q. But when you refer to features they at least include those four; is that correct? A. Yes. Well, at least includes, yeah. *Id.* at 109-110:

Element	'036 Claim Recitation	Exemplary Citations in Reference
		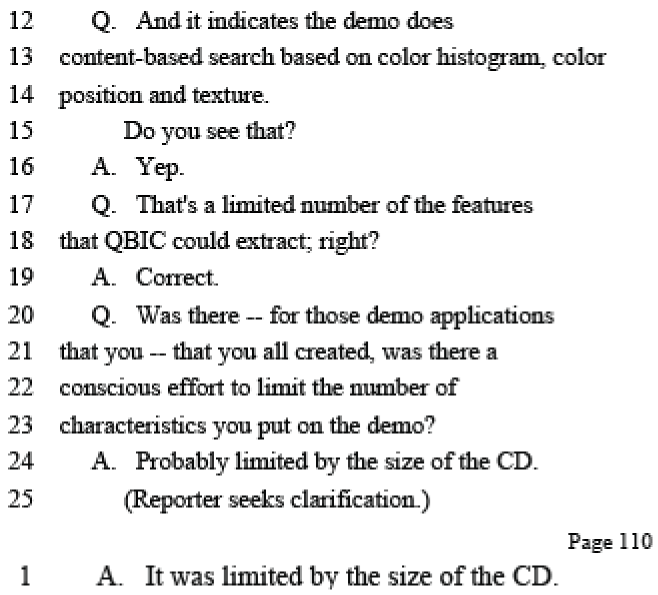 Q. And it indicates the demo does content-based search based on color histogram, color position and texture. Do you see that? A. Yep. Q. That's a limited number of the features that QBIC could extract; right? A. Correct. Q. Was there -- for those demo applications that you -- that you all created, was there a conscious effort to limit the number of characteristics you put on the demo? A. Probably limited by the size of the CD. (Reporter seeks clarification.) Page 110 A. It was limited by the size of the CD. *Id.* at 311: 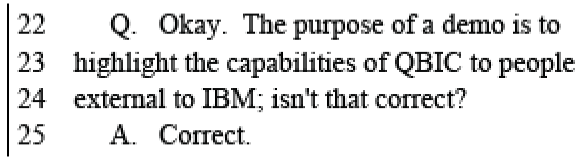 Q. Okay. The purpose of a demo is to highlight the capabilities of QBIC to people external to IBM; isn't that correct? A. Correct. *Id.* at 313-317:

Element	'036 Claim Recitation	Exemplary Citations in Reference
		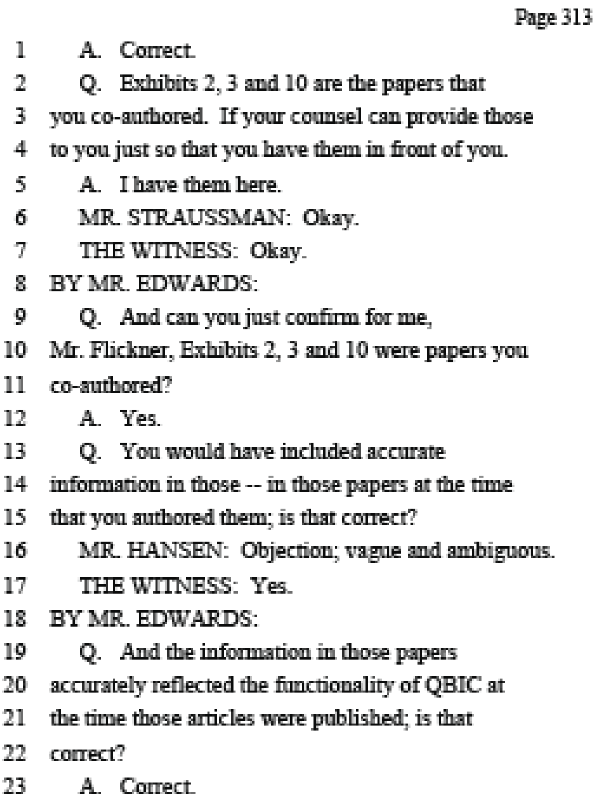 Page 313 A. Correct. Q. Exhibits 2, 3 and 10 are the papers that you co-authored. If your counsel can provide those to you just so that you have them in front of you. A. I have them here. MR. STRAUSSMAN: Okay. THE WITNESS: Okay. BY MR. EDWARDS: Q. And can you just confirm for me, Mr. Flickner, Exhibits 2, 3 and 10 were papers you co-authored? A. Yes. Q. You would have included accurate information in those -- in those papers at the time that you authored them; is that correct? MR. HANSEN: Objection; vague and ambiguous. THE WITNESS: Yes. BY MR. EDWARDS: Q. And the information in those papers accurately reflected the functionality of QBIC at the time those articles were published; is that correct? A. Correct.

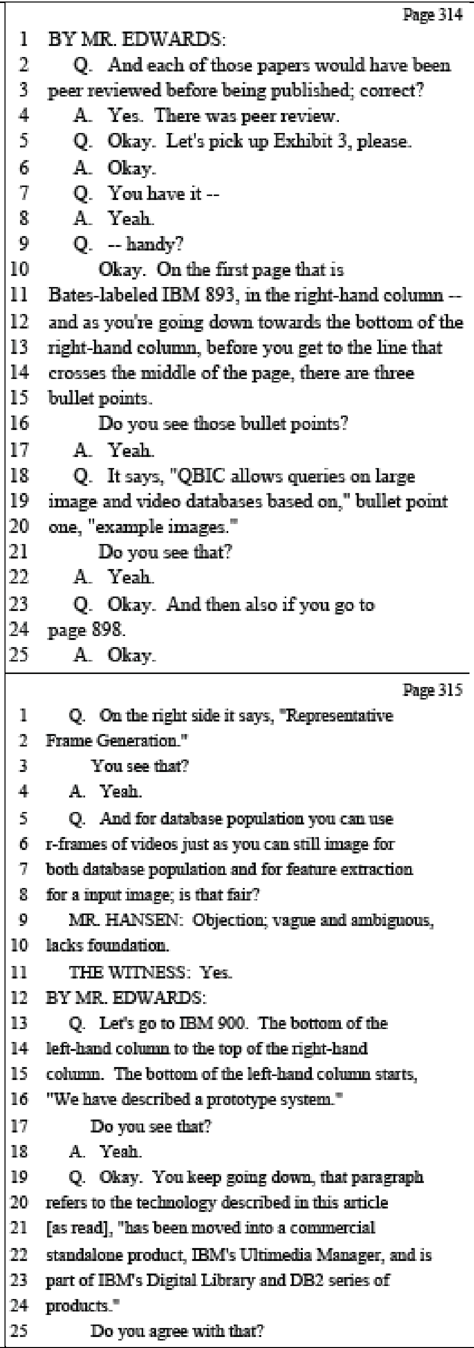

Page 314

BY MR. EDWARDS:
Q. And each of those papers would have been peer reviewed before being published; correct?
A. Yes. There was peer review.
Q. Okay. Let's pick up Exhibit 3, please.
A. Okay.
Q. You have it --
A. Yeah.
Q. -- handy?
Okay. On the first page that is Bates-labeled IBM 893, in the right-hand column -- and as you're going down towards the bottom of the right-hand column, before you get to the line that crosses the middle of the page, there are three bullet points.
Do you see those bullet points?
A. Yeah.
Q. It says, "QBIC allows queries on large image and video databases based on," bullet point one, "example images."
Do you see that?
A. Yeah.
Q. Okay. And then also if you go to page 898.
A. Okay.

Page 315

Q. On the right side it says, "Representative Frame Generation."
You see that?
A. Yeah.
Q. And for database population you can use r-frames of videos just as you can still image for both database population and for feature extraction for a input image; is that fair?
MR. HANSEN: Objection; vague and ambiguous, lacks foundation.
THE WITNESS: Yes.
BY MR. EDWARDS:
Q. Let's go to IBM 900. The bottom of the left-hand column to the top of the right-hand column. The bottom of the left-hand column starts, "We have described a prototype system."
Do you see that?
A. Yeah.
Q. Okay. You keep going down, that paragraph refers to the technology described in this article [as read], "has been moved into a commercial standalone product, IBM's Ultimedia Manager, and is part of IBM's Digital Library and DB2 series of products."
Do you agree with that?

Element	'036 Claim Recitation	Exemplary Citations in Reference
		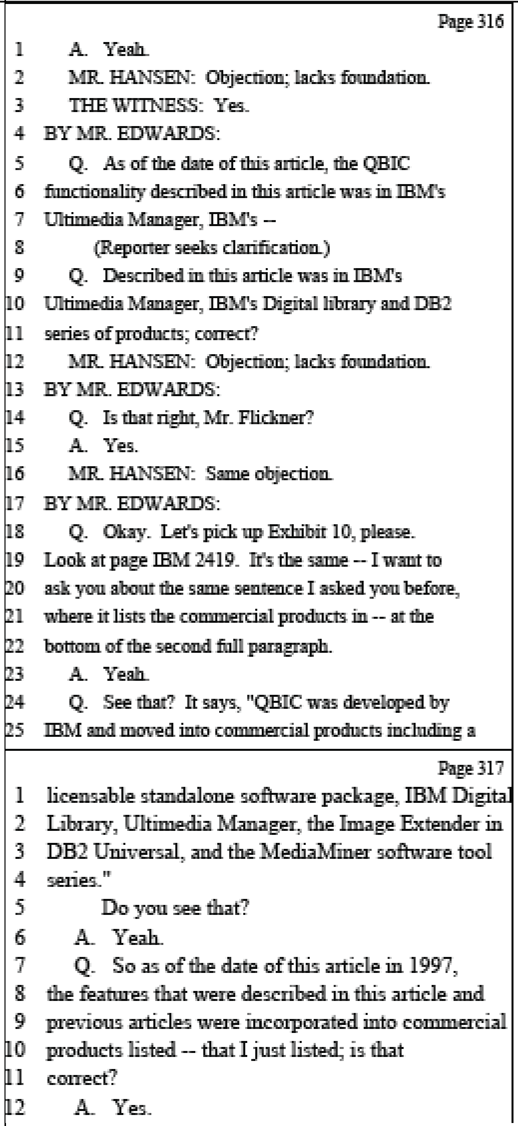 Page 316 A. Yeah. MR. HANSEN: Objection; lacks foundation. THE WITNESS: Yes. BY MR. EDWARDS: Q. As of the date of this article, the QBIC functionality described in this article was in IBM's Ultimedia Manager, IBM's -- (Reporter seeks clarification.) Q. Described in this article was in IBM's Ultimedia Manager, IBM's Digital library and DB2 series of products; correct? MR. HANSEN: Objection; lacks foundation. BY MR. EDWARDS: Q. Is that right, Mr. Flickner? A. Yes. MR. HANSEN: Same objection. BY MR. EDWARDS: Q. Okay. Let's pick up Exhibit 10, please. Look at page IBM 2419. It's the same -- I want to ask you about the same sentence I asked you before, where it lists the commercial products in -- at the bottom of the second full paragraph. A. Yeah. Q. See that? It says, "QBIC was developed by IBM and moved into commercial products including a Page 317 licensable standalone software package, IBM Digital Library, Ultimedia Manager, the Image Extender in DB2 Universal, and the MediaMiner software tool series." Do you see that? A. Yeah. Q. So as of the date of this article in 1997, the features that were described in this article and previous articles were incorporated into commercial products listed -- that I just listed; is that correct? A. Yes.

Element	'036 Claim Recitation	Exemplary Citations in Reference
		See also Flickner Depo at 22-25, 25-26, 54-55, 76, 78-79, 89-90, 159, 201-202, 209-210 To the extent it is found that QBIC does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that QBIC does not anticipate this claim, QBIC renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine QBIC with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1C.ii]	sends the content information to at least one of the identification platform and the mobile device.	QBIC discloses [a content service coupled with the identification platform] that sends the content information to at least one of the identification platform and the mobile device. The QBIC Project: Querying Images by Content Using Color, Texture, and Shape (IBM 0002390–404) at 8 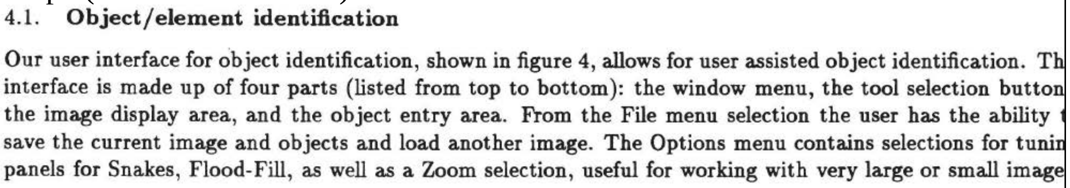 4.1. **Object/element identification** Our user interface for object identification, shown in figure 4, allows for user assisted object identification. Th interface is made up of four parts (listed from top to bottom): the window menu, the tool selection button the image display area, and the object entry area. From the File menu selection the user has the ability t save the current image and objects and load another image. The Options menu contains selections for tunin panels for Snakes, Flood-Fill, as well as a Zoom selection, useful for working with very large or small image The QBIC Project: Querying Images by Content Using Color, Texture, and Shape (IBM 0002390–404) at Fig. 4

Element	'036 Claim Recitation	Exemplary Citations in Reference
		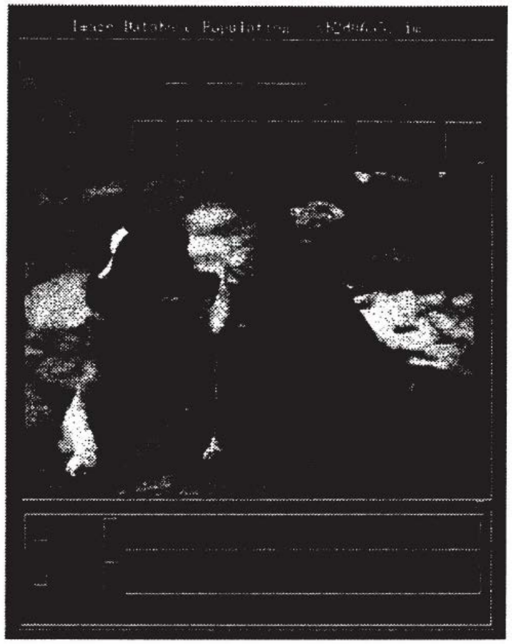 Figure 4: Object Identification Tool. Note for example, the right dog's ear, which has been outlined. The QBIC Project: Querying Images by Content Using Color, Texture, and Shape (IBM 0002390–404) at 8

Element	'036 Claim Recitation	Exemplary Citations in Reference
		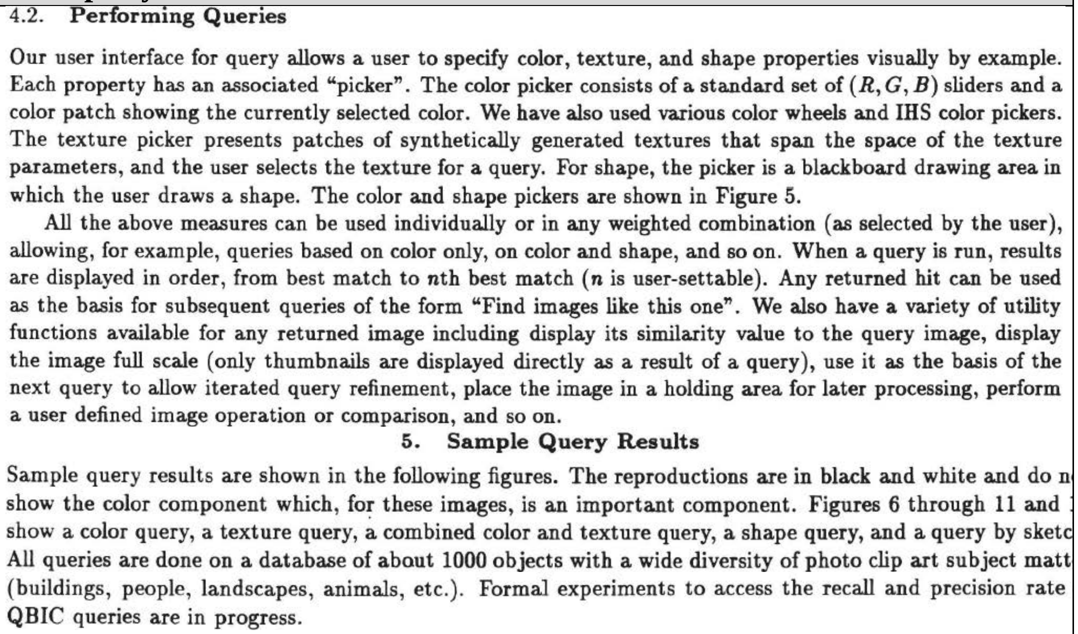 4.2. **Performing Queries** Our user interface for query allows a user to specify color, texture, and shape properties visually by example. Each property has an associated "picker". The color picker consists of a standard set of (R, G, B) sliders and a color patch showing the currently selected color. We have also used various color wheels and IHS color pickers. The texture picker presents patches of synthetically generated textures that span the space of the texture parameters, and the user selects the texture for a query. For shape, the picker is a blackboard drawing area in which the user draws a shape. The color and shape pickers are shown in Figure 5. All the above measures can be used individually or in any weighted combination (as selected by the user), allowing, for example, queries based on color only, on color and shape, and so on. When a query is run, results are displayed in order, from best match to nth best match (n is user-settable). Any returned hit can be used as the basis for subsequent queries of the form "Find images like this one". We also have a variety of utility functions available for any returned image including display its similarity value to the query image, display the image full scale (only thumbnails are displayed directly as a result of a query), use it as the basis of the next query to allow iterated query refinement, place the image in a holding area for later processing, perform a user defined image operation or comparison, and so on. 5. **Sample Query Results** Sample query results are shown in the following figures. The reproductions are in black and white and do n show the color component which, for these images, is an important component. Figures 6 through 11 and show a color query, a texture query, a combined color and texture query, a shape query, and a query by sketc All queries are done on a database of about 1000 objects with a wide diversity of photo clip art subject matt (buildings, people, landscapes, animals, etc.). Formal experiments to access the recall and precision rate QBIC queries are in progress. QBIC (Query by Image Content) (BOFA00001391–392) at 2:

Element	'036 Claim Recitation	Exemplary Citations in Reference
		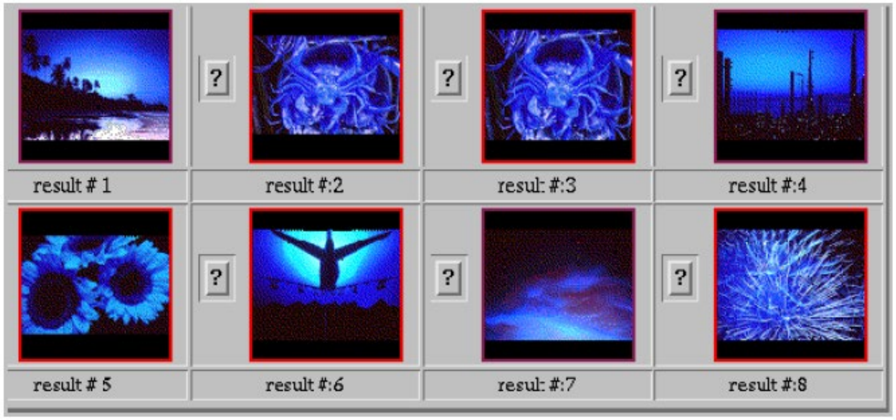 **Result of QBIC Search** Query by Image and Video Content: The QBIC System (IBM 0000893–902): "To achieve this functionality, QBIC has two main components: database population (the process of creating an image database) and database query. During the population, images and videos are processed to extract features describing their content-colors, textures, shapes, and camera and object motion-and the features are stored in a database. During the query, the user composes a query graphically. Features are generated from the graphical query and then input to a matching engine that finds images or videos from the database with similar features. Figure 2 shows the system architecture." *Id*. at p. 3.

Element	'036 Claim Recitation	Exemplary Citations in Reference
		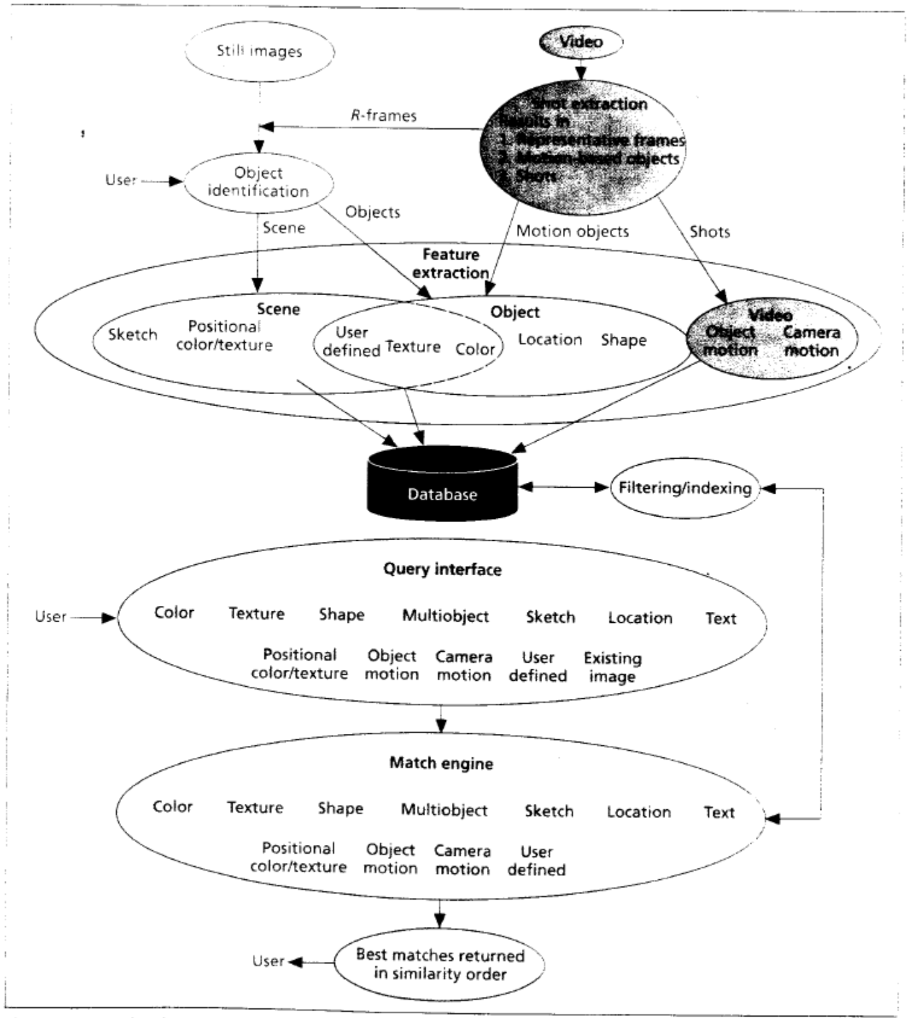 Figure 2. QBIC database population (top) and query (bottom) architecture. *Id.* at Fig. 2.

Element	'036 Claim Recitation	Exemplary Citations in Reference
		Image, Audio, and Video Extenders Administration and Programming (IBM 000001–617): "To make an image available for searching by content, you catalog the image. When you catalog an image, the Image Extender analyzes the image, by computing the feature values for the image, and stores the values in a QBIC catalog. When you search for an image by content, your query identifies one or more features for the search (such as average color), a source for each feature (such as an example image), and a target set of cataloged images. The Image Extender computes the feature value of the source and compares it to the cataloged feature values for the target images. It then computes a score that indicates how similar the feature values of the target images are to the source. You can have the Image Extender return the images whose features are most similar to the source. The Image Extender will return the handle of each image and the image score. You can also have the Image Extender return only the score of a single image." *Id.* at pp. 20–21. QBIC Demo (IBM000747): 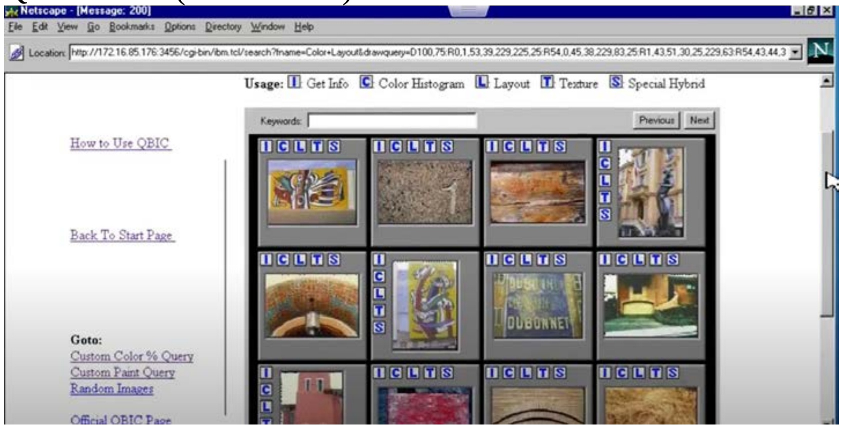

Element	’036 Claim Recitation	Exemplary Citations in Reference
		QBIC Demo (IBM000747):  Image france50.jpg Flickner Depo at 41-42: 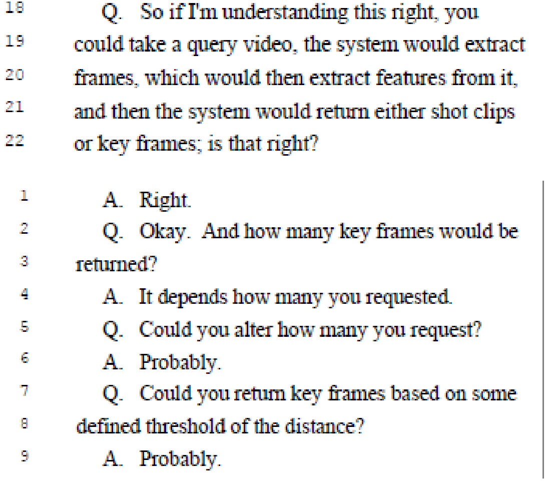 Q. So if I'm understanding this right, you could take a query video, the system would extract frames, which would then extract features from it, and then the system would return either shot clips or key frames; is that right? A. Right. Q. Okay. And how many key frames would be returned? A. It depends how many you requested. Q. Could you alter how many you request? A. Probably. Q. Could you return key frames based on some defined threshold of the distance? A. Probably.

Element	'036 Claim Recitation	Exemplary Citations in Reference
		Flickner Depo at 162-163 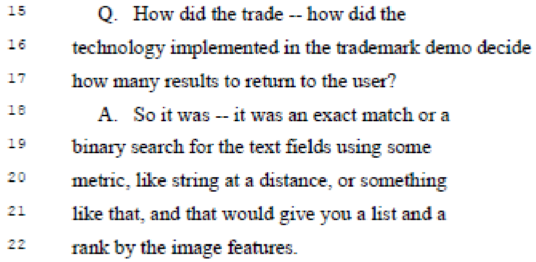 Q. How did the trade -- how did the technology implemented in the trademark demo decide how many results to return to the user? A. So it was -- it was an exact match or a binary search for the text fields using some metric, like string at a distance, or something like that, and that would give you a list and a rank by the image features. 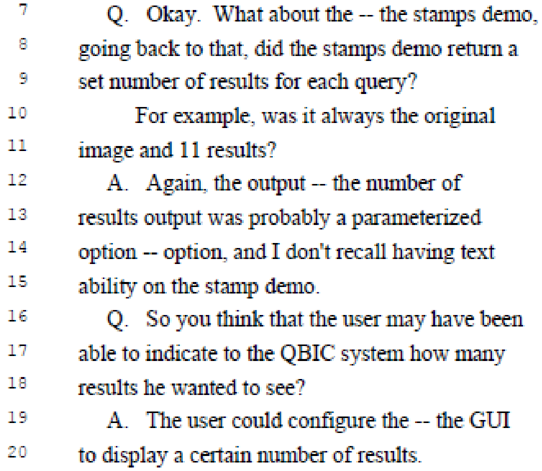 Q. Okay. What about the -- the stamps demo, going back to that, did the stamps demo return a set number of results for each query? For example, was it always the original image and 11 results? A. Again, the output -- the number of results output was probably a parameterized option -- option, and I don't recall having text ability on the stamp demo. Q. So you think that the user may have been able to indicate to the QBIC system how many results he wanted to see? A. The user could configure the -- the GUI to display a certain number of results.

Element	'036 Claim Recitation	Exemplary Citations in Reference
		Flickner Depo at 201-202: 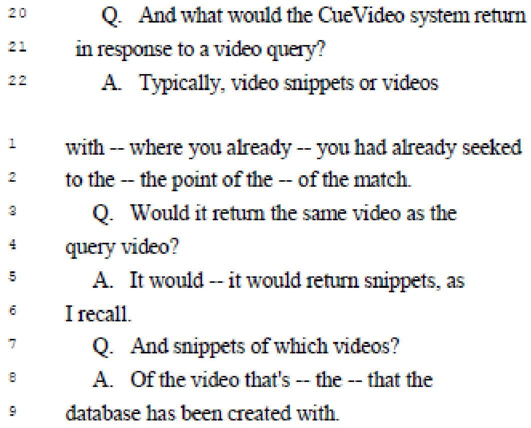 Q. And what would the CueVideo system return in response to a video query? A. Typically, video snippets or videos with -- where you already -- you had already seeked to the -- the point of the -- of the match. Q. Would it return the same video as the query video? A. It would -- it would return snippets, as I recall. Q. And snippets of which videos? A. Of the video that's -- the -- that the database has been created with. Flickner Depo at 131-132: 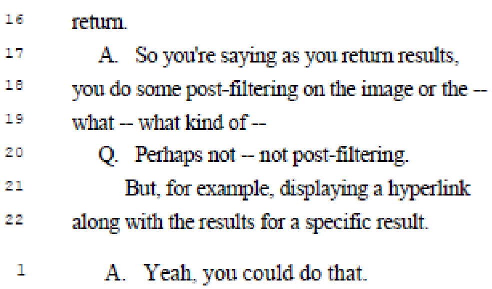 Q. An action that the QBIC system would automatically take based on the results it would return. A. So you're saying as you return results, you do some post-filtering on the image or the -- what -- what kind of -- Q. Perhaps not -- not post-filtering. But, for example, displaying a hyperlink along with the results for a specific result. A. Yeah, you could do that. Flickner 8/29/23 Depo at 68–72:

Element	'036 Claim Recitation	Exemplary Citations in Reference
		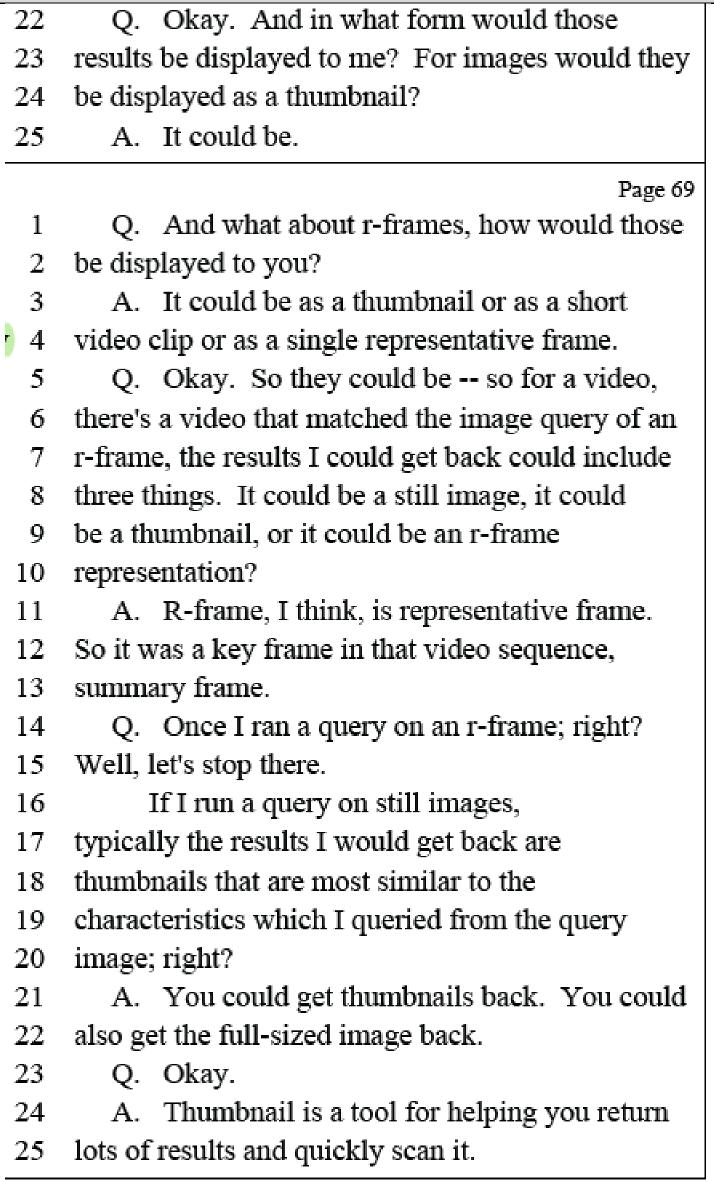 Q. Okay. And in what form would those results be displayed to me? For images would they be displayed as a thumbnail? A. It could be. Page 69 Q. And what about r-frames, how would those be displayed to you? A. It could be as a thumbnail or as a short video clip or as a single representative frame. Q. Okay. So they could be -- so for a video, there's a video that matched the image query of an r-frame, the results I could get back could include three things. It could be a still image, it could be a thumbnail, or it could be an r-frame representation? A. R-frame, I think, is representative frame. So it was a key frame in that video sequence, summary frame. Q. Once I ran a query on an r-frame; right? Well, let's stop there. If I run a query on still images, typically the results I would get back are thumbnails that are most similar to the characteristics which I queried from the query image; right? A. You could get thumbnails back. You could also get the full-sized image back. Q. Okay. A. Thumbnail is a tool for helping you return lots of results and quickly scan it.

Element	'036 Claim Recitation	Exemplary Citations in Reference
		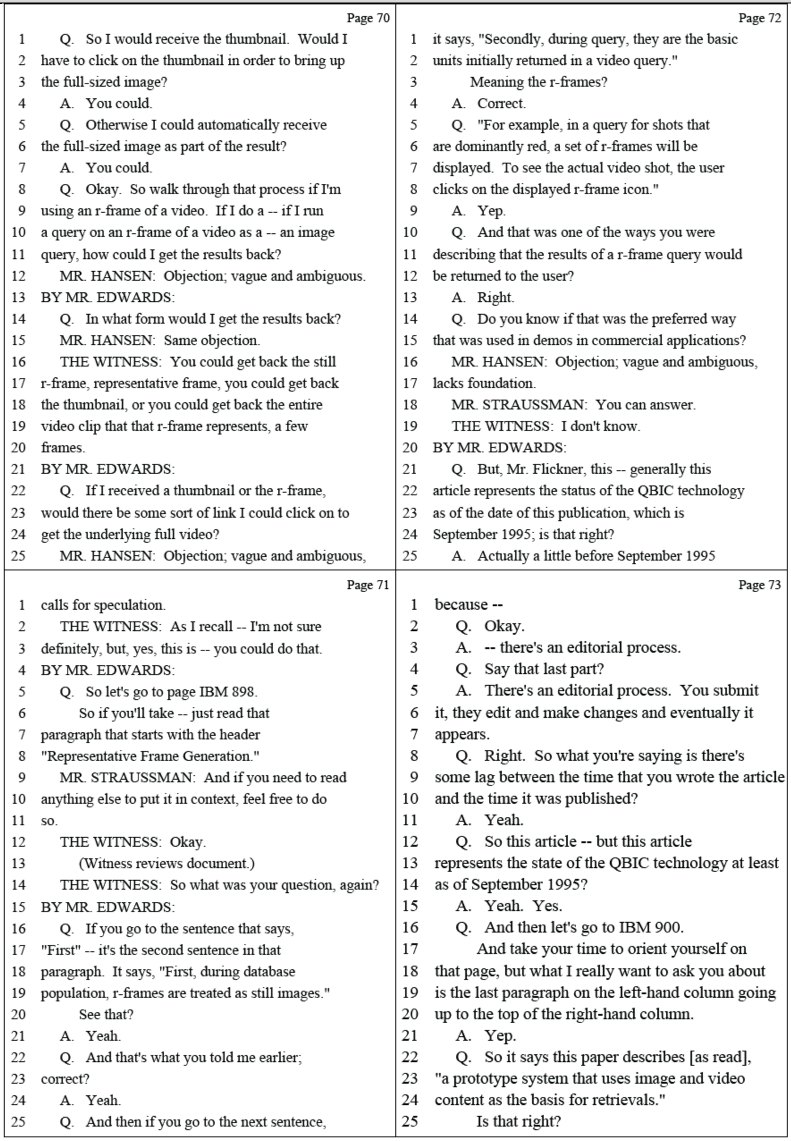 Page 70 Q. So I would receive the thumbnail. Would I have to click on the thumbnail in order to bring up the full-sized image? A. You could. Q. Otherwise I could automatically receive the full-sized image as part of the result? A. You could. Q. Okay. So walk through that process if I'm using an r-frame of a video. If I do a -- if I run a query on an r-frame of a video as a -- an image query, how could I get the results back? MR. HANSEN: Objection; vague and ambiguous. BY MR. EDWARDS: Q. In what form would I get the results back? MR. HANSEN: Same objection. THE WITNESS: You could get back the still r-frame, representative frame, you could get back the thumbnail, or you could get back the entire video clip that that r-frame represents, a few frames. BY MR. EDWARDS: Q. If I received a thumbnail or the r-frame, would there be some sort of link I could click on to get the underlying full video? MR. HANSEN: Objection; vague and ambiguous, Page 71 calls for speculation. THE WITNESS: As I recall -- I'm not sure definitely, but, yes, this is -- you could do that. BY MR. EDWARDS: Q. So let's go to page IBM 898. So if you'll take -- just read that paragraph that starts with the header "Representative Frame Generation." MR. STRAUSSMAN: And if you need to read anything else to put it in context, feel free to do so. THE WITNESS: Okay. (Witness reviews document.) THE WITNESS: So what was your question, again? BY MR. EDWARDS: Q. If you go to the sentence that says, "First" -- it's the second sentence in that paragraph. It says, "First, during database population, r-frames are treated as still images." See that? A. Yeah. Q. And that's what you told me earlier; correct? A. Yeah. Q. And then if you go to the next sentence, Page 72 it says, "Secondly, during query, they are the basic units initially returned in a video query." Meaning the r-frames? A. Correct. Q. "For example, in a query for shots that are dominantly red, a set of r-frames will be displayed. To see the actual video shot, the user clicks on the displayed r-frame icon." A. Yep. Q. And that was one of the ways you were describing that the results of a r-frame query would be returned to the user? A. Right. Q. Do you know if that was the preferred way that was used in demos in commercial applications? MR. HANSEN: Objection; vague and ambiguous, lacks foundation. MR. STRAUSSMAN: You can answer. THE WITNESS: I don't know. BY MR. EDWARDS: Q. But, Mr. Flickner, this -- generally this article represents the status of the QBIC technology as of the date of this publication, which is September 1995; is that right? A. Actually a little before September 1995 Page 73 because -- Q. Okay. A. -- there's an editorial process. Q. Say that last part? A. There's an editorial process. You submit it, they edit and make changes and eventually it appears. Q. Right. So what you're saying is there's some lag between the time that you wrote the article and the time it was published? A. Yeah. Q. So this article -- but this article represents the state of the QBIC technology at least as of September 1995? A. Yeah. Yes. Q. And then let's go to IBM 900. And take your time to orient yourself on that page, but what I really want to ask you about is the last paragraph on the left-hand column going up to the top of the right-hand column. A. Yep. Q. So it says this paper describes [as read], "a prototype system that uses image and video content as the basis for retrievals." Is that right?

Element	'036 Claim Recitation	Exemplary Citations in Reference
		Id. at 129–131: 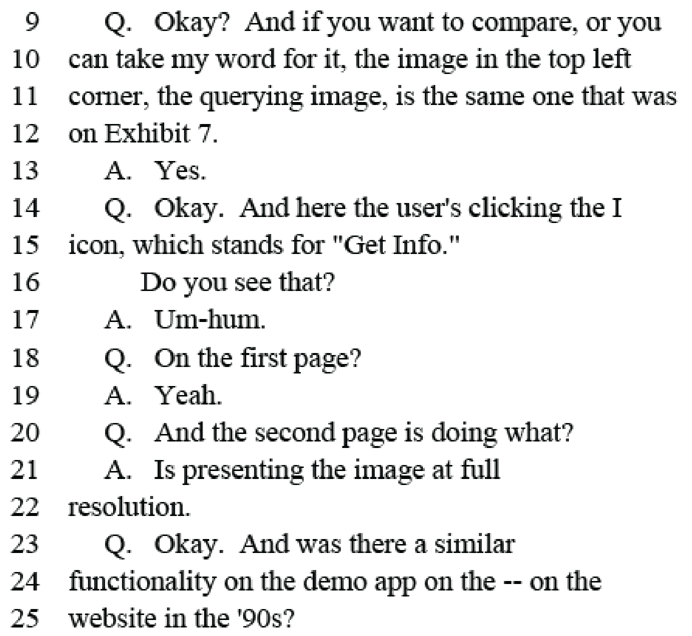 Q. Okay? And if you want to compare, or you can take my word for it, the image in the top left corner, the querying image, is the same one that was on Exhibit 7. A. Yes. Q. Okay. And here the user's clicking the I icon, which stands for "Get Info." Do you see that? A. Um-hum. Q. On the first page? A. Yeah. Q. And the second page is doing what? A. Is presenting the image at full resolution. Q. Okay. And was there a similar functionality on the demo app on the -- on the website in the '90s?

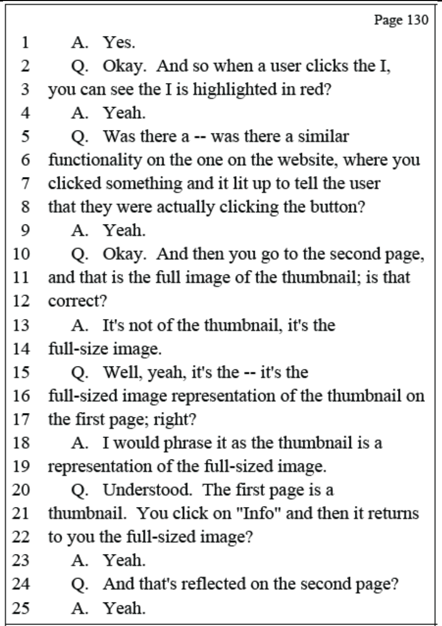

Page 130

A. Yes.
Q. Okay. And so when a user clicks the I, you can see the I is highlighted in red?
A. Yeah.
Q. Was there a -- was there a similar functionality on the one on the website, where you clicked something and it lit up to tell the user that they were actually clicking the button?
A. Yeah.
Q. Okay. And then you go to the second page, and that is the full image of the thumbnail; is that correct?
A. It's not of the thumbnail, it's the full-size image.
Q. Well, yeah, it's the -- it's the full-sized image representation of the thumbnail on the first page; right?
A. I would phrase it as the thumbnail is a representation of the full-sized image.
Q. Understood. The first page is a thumbnail. You click on "Info" and then it returns to you the full-sized image?
A. Yeah.
Q. And that's reflected on the second page?
A. Yeah.

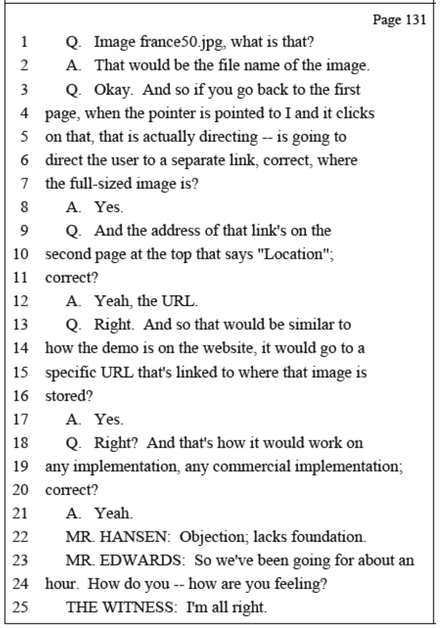

Page 131

Q. Image france50.jpg, what is that?
A. That would be the file name of the image.
Q. Okay. And so if you go back to the first page, when the pointer is pointed to I and it clicks on that, that is actually directing -- is going to direct the user to a separate link, correct, where the full-sized image is?
A. Yes.
Q. And the address of that link's on the second page at the top that says "Location"; correct?
A. Yeah, the URL.
Q. Right. And so that would be similar to how the demo is on the website, it would go to a specific URL that's linked to where that image is stored?
A. Yes.
Q. Right? And that's how it would work on any implementation, any commercial implementation; correct?
A. Yeah.
MR. HANSEN: Objection; lacks foundation.
MR. EDWARDS: So we've been going for about an hour. How do you -- how are you feeling?
THE WITNESS: I'm all right.

Element	'036 Claim Recitation	Exemplary Citations in Reference
		Id. at 161: 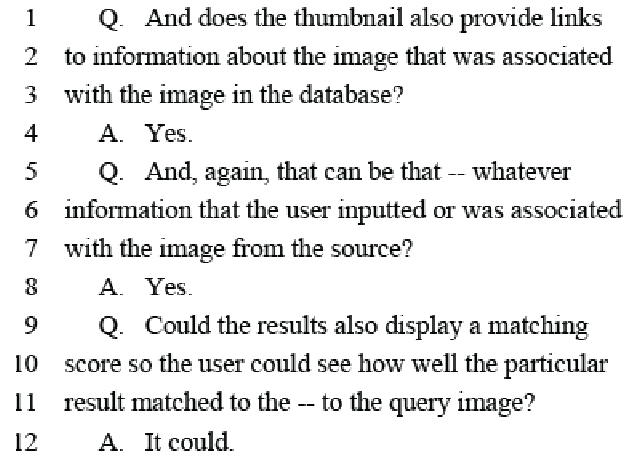 Q. And does the thumbnail also provide links to information about the image that was associated with the image in the database? A. Yes. Q. And, again, that can be that -- whatever information that the user inputted or was associated with the image from the source? A. Yes. Q. Could the results also display a matching score so the user could see how well the particular result matched to the -- to the query image? A. It could. *Id.* at 321: 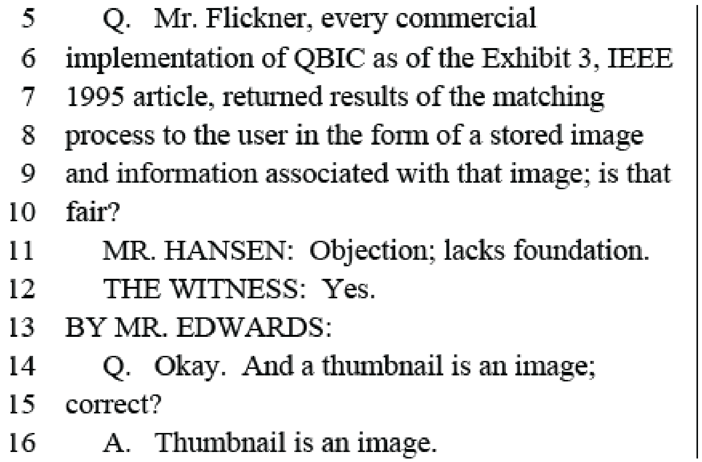 Q. Mr. Flickner, every commercial implementation of QBIC as of the Exhibit 3, IEEE 1995 article, returned results of the matching process to the user in the form of a stored image and information associated with that image; is that fair? MR. HANSEN: Objection; lacks foundation. THE WITNESS: Yes. BY MR. EDWARDS: Q. Okay. And a thumbnail is an image; correct? A. Thumbnail is an image. *Id.* at 305:

Element	’036 Claim Recitation	Exemplary Citations in Reference
		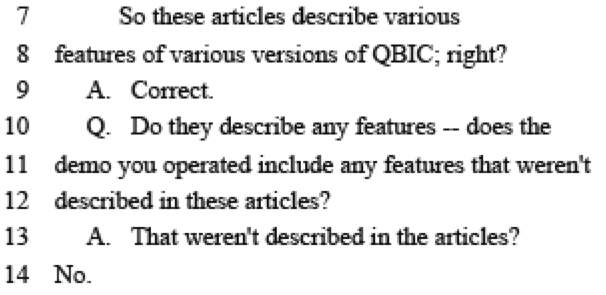 So these articles describe various features of various versions of QBIC; right? A. Correct. Q. Do they describe any features -- does the demo you operated include any features that weren't described in these articles? A. That weren't described in the articles? No. *Id.* at 35: 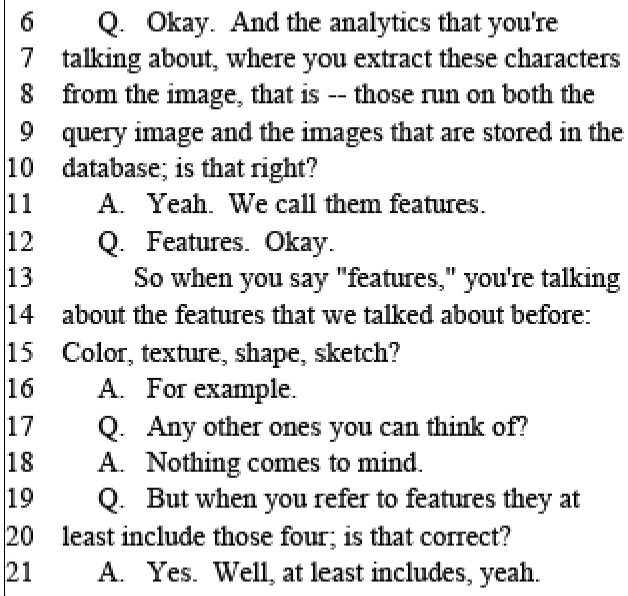 Q. Okay. And the analytics that you're talking about, where you extract these characters from the image, that is -- those run on both the query image and the images that are stored in the database; is that right? A. Yeah. We call them features. Q. Features. Okay. So when you say "features," you're talking about the features that we talked about before: Color, texture, shape, sketch? A. For example. Q. Any other ones you can think of? A. Nothing comes to mind. Q. But when you refer to features they at least include those four; is that correct? A. Yes. Well, at least includes, yeah. *Id.* at 109-110:

Element	'036 Claim Recitation	Exemplary Citations in Reference
		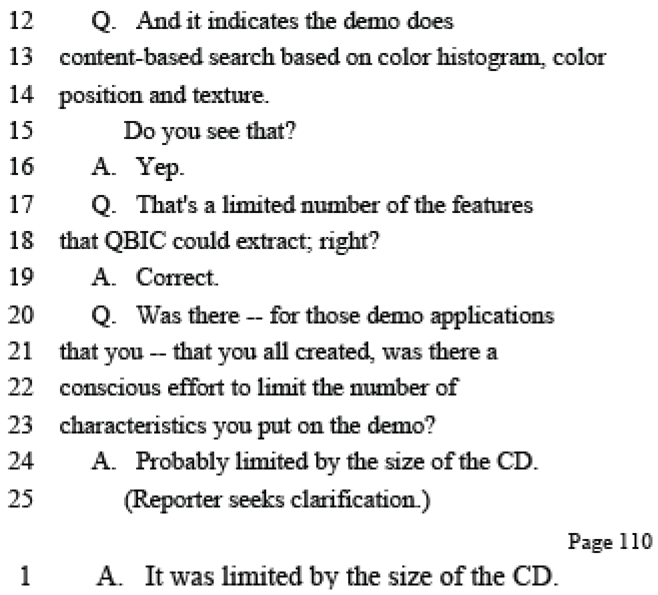 Q. And it indicates the demo does content-based search based on color histogram, color position and texture. Do you see that? A. Yep. Q. That's a limited number of the features that QBIC could extract; right? A. Correct. Q. Was there -- for those demo applications that you -- that you all created, was there a conscious effort to limit the number of characteristics you put on the demo? A. Probably limited by the size of the CD. (Reporter seeks clarification.) Page 110 A. It was limited by the size of the CD. *Id.* at 311: 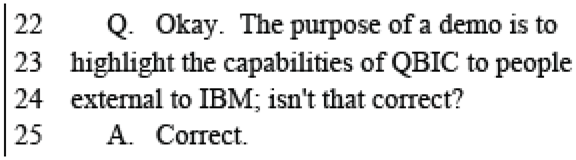 Q. Okay. The purpose of a demo is to highlight the capabilities of QBIC to people external to IBM; isn't that correct? A. Correct. *Id.* at 313–317:

Element	'036 Claim Recitation	Exemplary Citations in Those Reference
		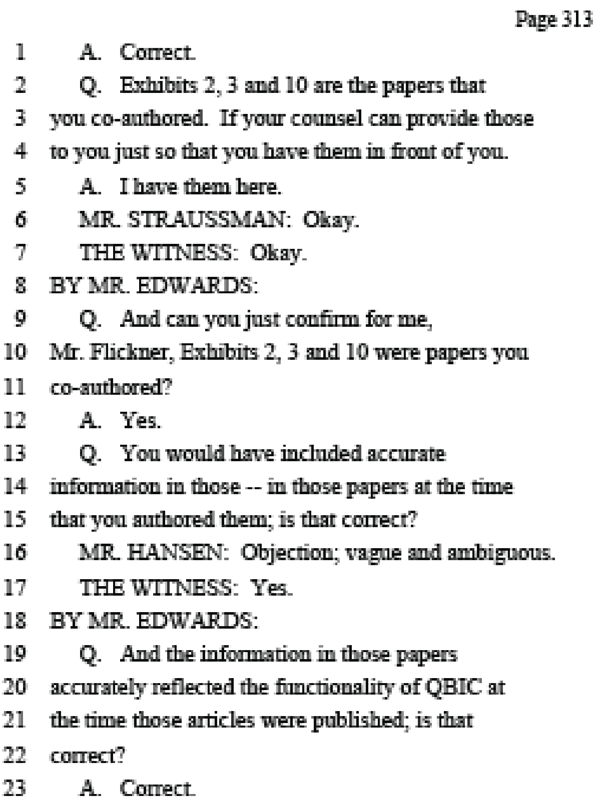 Page 313 A. Correct. Q. Exhibits 2, 3 and 10 are the papers that you co-authored. If your counsel can provide those to you just so that you have them in front of you. A. I have them here. MR. STRAUSSMAN: Okay. THE WITNESS: Okay. BY MR. EDWARDS: Q. And can you just confirm for me, Mr. Flickner, Exhibits 2, 3 and 10 were papers you co-authored? A. Yes. Q. You would have included accurate information in those -- in those papers at the time that you authored them; is that correct? MR. HANSEN: Objection; vague and ambiguous. THE WITNESS: Yes. BY MR. EDWARDS: Q. And the information in those papers accurately reflected the functionality of QBIC at the time those articles were published; is that correct? A. Correct.

		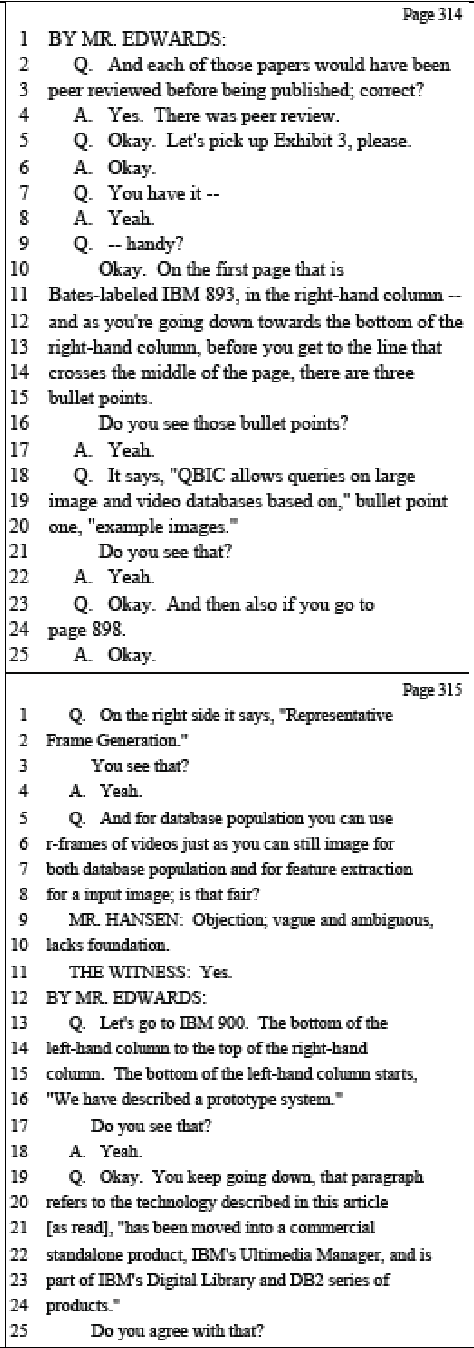 Page 314 BY MR. EDWARDS: Q. And each of those papers would have been peer reviewed before being published; correct? A. Yes. There was peer review. Q. Okay. Let's pick up Exhibit 3, please. A. Okay. Q. You have it -- A. Yeah. Q. -- handy? Okay. On the first page that is Bates-labeled IBM 893, in the right-hand column -- and as you're going down towards the bottom of the right-hand column, before you get to the line that crosses the middle of the page, there are three bullet points. Do you see those bullet points? A. Yeah. Q. It says, "QBIC allows queries on large image and video databases based on," bullet point one, "example images." Do you see that? A. Yeah. Q. Okay. And then also if you go to page 898. A. Okay. Page 315 Q. On the right side it says, "Representative Frame Generation." You see that? A. Yeah. Q. And for database population you can use r-frames of videos just as you can still image for both database population and for feature extraction for a input image; is that fair? MR. HANSEN: Objection; vague and ambiguous, lacks foundation. THE WITNESS: Yes. BY MR. EDWARDS: Q. Let's go to IBM 900. The bottom of the left-hand column to the top of the right-hand column. The bottom of the left-hand column starts, "We have described a prototype system." Do you see that? A. Yeah. Q. Okay. You keep going down, that paragraph refers to the technology described in this article [as read], "has been moved into a commercial standalone product, IBM's Ultimedia Manager, and is part of IBM's Digital Library and DB2 series of products." Do you agree with that?

Element	'036 Claim Recitation	Exemplary Citations in Reference
		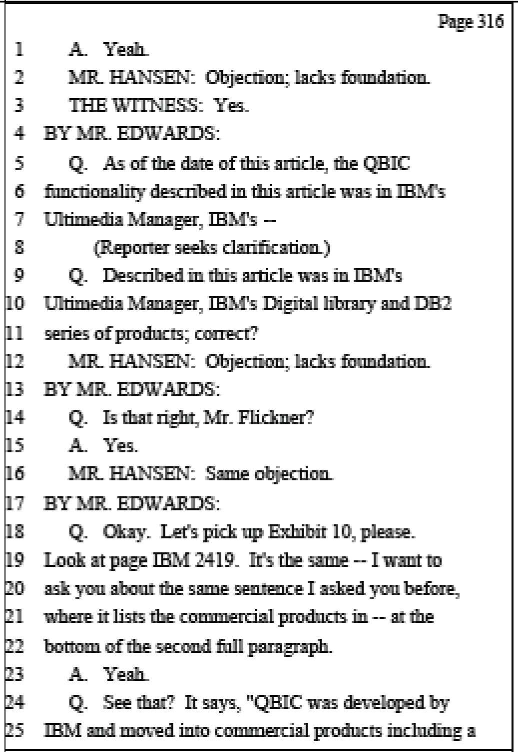 Page 316 A. Yeah. MR. HANSEN: Objection; lacks foundation. THE WITNESS: Yes. BY MR. EDWARDS: Q. As of the date of this article, the QBIC functionality described in this article was in IBM's Ultimedia Manager, IBM's -- (Reporter seeks clarification.) Q. Described in this article was in IBM's Ultimedia Manager, IBM's Digital library and DB2 series of products; correct? MR. HANSEN: Objection; lacks foundation. BY MR. EDWARDS: Q. Is that right, Mr. Flickner? A. Yes. MR. HANSEN: Same objection. BY MR. EDWARDS: Q. Okay. Let's pick up Exhibit 10, please. Look at page IBM 2419. It's the same -- I want to ask you about the same sentence I asked you before, where it lists the commercial products in -- at the bottom of the second full paragraph. A. Yeah. Q. See that? It says, "QBIC was developed by IBM and moved into commercial products including a 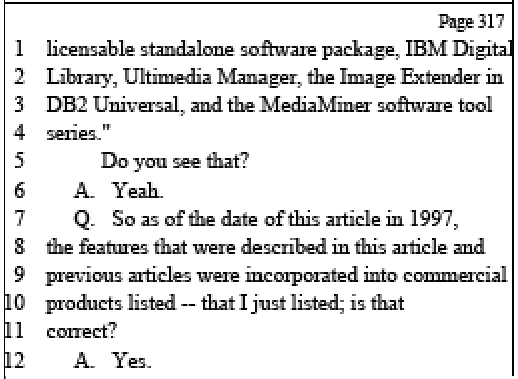 Page 317 licensable standalone software package, IBM Digital Library, Ultimedia Manager, the Image Extender in DB2 Universal, and the MediaMiner software tool series." Do you see that? A. Yeah. Q. So as of the date of this article in 1997, the features that were described in this article and previous articles were incorporated into commercial products listed -- that I just listed; is that correct? A. Yes.

Element	'036 Claim Recitation	Exemplary Citations in Reference
		See also Flickner Depo at 25-26, 50, 54-55, 209-210 To the extent it is found that QBIC does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that QBIC does not anticipate this claim, QBIC renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine QBIC with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

Nantworks, LLC v. Bank of America Corporation
Central District of California, No. 2:20-cv-7872
Invalidity Chart for U.S. Patent 8,478,036 based on Hewlett-Packard CoolTown System ("CoolTown")

Exhibit C-31

Bank of America Corporation ("BofA" or "Defendant") provides this chart subject to all reservations, objections, statements, and disclaimers set forth herein and in BofA's Final Invalidity Contentions Cover Pleading, as well as any amendment, supplement, or modification thereof, which are incorporated herein by reference in their entirety.

On information and belief, CoolTown was made, known or used by others in the United States no later than July 10, 2000. *See generally* T. Kindberg et al., A Web-Based Nomadic Computing System; J. Barton et al., The Cooltown User Experience; cooltown developer's network, coolbase overview; cooltown developer's network, people, places, things: web presence for the real world. CoolTown is available as prior art at least under 35 U.S.C. § 102(a). On information and belief, HP was using CoolTown and made it available to the public in the United States by releasing the website on the Internet (and accessible to U.S. users) by at least March 2000, making the software available on the HP website by March 2000, and making available on the Internet (and accessible to U.S. users) the open source code, which reflects the functionality of CoolTown by at least March 2000. *See e.g.*, HP000043–044; HP000045–046; HP000060–061; *see also* https://www.zdnet.com/article/hp-brings-open-source-to-cooltown/. Therefore, HP Cooltown was at least known, used, and/or sold to the public by at least 2000. On information and belief, CoolTown is also available as prior art under 102(g). BofA is diligently investigating CoolTown and will supplement with additional information, if necessary.

As shown in the chart below, CoolTown anticipates asserted claim 1 of the '036 Patent. To the extent CoolTown is found not to expressly disclose certain limitations in the asserted claims, such limitations are inherent. To the extent CoolTown is found not to anticipate any asserted claims or claim elements of the '036 Patent, the reference nevertheless renders those claims or claim elements obvious under 35 U.S.C. § 103, either alone or in combination with other art identified herein or in the cover pleading of BofA's Invalidity Contentions, for the reasons identified in the cover pleading. These Final Invalidity Contentions are not an admission by BofA that the accused products, including any current or past versions of the accused products, are covered by, or infringe the asserted claims, but are based instead on the recognition that if the claims are interpreted to be broad enough to encompass the accused products, the claims must also be construed to have that same scope when considering whether they are invalid.

The following chart is based on the Court's claim construction rulings (Dkt. Nos. 145, 153, 236), and to the extent not contradictory, Plaintiffs' apparent interpretations of the claims in Plaintiffs' "Preliminary" Final Infringement Contentions. BofA notes that in many instances, Plaintiffs' "Preliminary" Final Infringement Contentions ignore and/or contradict the Court's claim construction rulings. BofA does not accept the assumptions concerning the scope and meaning implicit in Plaintiffs' Preliminary Final Infringement Contentions, to the extent those assumptions are discernible, and reserves the right to challenge Plaintiffs' proposed (or implied) constructions. Furthermore, Plaintiffs' "Preliminary" Final Contentions fail to comply with the level of specificity required by SPR 2.1.2 and 2.1.3, and include improper reservation of rights to supplement with "additional ways in which BoA Accused Products infringe." Any alleged deficient discovery is of Plaintiffs own making. Most if not all of the deficiencies in Plaintiffs' Preliminary Final Contentions have been documented by Defendants through multiple letters since at least May 18, 2021 and up to Bank of America's most recent letter of August 19, 2023. Plaintiffs' shifting infringement contentions, deficiencies in its infringement contentions, and changing claim scope have prejudiced BofA from being able to reasonably prepare its defenses, including its invalidity contentions. BofA also reserves the right to revise and supplement these charts if and when Plaintiffs are permitted to provide revised or supplemental Infringement Contentions. Citations given in the chart below are merely representative of the respective elements to provide notice to Plaintiffs of BofA's invalidity theories and are not meant to be exhaustive.

1. Claim 1

Element	'036 Claim Recitation	Exemplary Citations in Reference
[1Pre]	A content provisioning system comprising:	CoolTown discloses a content provisioning system. The Cooltown User Experience (HP000088 – 093) at 1: 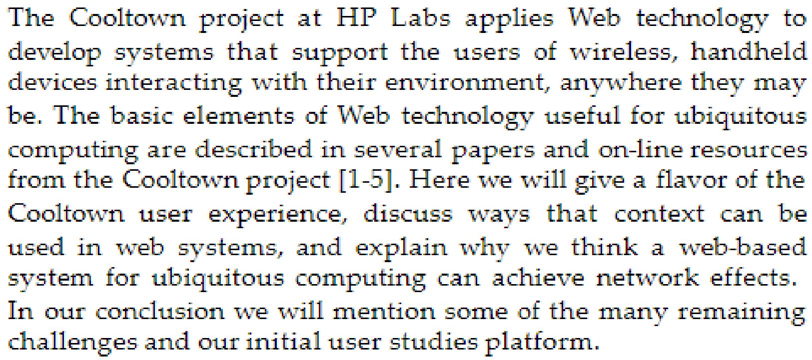 The Cooltown project at HP Labs applies Web technology to develop systems that support the users of wireless, handheld devices interacting with their environment, anywhere they may be. The basic elements of Web technology useful for ubiquitous computing are described in several papers and on-line resources from the Cooltown project [1-5]. Here we will give a flavor of the Cooltown user experience, discuss ways that context can be used in web systems, and explain why we think a web-based system for ubiquitous computing can achieve network effects. In our conclusion we will mention some of the many remaining challenges and our initial user studies platform. A Web-Based Nomadic Computing System (HP000074 – 87) at 1 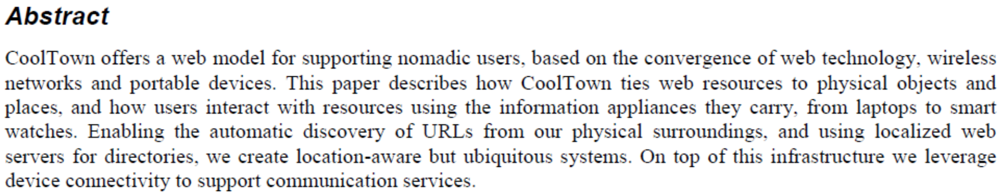 ***Abstract*** CoolTown offers a web model for supporting nomadic users, based on the convergence of web technology, wireless networks and portable devices. This paper describes how CoolTown ties web resources to physical objects and places, and how users interact with resources using the information appliances they carry, from laptops to smart watches. Enabling the automatic discovery of URLs from our physical surroundings, and using localized web servers for directories, we create location-aware but ubiquitous systems. On top of this infrastructure we leverage device connectivity to support communication services.

Element	'036 Claim Recitation	Exemplary Citations in Reference
		Coolbase[1] Overview at 1 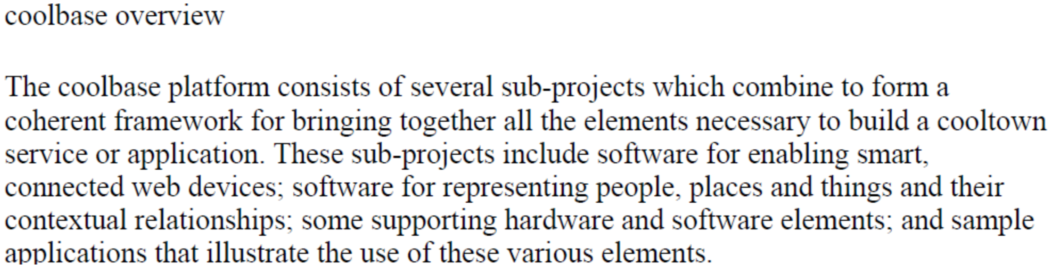 coolbase overview The coolbase platform consists of several sub-projects which combine to form a coherent framework for bringing together all the elements necessary to build a cooltown service or application. These sub-projects include software for enabling smart, connected web devices; software for representing people, places and things and their contextual relationships; some supporting hardware and software elements; and sample applications that illustrate the use of these various elements.
[1A]	a target database storing known targets of different types and recognition parameters associated with the known targets;	CoolTown discloses a target database storing known targets of barcodes. A Web-Based Nomadic Computing System (HP000074 – 87) at 4 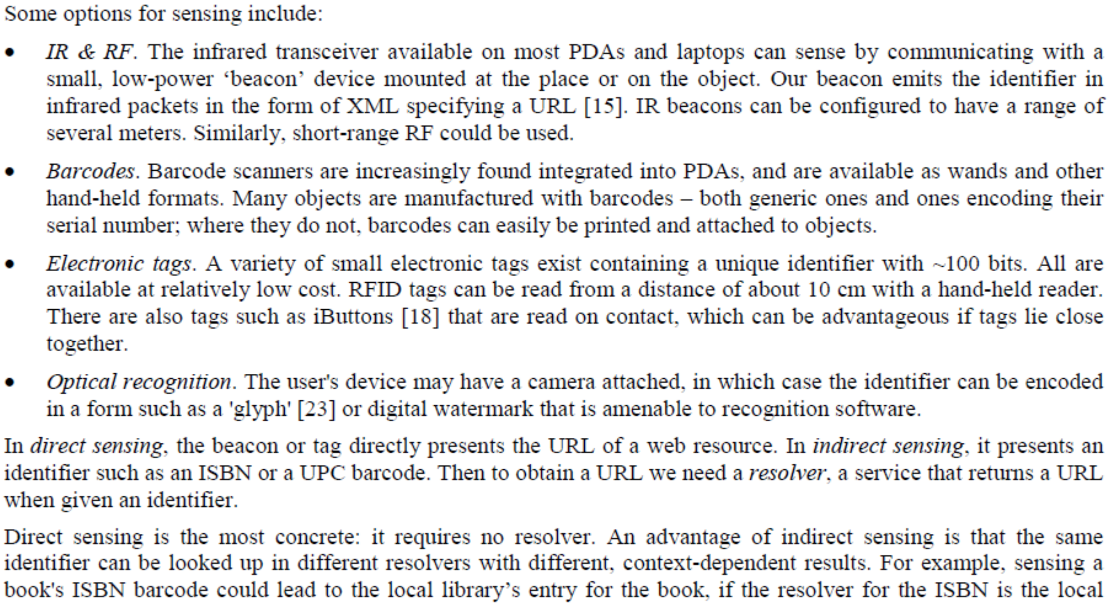 Some options for sensing include: • *IR & RF*. The infrared transceiver available on most PDAs and laptops can sense by communicating with a small, low-power 'beacon' device mounted at the place or on the object. Our beacon emits the identifier in infrared packets in the form of XML specifying a URL [15]. IR beacons can be configured to have a range of several meters. Similarly, short-range RF could be used. • *Barcodes*. Barcode scanners are increasingly found integrated into PDAs, and are available as wands and other hand-held formats. Many objects are manufactured with barcodes – both generic ones and ones encoding their serial number; where they do not, barcodes can easily be printed and attached to objects. • *Electronic tags*. A variety of small electronic tags exist containing a unique identifier with ~100 bits. All are available at relatively low cost. RFID tags can be read from a distance of about 10 cm with a hand-held reader. There are also tags such as iButtons [18] that are read on contact, which can be advantageous if tags lie close together. • *Optical recognition*. The user's device may have a camera attached, in which case the identifier can be encoded in a form such as a 'glyph' [23] or digital watermark that is amenable to recognition software. In *direct sensing*, the beacon or tag directly presents the URL of a web resource. In *indirect sensing*, it presents an identifier such as an ISBN or a UPC barcode. Then to obtain a URL we need a *resolver*, a service that returns a URL when given an identifier. Direct sensing is the most concrete: it requires no resolver. An advantage of indirect sensing is that the same identifier can be looked up in different resolvers with different, context-dependent results. For example, sensing a book's ISBN barcode could lead to the local library's entry for the book, if the resolver for the ISBN is the local

---

[1] *See* https://web.archive.org/web/20011030090650/http://cooltown.hp.com/dev/coolbase-overview.asp

Element	'036 Claim Recitation	Exemplary Citations in Reference
		A Web-Based Nomadic Computing System (HP000074 – 87) at 5 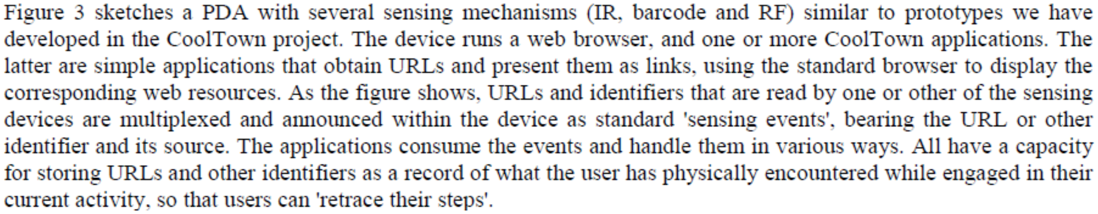 Figure 3 sketches a PDA with several sensing mechanisms (IR, barcode and RF) similar to prototypes we have developed in the CoolTown project. The device runs a web browser, and one or more CoolTown applications. The latter are simple applications that obtain URLs and present them as links, using the standard browser to display the corresponding web resources. As the figure shows, URLs and identifiers that are read by one or other of the sensing devices are multiplexed and announced within the device as standard 'sensing events', bearing the URL or other identifier and its source. The applications consume the events and handle them in various ways. All have a capacity for storing URLs and other identifiers as a record of what the user has physically encountered while engaged in their current activity, so that users can 'retrace their steps'. A Web-Based Nomadic Computing System (HP000074 – 87) at 6 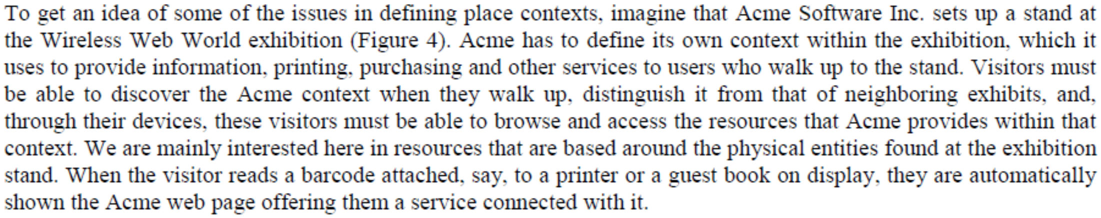 To get an idea of some of the issues in defining place contexts, imagine that Acme Software Inc. sets up a stand at the Wireless Web World exhibition (Figure 4). Acme has to define its own context within the exhibition, which it uses to provide information, printing, purchasing and other services to users who walk up to the stand. Visitors must be able to discover the Acme context when they walk up, distinguish it from that of neighboring exhibits, and, through their devices, these visitors must be able to browse and access the resources that Acme provides within that context. We are mainly interested here in resources that are based around the physical entities found at the exhibition stand. When the visitor reads a barcode attached, say, to a printer or a guest book on display, they are automatically shown the Acme web page offering them a service connected with it. A Web-Based Nomadic Computing System (HP000074 – 87) at 6 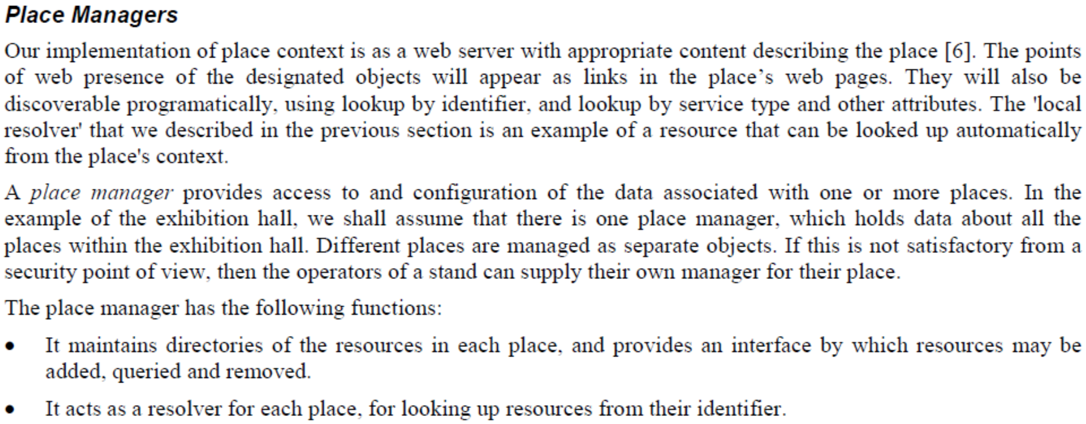 ***Place Managers*** Our implementation of place context is as a web server with appropriate content describing the place [6]. The points of web presence of the designated objects will appear as links in the place's web pages. They will also be discoverable programatically, using lookup by identifier, and lookup by service type and other attributes. The 'local resolver' that we described in the previous section is an example of a resource that can be looked up automatically from the place's context. A *place manager* provides access to and configuration of the data associated with one or more places. In the example of the exhibition hall, we shall assume that there is one place manager, which holds data about all the places within the exhibition hall. Different places are managed as separate objects. If this is not satisfactory from a security point of view, then the operators of a stand can supply their own manager for their place. The place manager has the following functions: • It maintains directories of the resources in each place, and provides an interface by which resources may be added, queried and removed. • It acts as a resolver for each place, for looking up resources from their identifier. Cooltown Developer's Network, People, Place, and Things: web presence for the real world (HP000025–34) (HP000025–34) at 8

Element	'036 Claim Recitation	Exemplary Citations in Reference
		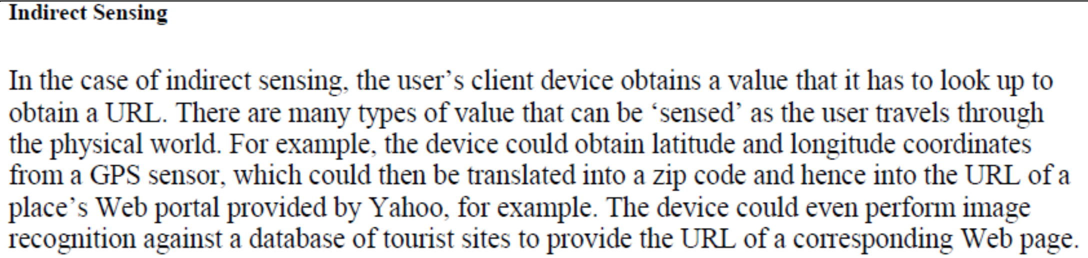 **Indirect Sensing** In the case of indirect sensing, the user's client device obtains a value that it has to look up to obtain a URL. There are many types of value that can be 'sensed' as the user travels through the physical world. For example, the device could obtain latitude and longitude coordinates from a GPS sensor, which could then be translated into a zip code and hence into the URL of a place's Web portal provided by Yahoo, for example. The device could even perform image recognition against a database of tourist sites to provide the URL of a corresponding Web page. Cooltown Developer's Network, People, Place, and Things: web presence for the real world (HP000025–34) (HP000025–34) at 23 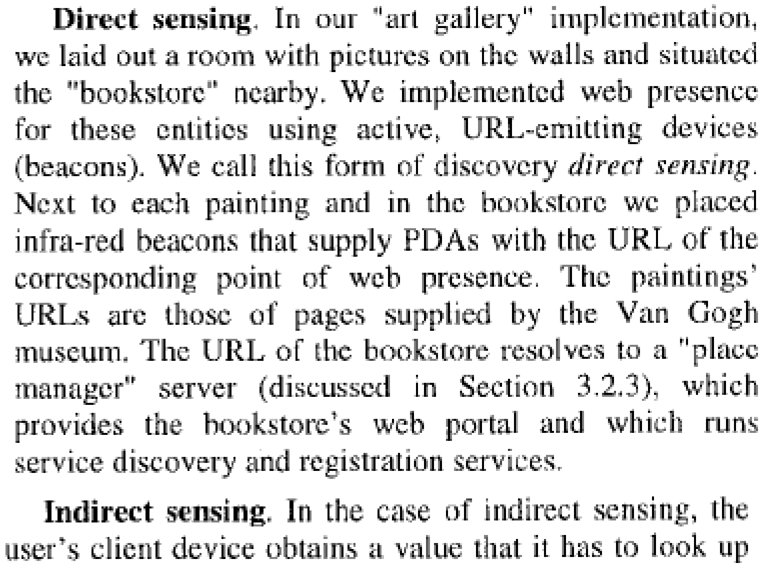 **Direct sensing**. In our "art gallery" implementation, we laid out a room with pictures on the walls and situated the "bookstore" nearby. We implemented web presence for these entities using active, URL-emitting devices (beacons). We call this form of discovery *direct sensing*. Next to each painting and in the bookstore we placed infra-red beacons that supply PDAs with the URL of the corresponding point of web presence. The paintings' URLs are those of pages supplied by the Van Gogh museum. The URL of the bookstore resolves to a "place manager" server (discussed in Section 3.2.3), which provides the bookstore's web portal and which runs service discovery and registration services. **Indirect sensing**. In the case of indirect sensing, the user's client device obtains a value that it has to look up

Element	'036 Claim Recitation	Exemplary Citations in Reference
		to obtain a URL. There are many types of value that can be 'sensed' as the user travels through the physical world. For example, the device could obtain latitude and longitude coordinates from a GPS sensor, which could then be translated into a zip code and hence into the URL of a place's web portal provided by Yahoo, for example. The device could even perform image recognition against a database of tourist sites to provide the URL of a corresponding web page. We have been experimenting with indirect sensing where the sensed value is obtained from a *tag*. There are electronic tags like RFID tags [33] and iButtons [22] – small devices that supply a unique identification string to a sensor placed near them or in contact with them. 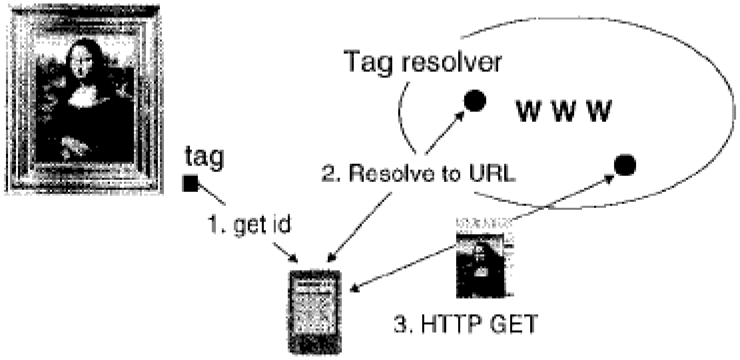 **Figure 4. Indirect sensing: the web presence for a painting.**

Element	'036 Claim Recitation	Exemplary Citations in Reference
		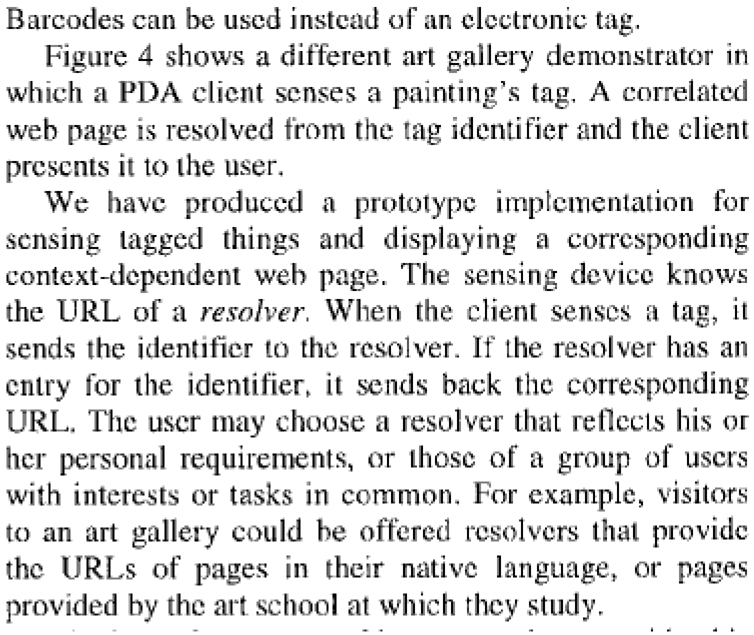 Barcodes can be used instead of an electronic tag. Figure 4 shows a different art gallery demonstrator in which a PDA client senses a painting's tag. A correlated web page is resolved from the tag identifier and the client presents it to the user. We have produced a prototype implementation for sensing tagged things and displaying a corresponding context-dependent web page. The sensing device knows the URL of a *resolver*. When the client senses a tag, it sends the identifier to the resolver. If the resolver has an entry for the identifier, it sends back the corresponding URL. The user may choose a resolver that reflects his or her personal requirements, or those of a group of users with interests or tasks in common. For example, visitors to an art gallery could be offered resolvers that provide the URLs of pages in their native language, or pages provided by the art school at which they study. To the extent it is found that CoolTown does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that CoolTown does not anticipate this claim, CoolTown renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine CoolTown with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1B]	an identification platform coupled with the target	CoolTown discloses an identification platform coupled with the target database.

Element	’036 Claim Recitation	Exemplary Citations in Reference
	database, and that	The Cooltown: People, Places, Things Web Presence for the Real World (HP000025 – 34) at 20: 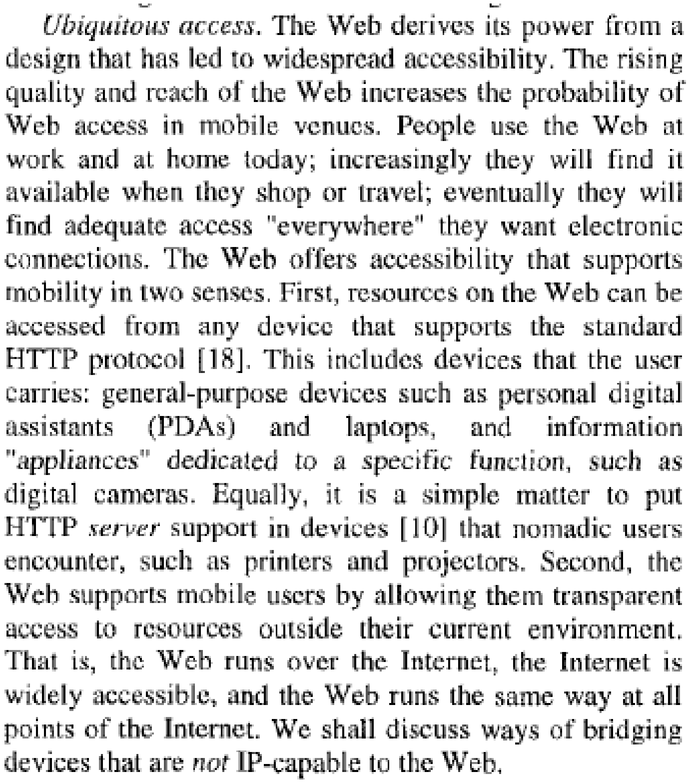 *Ubiquitous access*. The Web derives its power from a design that has led to widespread accessibility. The rising quality and reach of the Web increases the probability of Web access in mobile venues. People use the Web at work and at home today; increasingly they will find it available when they shop or travel; eventually they will find adequate access "everywhere" they want electronic connections. The Web offers accessibility that supports mobility in two senses. First, resources on the Web can be accessed from any device that supports the standard HTTP protocol [18]. This includes devices that the user carries: general-purpose devices such as personal digital assistants (PDAs) and laptops, and information "appliances" dedicated to a specific function, such as digital cameras. Equally, it is a simple matter to put HTTP *server* support in devices [10] that nomadic users encounter, such as printers and projectors. Second, the Web supports mobile users by allowing them transparent access to resources outside their current environment. That is, the Web runs over the Internet, the Internet is widely accessible, and the Web runs the same way at all points of the Internet. We shall discuss ways of bridging devices that are *not* IP-capable to the Web. The Cooltown: People, Places, Things Web Presence for the Real World (HP000025 – 34) at 24:

Element	'036 Claim Recitation	Exemplary Citations in Reference
		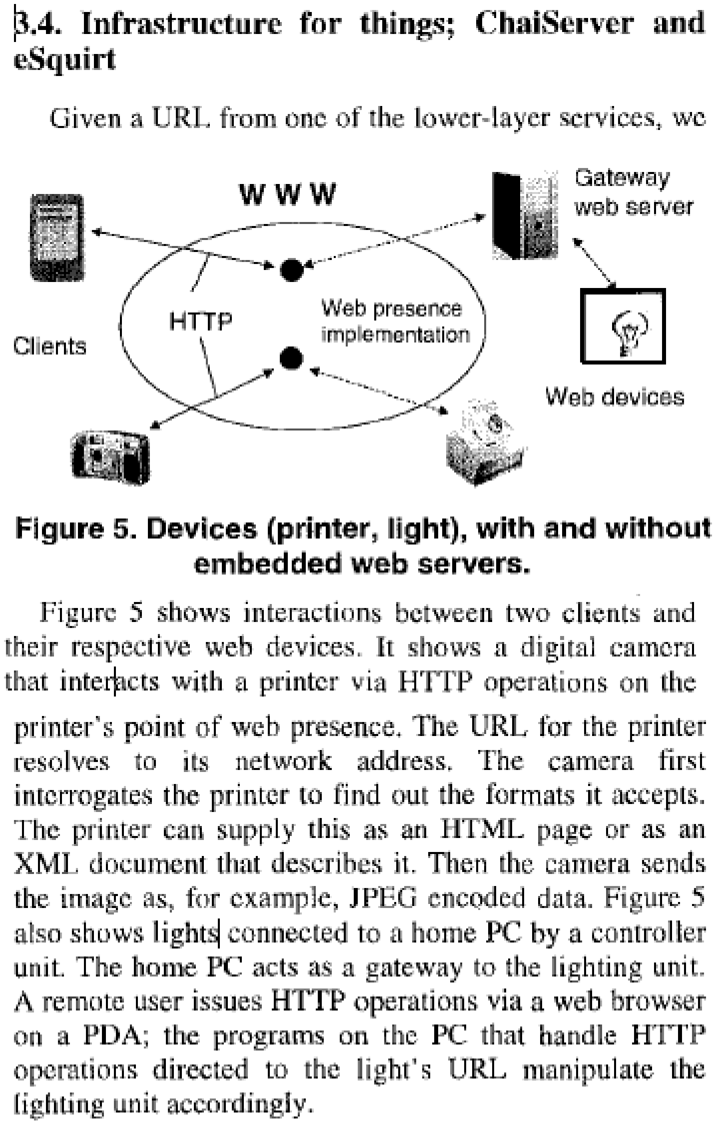 **3.4. Infrastructure for things; ChaiServer and eSquirt** Given a URL from one of the lower-layer services, we **Figure 5. Devices (printer, light), with and without embedded web servers.** Figure 5 shows interactions between two clients and their respective web devices. It shows a digital camera that interacts with a printer via HTTP operations on the printer's point of web presence. The URL for the printer resolves to its network address. The camera first interrogates the printer to find out the formats it accepts. The printer can supply this as an HTML page or as an XML document that describes it. Then the camera sends the image as, for example, JPEG encoded data. Figure 5 also shows lights connected to a home PC by a controller unit. The home PC acts as a gateway to the lighting unit. A remote user issues HTTP operations via a web browser on a PDA; the programs on the PC that handle HTTP operations directed to the light's URL manipulate the lighting unit accordingly. The Cooltown User Experience (HP000088 – 093) at 1:

Element	’036 Claim Recitation	Exemplary Citations in Reference
		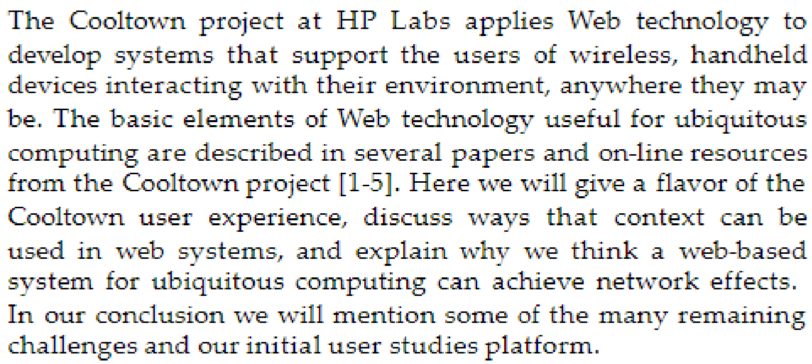 The Cooltown project at HP Labs applies Web technology to develop systems that support the users of wireless, handheld devices interacting with their environment, anywhere they may be. The basic elements of Web technology useful for ubiquitous computing are described in several papers and on-line resources from the Cooltown project [1-5]. Here we will give a flavor of the Cooltown user experience, discuss ways that context can be used in web systems, and explain why we think a web-based system for ubiquitous computing can achieve network effects. In our conclusion we will mention some of the many remaining challenges and our initial user studies platform. A Web-Based Nomadic Computing System (HP000074 – 87) at 1 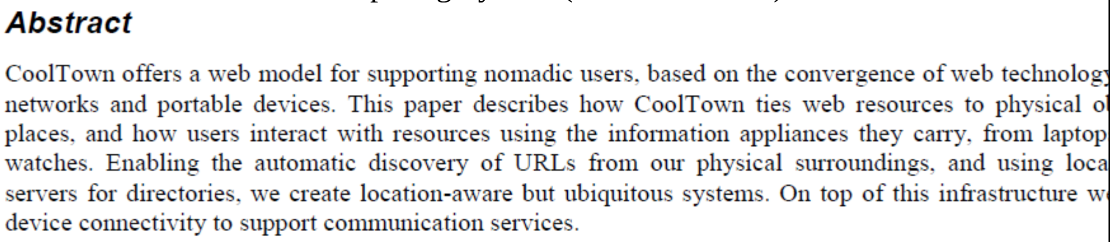 ***Abstract*** CoolTown offers a web model for supporting nomadic users, based on the convergence of web technolog networks and portable devices. This paper describes how CoolTown ties web resources to physical o places, and how users interact with resources using the information appliances they carry, from laptop watches. Enabling the automatic discovery of URLs from our physical surroundings, and using loca servers for directories, we create location-aware but ubiquitous systems. On top of this infrastructure w device connectivity to support communication services. A Web-Based Nomadic Computing System (HP000074 – 87) at 2 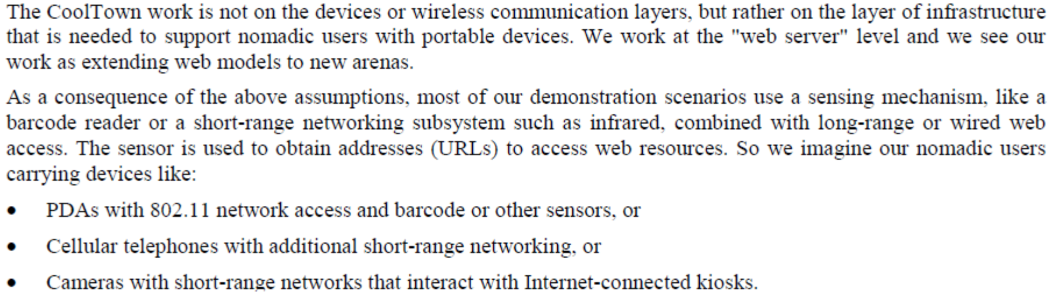 The CoolTown work is not on the devices or wireless communication layers, but rather on the layer of infrastructure that is needed to support nomadic users with portable devices. We work at the "web server" level and we see our work as extending web models to new arenas. As a consequence of the above assumptions, most of our demonstration scenarios use a sensing mechanism, like a barcode reader or a short-range networking subsystem such as infrared, combined with long-range or wired web access. The sensor is used to obtain addresses (URLs) to access web resources. So we imagine our nomadic users carrying devices like: • PDAs with 802.11 network access and barcode or other sensors, or • Cellular telephones with additional short-range networking, or • Cameras with short-range networks that interact with Internet-connected kiosks. A Web-Based Nomadic Computing System (HP000074 – 87) at 3

Element	'036 Claim Recitation	Exemplary Citations in Reference
		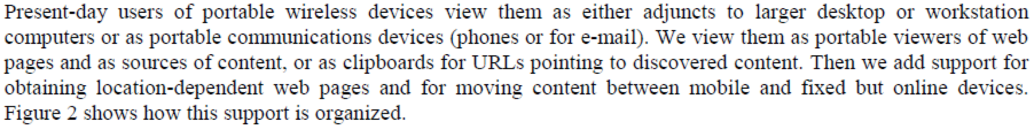 Present-day users of portable wireless devices view them as either adjuncts to larger desktop or workstation computers or as portable communications devices (phones or for e-mail). We view them as portable viewers of web pages and as sources of content, or as clipboards for URLs pointing to discovered content. Then we add support for obtaining location-dependent web pages and for moving content between mobile and fixed but online devices. Figure 2 shows how this support is organized. A Web-Based Nomadic Computing System (HP000074 – 87) at 3 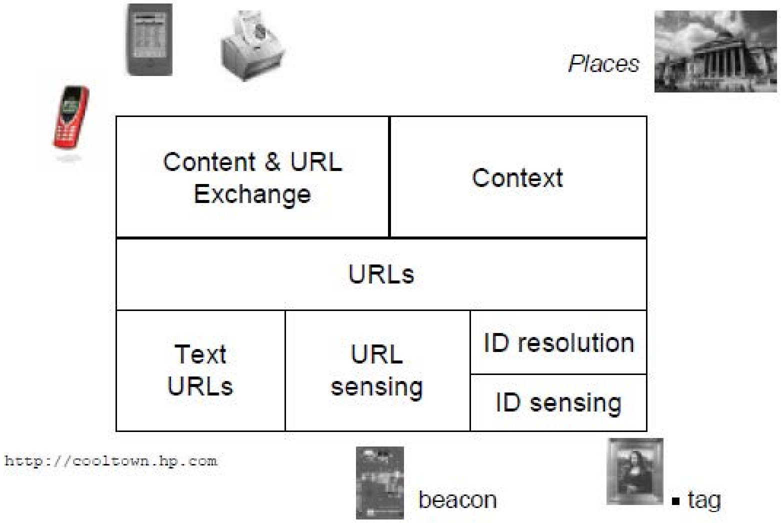 Figure 2. Component relationships. Sources of URLs, including keyboards, beacons, and resolved tags point to context or content. A Web-Based Nomadic Computing System (HP000074 – 87) at 4

Element	'036 Claim Recitation	Exemplary Citations in Reference
		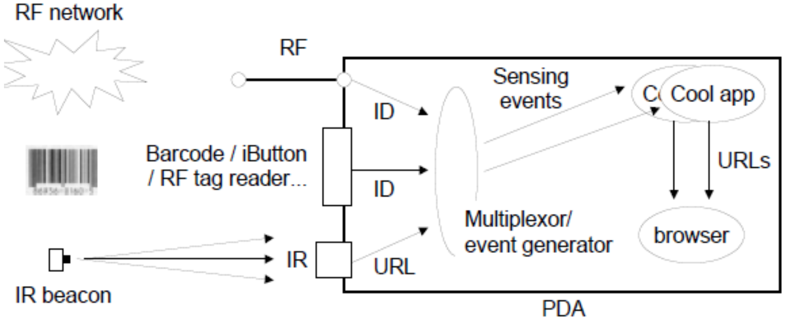 Figure 3. A sketch of a PDA equipped for sensing URLs and identifiers. 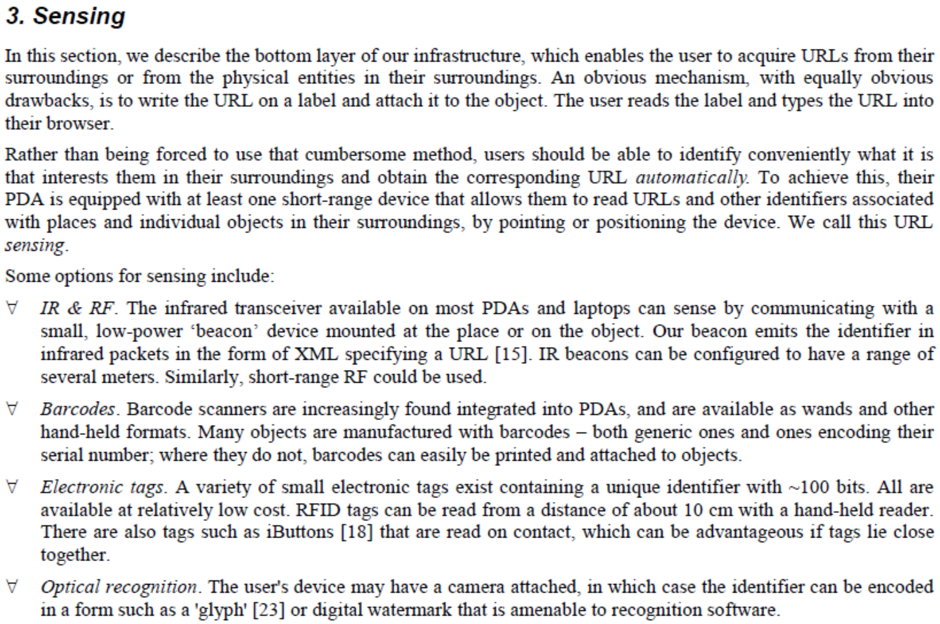 ***3. Sensing*** In this section, we describe the bottom layer of our infrastructure, which enables the user to acquire URLs from their surroundings or from the physical entities in their surroundings. An obvious mechanism, with equally obvious drawbacks, is to write the URL on a label and attach it to the object. The user reads the label and types the URL into their browser. Rather than being forced to use that cumbersome method, users should be able to identify conveniently what it is that interests them in their surroundings and obtain the corresponding URL *automatically*. To achieve this, their PDA is equipped with at least one short-range device that allows them to read URLs and other identifiers associated with places and individual objects in their surroundings, by pointing or positioning the device. We call this URL *sensing*. Some options for sensing include: ∀ *IR & RF*. The infrared transceiver available on most PDAs and laptops can sense by communicating with a small, low-power 'beacon' device mounted at the place or on the object. Our beacon emits the identifier in infrared packets in the form of XML specifying a URL [15]. IR beacons can be configured to have a range of several meters. Similarly, short-range RF could be used. ∀ *Barcodes*. Barcode scanners are increasingly found integrated into PDAs, and are available as wands and other hand-held formats. Many objects are manufactured with barcodes – both generic ones and ones encoding their serial number; where they do not, barcodes can easily be printed and attached to objects. ∀ *Electronic tags*. A variety of small electronic tags exist containing a unique identifier with ~100 bits. All are available at relatively low cost. RFID tags can be read from a distance of about 10 cm with a hand-held reader. There are also tags such as iButtons [18] that are read on contact, which can be advantageous if tags lie close together. ∀ *Optical recognition*. The user's device may have a camera attached, in which case the identifier can be encoded in a form such as a 'glyph' [23] or digital watermark that is amenable to recognition software. A Web-Based Nomadic Computing System (HP000074 – 87) at 4 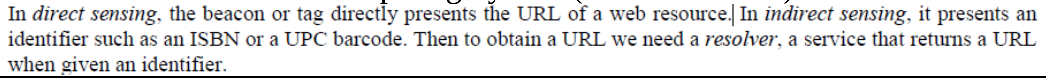 In *direct sensing*, the beacon or tag directly presents the URL of a web resource. In *indirect sensing*, it presents an identifier such as an ISBN or a UPC barcode. Then to obtain a URL we need a *resolver*, a service that returns a URL when given an identifier.

Element	'036 Claim Recitation	Exemplary Citations in Reference
		A Web-Based Nomadic Computing System (HP000074 – 87) at 8 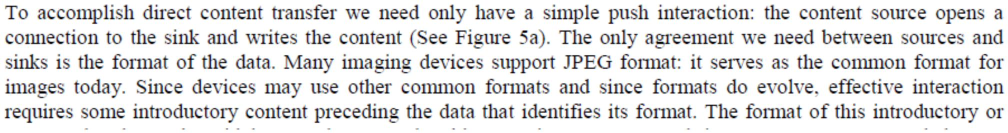 To accomplish direct content transfer we need only have a simple push interaction: the content source opens a connection to the sink and writes the content (See Figure 5a). The only agreement we need between sources and sinks is the format of the data. Many imaging devices support JPEG format: it serves as the common format for images today. Since devices may use other common formats and since formats do evolve, effective interaction requires some introductory content preceding the data that identifies its format. The format of this introductory or The Cooltown: People, Places, Things Web Presence for the Real World (HP000025 – 34) at 20: 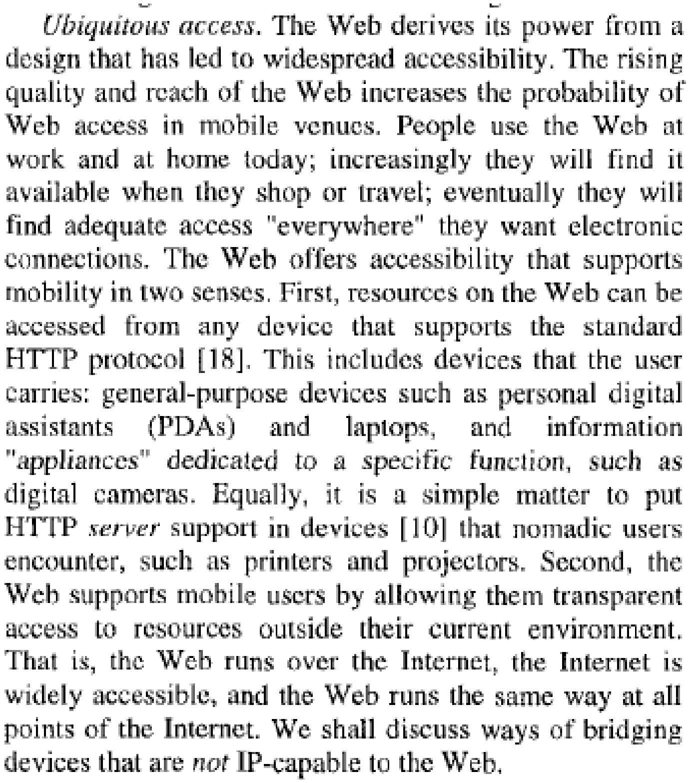 *Ubiquitous access*. The Web derives its power from a design that has led to widespread accessibility. The rising quality and reach of the Web increases the probability of Web access in mobile venues. People use the Web at work and at home today; increasingly they will find it available when they shop or travel; eventually they will find adequate access "everywhere" they want electronic connections. The Web offers accessibility that supports mobility in two senses. First, resources on the Web can be accessed from any device that supports the standard HTTP protocol [18]. This includes devices that the user carries: general-purpose devices such as personal digital assistants (PDAs) and laptops, and information "appliances" dedicated to a specific function, such as digital cameras. Equally, it is a simple matter to put HTTP *server* support in devices [10] that nomadic users encounter, such as printers and projectors. Second, the Web supports mobile users by allowing them transparent access to resources outside their current environment. That is, the Web runs over the Internet, the Internet is widely accessible, and the Web runs the same way at all points of the Internet. We shall discuss ways of bridging devices that are *not* IP-capable to the Web.

Element	'036 Claim Recitation	Exemplary Citations in Reference
		To the extent it is found that CoolTown does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that CoolTown does not anticipate this claim, CoolTown renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine CoolTown with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1B.i]	communicates with a mobile device capable of acquiring a digital representation of a scene containing at least a portion of a target;	CoolTown discloses an identification platform coupled with the target database that communicates with a mobile device capable of acquiring a digital representation of a scene containing at least a portion of a target. A Web-Based Nomadic Computing System (HP000074 – 87) at 1 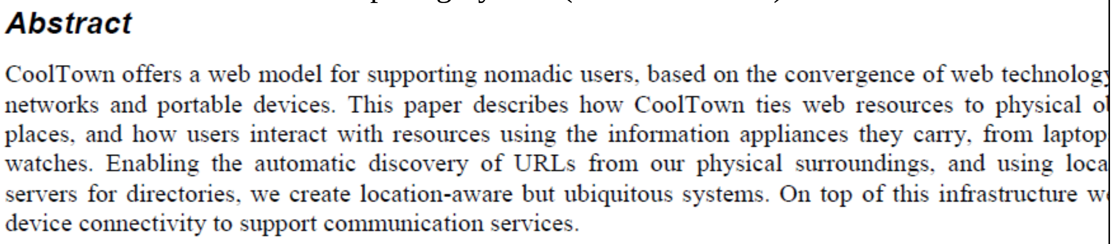 **Abstract** CoolTown offers a web model for supporting nomadic users, based on the convergence of web technolog networks and portable devices. This paper describes how CoolTown ties web resources to physical o places, and how users interact with resources using the information appliances they carry, from laptop watches. Enabling the automatic discovery of URLs from our physical surroundings, and using loca servers for directories, we create location-aware but ubiquitous systems. On top of this infrastructure w device connectivity to support communication services. The Cooltown: People, Places, Things Web Presence for the Real World (HP000025 – 34) at 20:

Element	'036 Claim Recitation	Exemplary Citations in Reference
		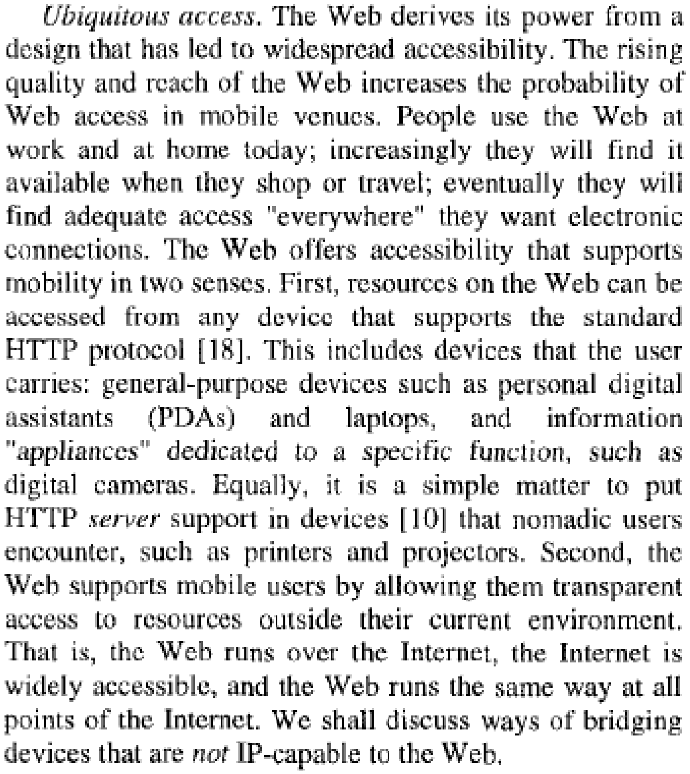 *Ubiquitous access.* The Web derives its power from a design that has led to widespread accessibility. The rising quality and reach of the Web increases the probability of Web access in mobile venues. People use the Web at work and at home today; increasingly they will find it available when they shop or travel; eventually they will find adequate access "everywhere" they want electronic connections. The Web offers accessibility that supports mobility in two senses. First, resources on the Web can be accessed from any device that supports the standard HTTP protocol [18]. This includes devices that the user carries: general-purpose devices such as personal digital assistants (PDAs) and laptops, and information "appliances" dedicated to a specific function, such as digital cameras. Equally, it is a simple matter to put HTTP *server* support in devices [10] that nomadic users encounter, such as printers and projectors. Second, the Web supports mobile users by allowing them transparent access to resources outside their current environment. That is, the Web runs over the Internet, the Internet is widely accessible, and the Web runs the same way at all points of the Internet. We shall discuss ways of bridging devices that are *not* IP-capable to the Web. A Web-Based Nomadic Computing System (HP000074 – 87) at 2 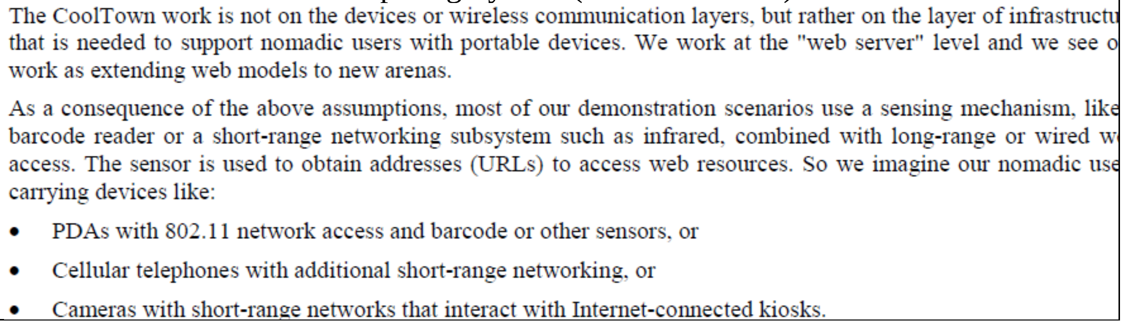 The CoolTown work is not on the devices or wireless communication layers, but rather on the layer of infrastructu that is needed to support nomadic users with portable devices. We work at the "web server" level and we see o work as extending web models to new arenas. As a consequence of the above assumptions, most of our demonstration scenarios use a sensing mechanism, like barcode reader or a short-range networking subsystem such as infrared, combined with long-range or wired w access. The sensor is used to obtain addresses (URLs) to access web resources. So we imagine our nomadic use carrying devices like: • PDAs with 802.11 network access and barcode or other sensors, or • Cellular telephones with additional short-range networking, or • Cameras with short-range networks that interact with Internet-connected kiosks.

Element	'036 Claim Recitation	Exemplary Citations in Reference
		The Cooltown: People, Places, Things Web Presence for the Real World (HP000025 – 34) at 24: 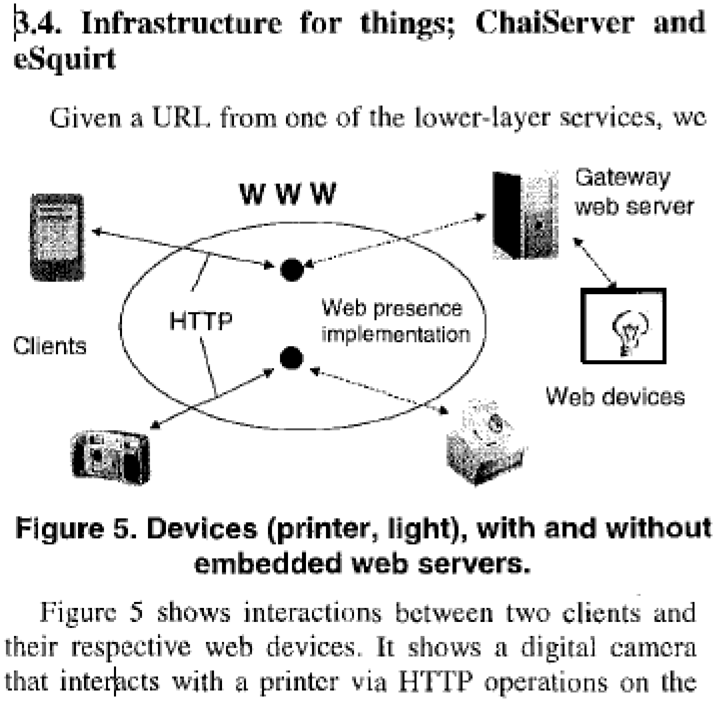 **3.4. Infrastructure for things; ChaiServer and eSquirt** Given a URL from one of the lower-layer services, we **Figure 5. Devices (printer, light), with and without embedded web servers.** Figure 5 shows interactions between two clients and their respective web devices. It shows a digital camera that interacts with a printer via HTTP operations on the

Element	'036 Claim Recitation	Exemplary Citations in Reference
		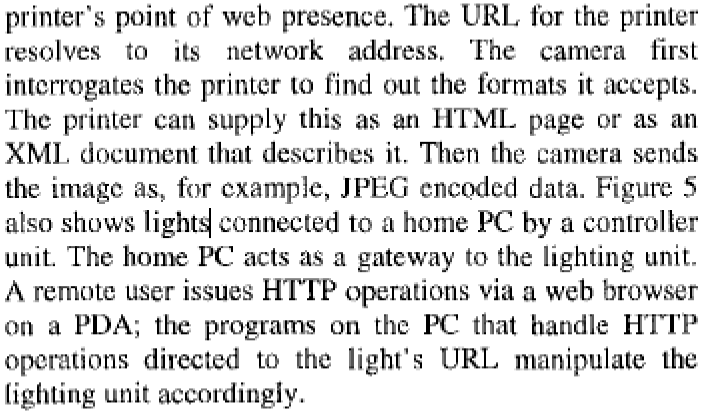 printer's point of web presence. The URL for the printer resolves to its network address. The camera first interrogates the printer to find out the formats it accepts. The printer can supply this as an HTML page or as an XML document that describes it. Then the camera sends the image as, for example, JPEG encoded data. Figure 5 also shows lights connected to a home PC by a controller unit. The home PC acts as a gateway to the lighting unit. A remote user issues HTTP operations via a web browser on a PDA; the programs on the PC that handle HTTP operations directed to the light's URL manipulate the lighting unit accordingly. The Cooltown User Experience (HP000088 – 093) at 1: 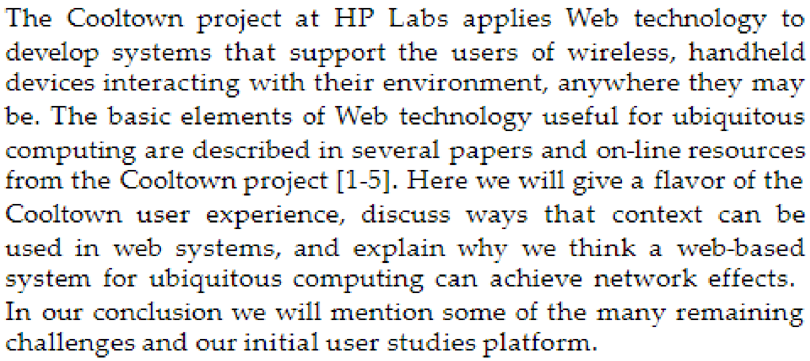 The Cooltown project at HP Labs applies Web technology to develop systems that support the users of wireless, handheld devices interacting with their environment, anywhere they may be. The basic elements of Web technology useful for ubiquitous computing are described in several papers and on-line resources from the Cooltown project [1-5]. Here we will give a flavor of the Cooltown user experience, discuss ways that context can be used in web systems, and explain why we think a web-based system for ubiquitous computing can achieve network effects. In our conclusion we will mention some of the many remaining challenges and our initial user studies platform. A Web-Based Nomadic Computing System (HP000074 – 87) at 2

Element	'036 Claim Recitation	Exemplary Citations in Reference
		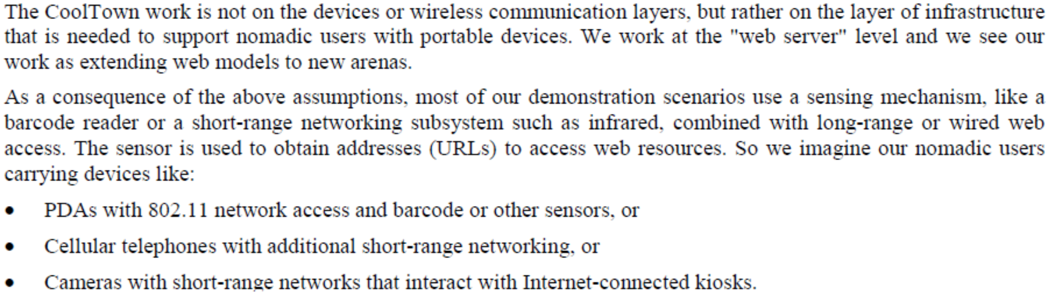 The CoolTown work is not on the devices or wireless communication layers, but rather on the layer of infrastructure that is needed to support nomadic users with portable devices. We work at the "web server" level and we see our work as extending web models to new arenas. As a consequence of the above assumptions, most of our demonstration scenarios use a sensing mechanism, like a barcode reader or a short-range networking subsystem such as infrared, combined with long-range or wired web access. The sensor is used to obtain addresses (URLs) to access web resources. So we imagine our nomadic users carrying devices like: • PDAs with 802.11 network access and barcode or other sensors, or • Cellular telephones with additional short-range networking, or • Cameras with short-range networks that interact with Internet-connected kiosks. A Web-Based Nomadic Computing System (HP000074 – 87) at 3 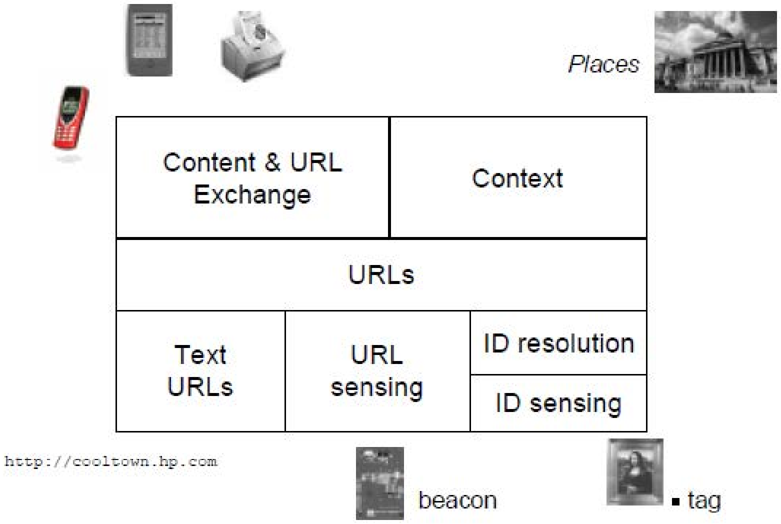 Figure 2. Component relationships. Sources of URLs, including keyboards, beacons, and resolved tags point to context or content. A Web-Based Nomadic Computing System (HP000074 – 87) at 4

Element	'036 Claim Recitation	Exemplary Citations in Reference
		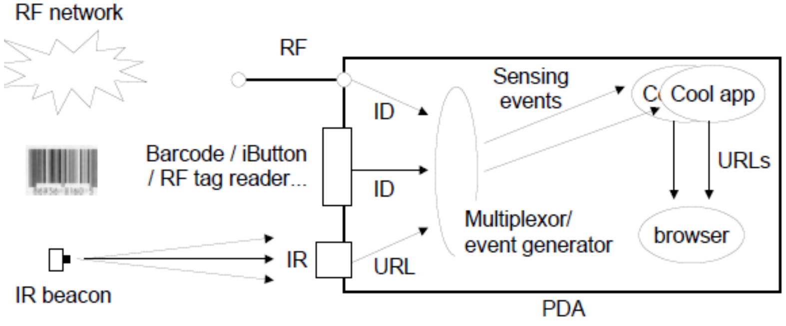 Figure 3. A sketch of a PDA equipped for sensing URLs and identifiers. To the extent it is found that CoolTown does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that CoolTown does not anticipate this claim, CoolTown renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine CoolTown with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1B.ii]	receives the digital representation from the mobile device; and	CoolTown discloses an identification platform coupled with the target database that receives the digital representation from the mobile device. A Web-Based Nomadic Computing System (HP000074 – 87) at 2

Element	'036 Claim Recitation	Exemplary Citations in Reference
		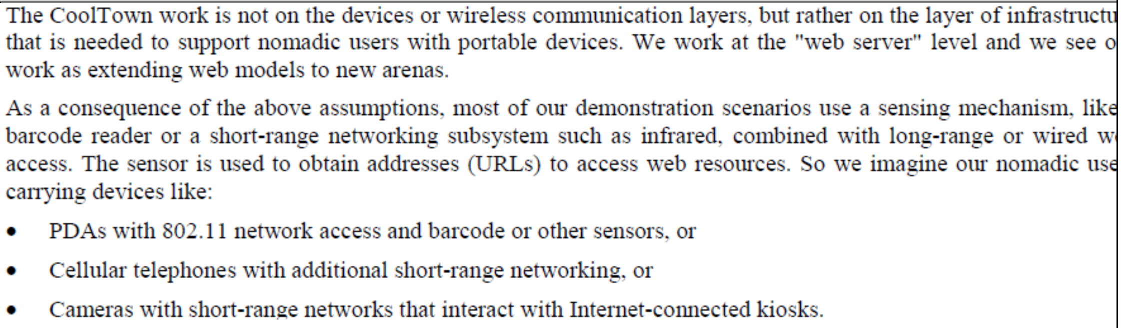 The CoolTown work is not on the devices or wireless communication layers, but rather on the layer of infrastructu that is needed to support nomadic users with portable devices. We work at the "web server" level and we see o work as extending web models to new arenas. As a consequence of the above assumptions, most of our demonstration scenarios use a sensing mechanism, like barcode reader or a short-range networking subsystem such as infrared, combined with long-range or wired w access. The sensor is used to obtain addresses (URLs) to access web resources. So we imagine our nomadic use carrying devices like: • PDAs with 802.11 network access and barcode or other sensors, or • Cellular telephones with additional short-range networking, or • Cameras with short-range networks that interact with Internet-connected kiosks. A Web-Based Nomadic Computing System (HP000074 – 87) at 4 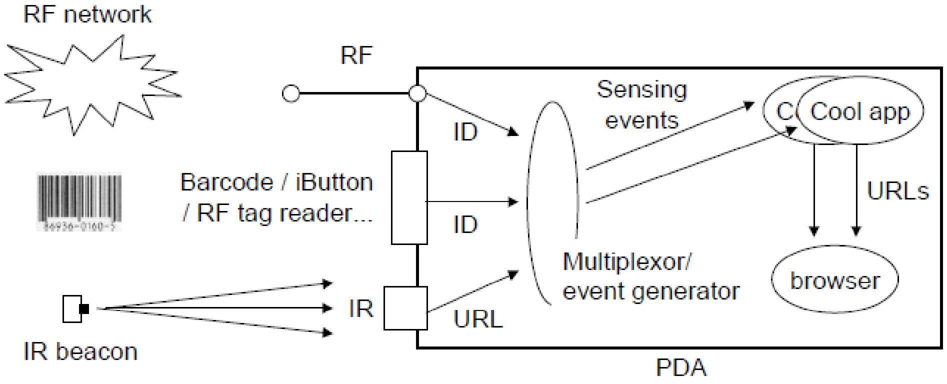 Figure 3. A sketch of a PDA equipped for sensing URLs and identifiers.

Element	'036 Claim Recitation	Exemplary Citations in Reference
		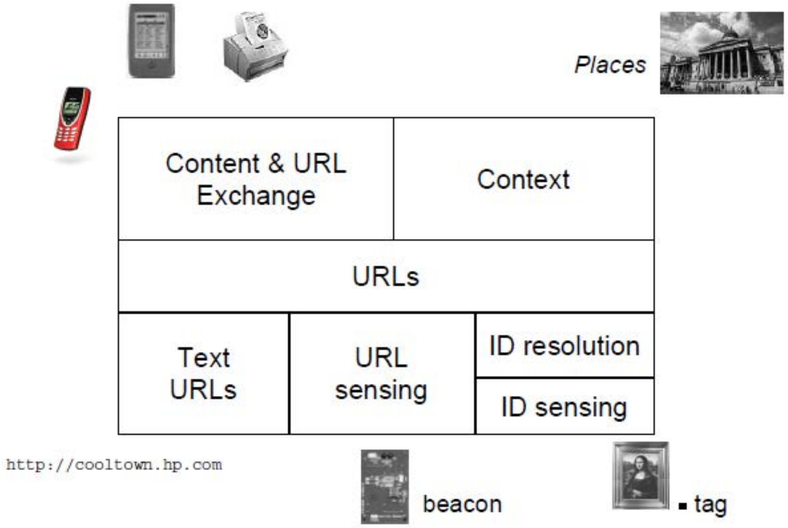 Figure 2. Component relationships. Sources of URLs, including keyboards, beacons, and resolved tags point to context or content. A Web-Based Nomadic Computing System (HP000074 – 87) at 4 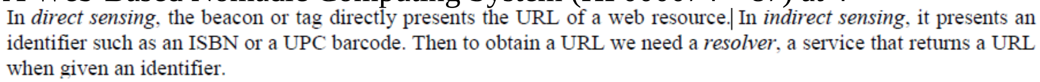 In *direct sensing*, the beacon or tag directly presents the URL of a web resource. In *indirect sensing*, it presents an identifier such as an ISBN or a UPC barcode. Then to obtain a URL we need a *resolver*, a service that returns a URL when given an identifier. A Web-Based Nomadic Computing System (HP000074 – 87) at 5 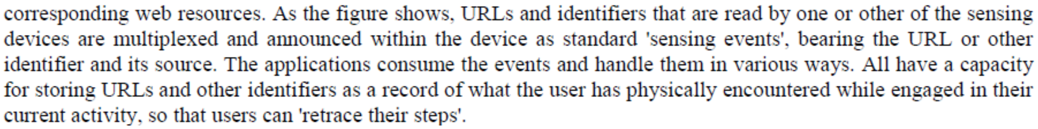 Figure 3 sketches a PDA with several sensing mechanisms (IR, barcode and RF) similar to prototypes we have developed in the CoolTown project. The device runs a web browser, and one or more CoolTown applications. The latter are simple applications that obtain URLs and present them as links, using the standard browser to display the corresponding web resources. As the figure shows, URLs and identifiers that are read by one or other of the sensing devices are multiplexed and announced within the device as standard 'sensing events', bearing the URL or other identifier and its source. The applications consume the events and handle them in various ways. All have a capacity for storing URLs and other identifiers as a record of what the user has physically encountered while engaged in their current activity, so that users can 'retrace their steps'. A Web-Based Nomadic Computing System (HP000074 – 87) at 4

Element	'036 Claim Recitation	Exemplary Citations in Reference
		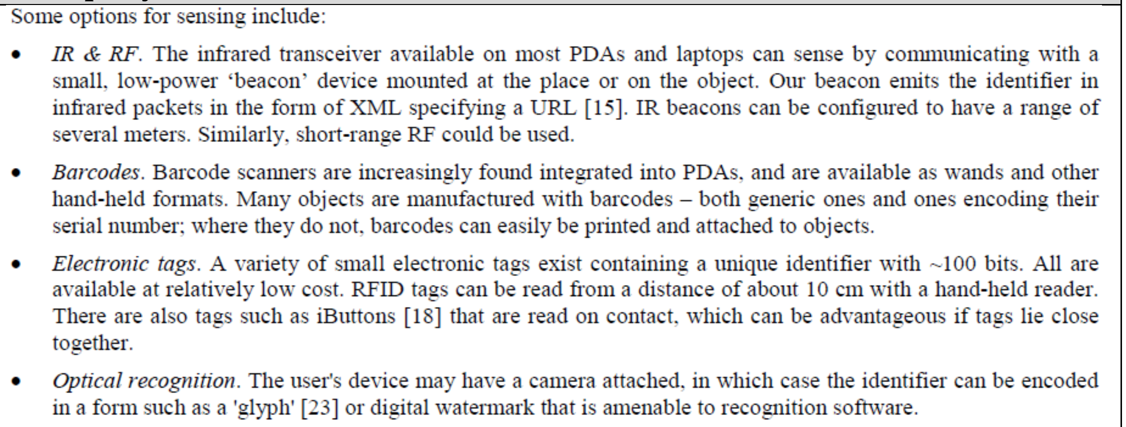 Some options for sensing include: - *IR & RF*. The infrared transceiver available on most PDAs and laptops can sense by communicating with a small, low-power 'beacon' device mounted at the place or on the object. Our beacon emits the identifier in infrared packets in the form of XML specifying a URL [15]. IR beacons can be configured to have a range of several meters. Similarly, short-range RF could be used. - *Barcodes*. Barcode scanners are increasingly found integrated into PDAs, and are available as wands and other hand-held formats. Many objects are manufactured with barcodes – both generic ones and ones encoding their serial number; where they do not, barcodes can easily be printed and attached to objects. - *Electronic tags*. A variety of small electronic tags exist containing a unique identifier with ~100 bits. All are available at relatively low cost. RFID tags can be read from a distance of about 10 cm with a hand-held reader. There are also tags such as iButtons [18] that are read on contact, which can be advantageous if tags lie close together. - *Optical recognition*. The user's device may have a camera attached, in which case the identifier can be encoded in a form such as a 'glyph' [23] or digital watermark that is amenable to recognition software. To the extent it is found that CoolTown does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that CoolTown does not anticipate this claim, CoolTown renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine CoolTown with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1B.iii]	recognizes the target as a known target from the target database based on comparing parameters derived from the digital representation to recognition parameters associated with the known targets; and	CoolTown discloses recognizes the target as a known target from the target database based on comparing parameters derived from the digital representation to recognition parameters associated with the known targets. A Web-Based Nomadic Computing System (HP000074 – 87) at 4

Element	'036 Claim Recitation	Exemplary Citations in Reference
		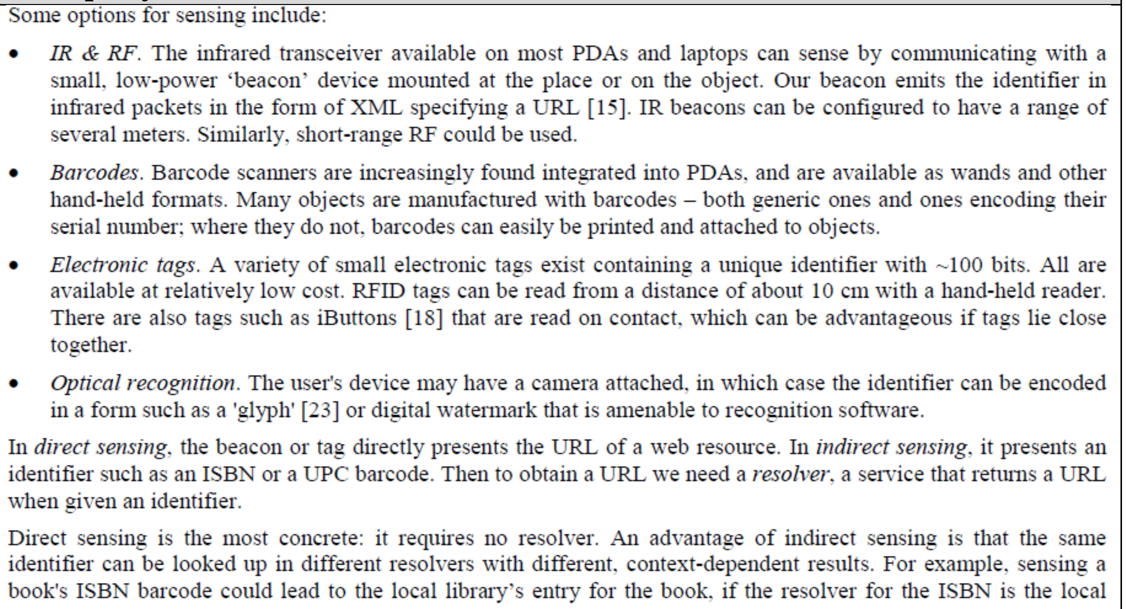 Some options for sensing include: • *IR & RF*. The infrared transceiver available on most PDAs and laptops can sense by communicating with a small, low-power 'beacon' device mounted at the place or on the object. Our beacon emits the identifier in infrared packets in the form of XML specifying a URL [15]. IR beacons can be configured to have a range of several meters. Similarly, short-range RF could be used. • *Barcodes*. Barcode scanners are increasingly found integrated into PDAs, and are available as wands and other hand-held formats. Many objects are manufactured with barcodes – both generic ones and ones encoding their serial number; where they do not, barcodes can easily be printed and attached to objects. • *Electronic tags*. A variety of small electronic tags exist containing a unique identifier with ~100 bits. All are available at relatively low cost. RFID tags can be read from a distance of about 10 cm with a hand-held reader. There are also tags such as iButtons [18] that are read on contact, which can be advantageous if tags lie close together. • *Optical recognition*. The user's device may have a camera attached, in which case the identifier can be encoded in a form such as a 'glyph' [23] or digital watermark that is amenable to recognition software. In *direct sensing*, the beacon or tag directly presents the URL of a web resource. In *indirect sensing*, it presents an identifier such as an ISBN or a UPC barcode. Then to obtain a URL we need a *resolver*, a service that returns a URL when given an identifier. Direct sensing is the most concrete: it requires no resolver. An advantage of indirect sensing is that the same identifier can be looked up in different resolvers with different, context-dependent results. For example, sensing a book's ISBN barcode could lead to the local library's entry for the book, if the resolver for the ISBN is the local A Web-Based Nomadic Computing System (HP000074 – 87) at 5 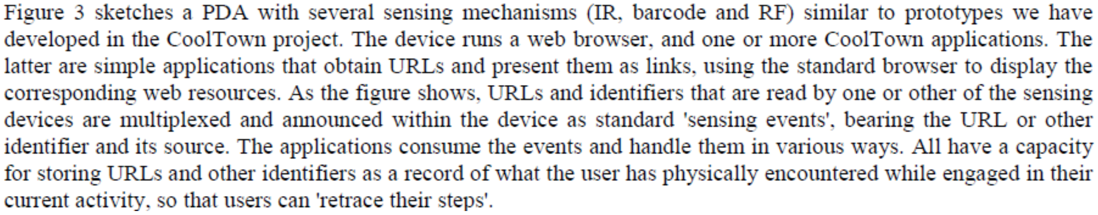 Figure 3 sketches a PDA with several sensing mechanisms (IR, barcode and RF) similar to prototypes we have developed in the CoolTown project. The device runs a web browser, and one or more CoolTown applications. The latter are simple applications that obtain URLs and present them as links, using the standard browser to display the corresponding web resources. As the figure shows, URLs and identifiers that are read by one or other of the sensing devices are multiplexed and announced within the device as standard 'sensing events', bearing the URL or other identifier and its source. The applications consume the events and handle them in various ways. All have a capacity for storing URLs and other identifiers as a record of what the user has physically encountered while engaged in their current activity, so that users can 'retrace their steps'. A Web-Based Nomadic Computing System (HP000074 – 87) at 6 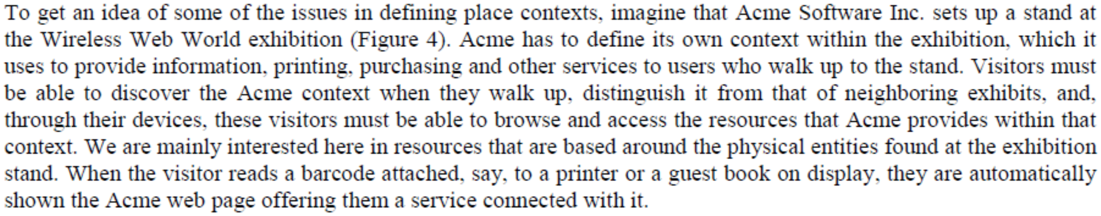 To get an idea of some of the issues in defining place contexts, imagine that Acme Software Inc. sets up a stand at the Wireless Web World exhibition (Figure 4). Acme has to define its own context within the exhibition, which it uses to provide information, printing, purchasing and other services to users who walk up to the stand. Visitors must be able to discover the Acme context when they walk up, distinguish it from that of neighboring exhibits, and, through their devices, these visitors must be able to browse and access the resources that Acme provides within that context. We are mainly interested here in resources that are based around the physical entities found at the exhibition stand. When the visitor reads a barcode attached, say, to a printer or a guest book on display, they are automatically shown the Acme web page offering them a service connected with it.

Element	'036 Claim Recitation	Exemplary Citations in Reference
		A Web-Based Nomadic Computing System (HP000074 – 87) at 6 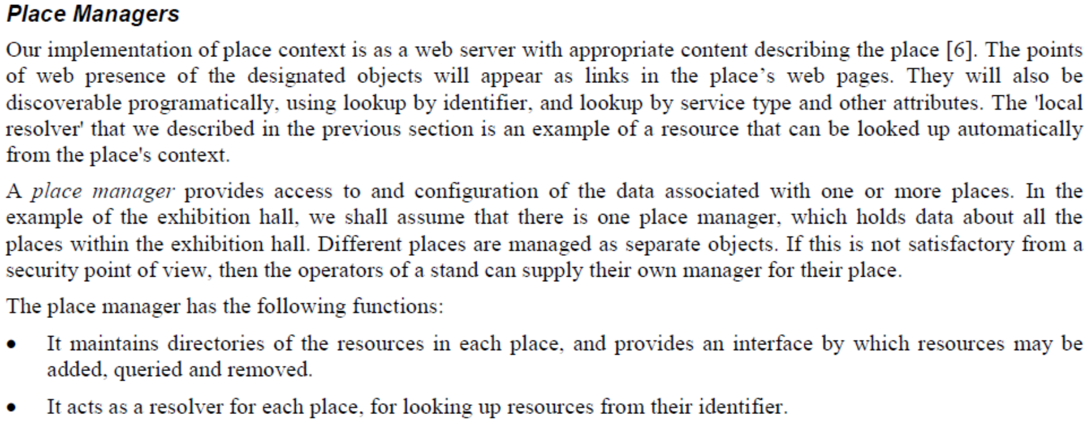 ***Place Managers*** Our implementation of place context is as a web server with appropriate content describing the place [6]. The points of web presence of the designated objects will appear as links in the place's web pages. They will also be discoverable programatically, using lookup by identifier, and lookup by service type and other attributes. The 'local resolver' that we described in the previous section is an example of a resource that can be looked up automatically from the place's context. A *place manager* provides access to and configuration of the data associated with one or more places. In the example of the exhibition hall, we shall assume that there is one place manager, which holds data about all the places within the exhibition hall. Different places are managed as separate objects. If this is not satisfactory from a security point of view, then the operators of a stand can supply their own manager for their place. The place manager has the following functions: • It maintains directories of the resources in each place, and provides an interface by which resources may be added, queried and removed. • It acts as a resolver for each place, for looking up resources from their identifier. Cooltown Developer's Network, People, Place, and Things: web presence for the real world (HP000025–34) (HP000025–34) at 8 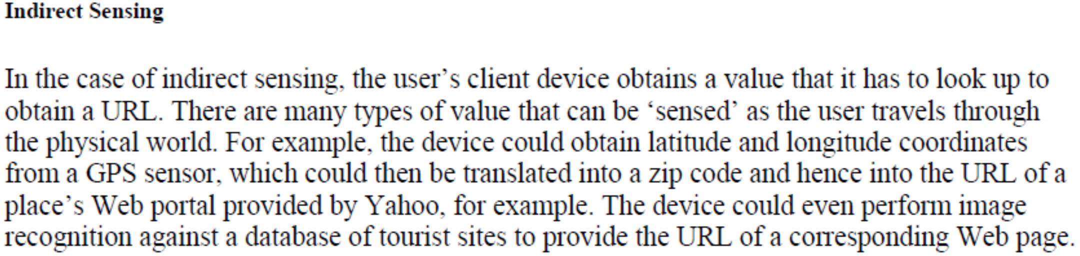 **Indirect Sensing** In the case of indirect sensing, the user's client device obtains a value that it has to look up to obtain a URL. There are many types of value that can be 'sensed' as the user travels through the physical world. For example, the device could obtain latitude and longitude coordinates from a GPS sensor, which could then be translated into a zip code and hence into the URL of a place's Web portal provided by Yahoo, for example. The device could even perform image recognition against a database of tourist sites to provide the URL of a corresponding Web page. Cooltown Developer's Network, People, Place, and Things: web presence for the real world (HP000025–34) (HP000025–34) at 23

Element	'036 Claim Recitation	Exemplary Citations in Reference
		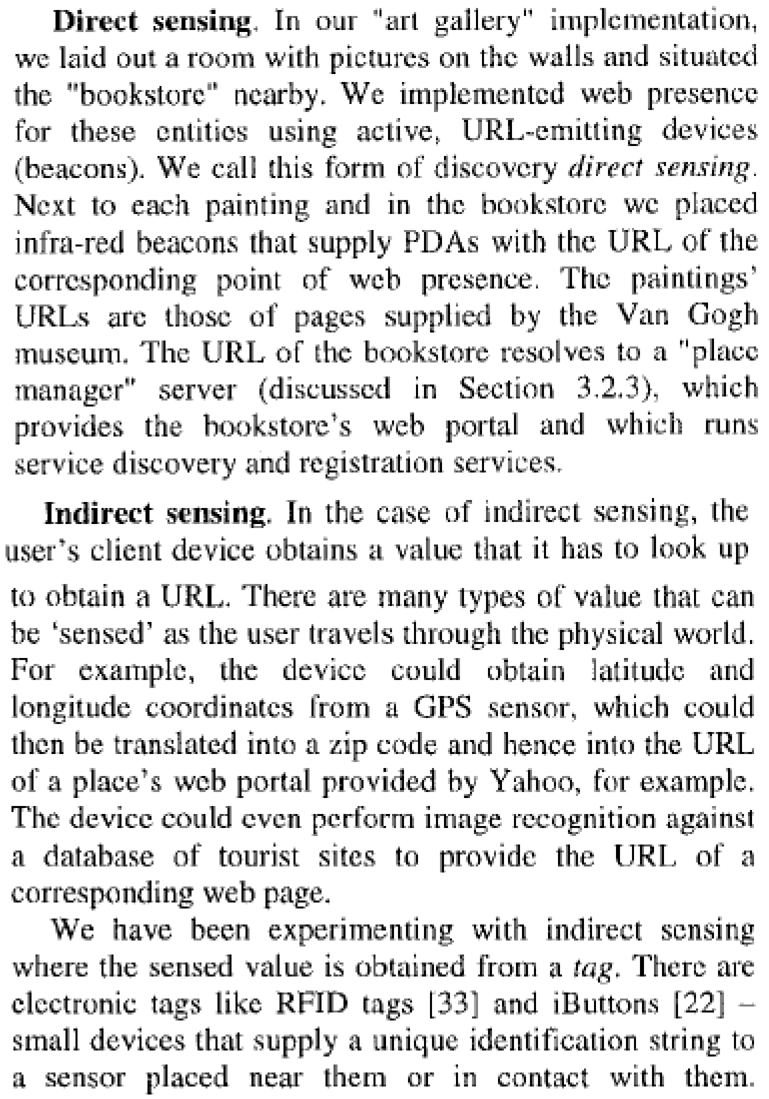 **Direct sensing**. In our "art gallery" implementation, we laid out a room with pictures on the walls and situated the "bookstore" nearby. We implemented web presence for these entities using active, URL-emitting devices (beacons). We call this form of discovery *direct sensing*. Next to each painting and in the bookstore we placed infra-red beacons that supply PDAs with the URL of the corresponding point of web presence. The paintings' URLs are those of pages supplied by the Van Gogh museum. The URL of the bookstore resolves to a "place manager" server (discussed in Section 3.2.3), which provides the bookstore's web portal and which runs service discovery and registration services. **Indirect sensing**. In the case of indirect sensing, the user's client device obtains a value that it has to look up to obtain a URL. There are many types of value that can be 'sensed' as the user travels through the physical world. For example, the device could obtain latitude and longitude coordinates from a GPS sensor, which could then be translated into a zip code and hence into the URL of a place's web portal provided by Yahoo, for example. The device could even perform image recognition against a database of tourist sites to provide the URL of a corresponding web page. We have been experimenting with indirect sensing where the sensed value is obtained from a *tag*. There are electronic tags like RFID tags [33] and iButtons [22] – small devices that supply a unique identification string to a sensor placed near them or in contact with them.

Element	'036 Claim Recitation	Exemplary Citations in Reference
		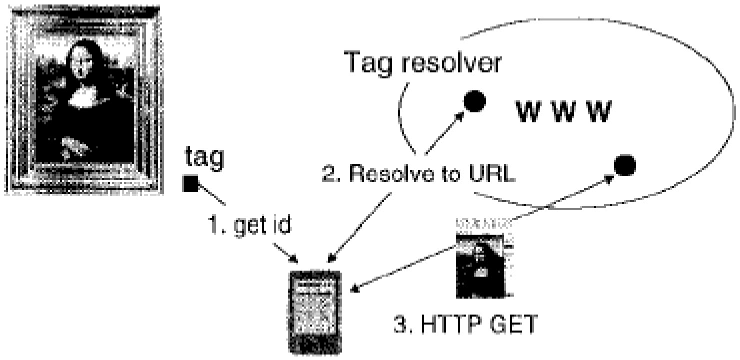 **Figure 4. Indirect sensing: the web presence for a painting.** Barcodes can be used instead of an electronic tag. Figure 4 shows a different art gallery demonstrator in which a PDA client senses a painting's tag. A correlated web page is resolved from the tag identifier and the client presents it to the user. We have produced a prototype implementation for sensing tagged things and displaying a corresponding context-dependent web page. The sensing device knows the URL of a *resolver*. When the client senses a tag, it sends the identifier to the resolver. If the resolver has an entry for the identifier, it sends back the corresponding URL. The user may choose a resolver that reflects his or her personal requirements, or those of a group of users with interests or tasks in common. For example, visitors to an art gallery could be offered resolvers that provide the URLs of pages in their native language, or pages provided by the art school at which they study.

Element	**'036 Claim Recitation**	**Exemplary Citations in Reference**
		To the extent it is found that CoolTown does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that CoolTown does not anticipate this claim, CoolTown renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine CoolTown with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1C]	a content service coupled with the identification platform, and that:	CoolTown discloses a content service coupled with the identification platform. A Web-Based Nomadic Computing System (HP000074 – 87) at Fig. 2

Element	'036 Claim Recitation	Exemplary Citations in Reference
		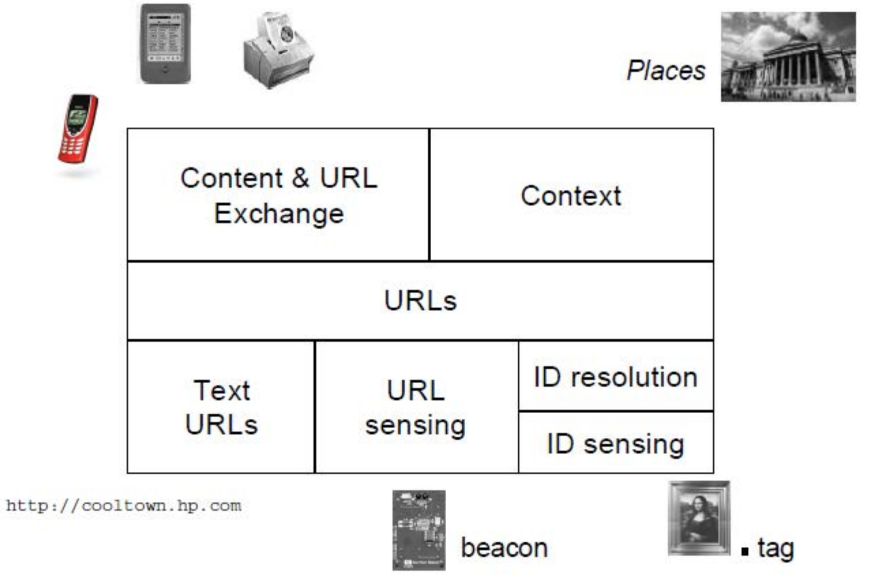 Figure 2. Component relationships. Sources of URLs, including keyboards, beacons, and resolved tags point to context or content. A Web-Based Nomadic Computing System (HP000074 – 87) at 3 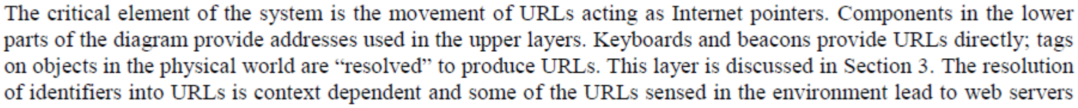 The critical element of the system is the movement of URLs acting as Internet pointers. Components in the lower parts of the diagram provide addresses used in the upper layers. Keyboards and beacons provide URLs directly; tags on objects in the physical world are "resolved" to produce URLs. This layer is discussed in Section 3. The resolution of identifiers into URLs is context dependent and some of the URLs sensed in the environment lead to web servers A Web-Based Nomadic Computing System (HP000074 – 87) at 4

Element	’036 Claim Recitation	Exemplary Citations in Reference
		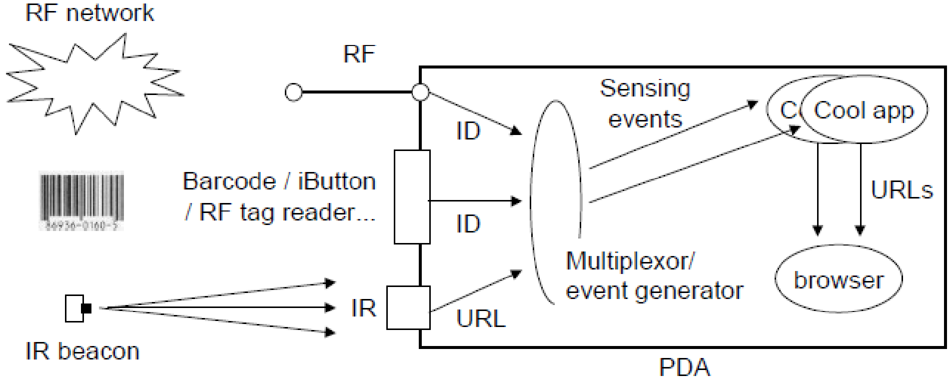 Figure 3. A sketch of a PDA equipped for sensing URLs and identifiers. A Web-Based Nomadic Computing System (HP000074 – 87) at 4–5 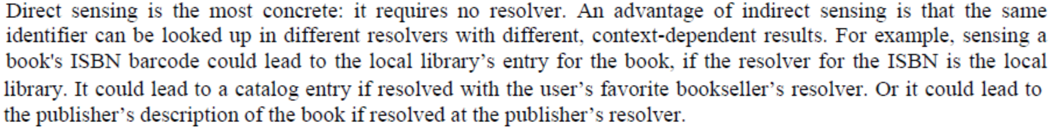Direct sensing is the most concrete: it requires no resolver. An advantage of indirect sensing is that the same identifier can be looked up in different resolvers with different, context-dependent results. For example, sensing a book's ISBN barcode could lead to the local library’s entry for the book, if the resolver for the ISBN is the local library. It could lead to a catalog entry if resolved with the user’s favorite bookseller’s resolver. Or it could lead to the publisher’s description of the book if resolved at the publisher’s resolver. A Web-Based Nomadic Computing System (HP000074 – 87) at 5 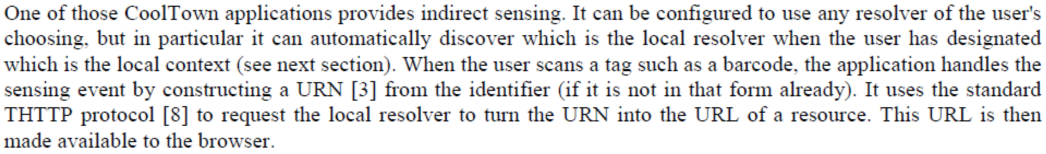One of those CoolTown applications provides indirect sensing. It can be configured to use any resolver of the user's choosing, but in particular it can automatically discover which is the local resolver when the user has designated which is the local context (see next section). When the user scans a tag such as a barcode, the application handles the sensing event by constructing a URN [3] from the identifier (if it is not in that form already). It uses the standard THTTP protocol [8] to request the local resolver to turn the URN into the URL of a resource. This URL is then made available to the browser. A Web-Based Nomadic Computing System (HP000074 – 87) at 6 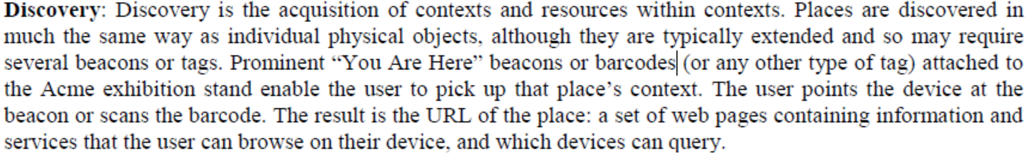**Discovery**: Discovery is the acquisition of contexts and resources within contexts. Places are discovered in much the same way as individual physical objects, although they are typically extended and so may require several beacons or tags. Prominent “You Are Here” beacons or barcodes (or any other type of tag) attached to the Acme exhibition stand enable the user to pick up that place’s context. The user points the device at the beacon or scans the barcode. The result is the URL of the place: a set of web pages containing information and services that the user can browse on their device, and which devices can query. A Web-Based Nomadic Computing System (HP000074 – 87) at 6

Element	'036 Claim Recitation	Exemplary Citations in Reference
		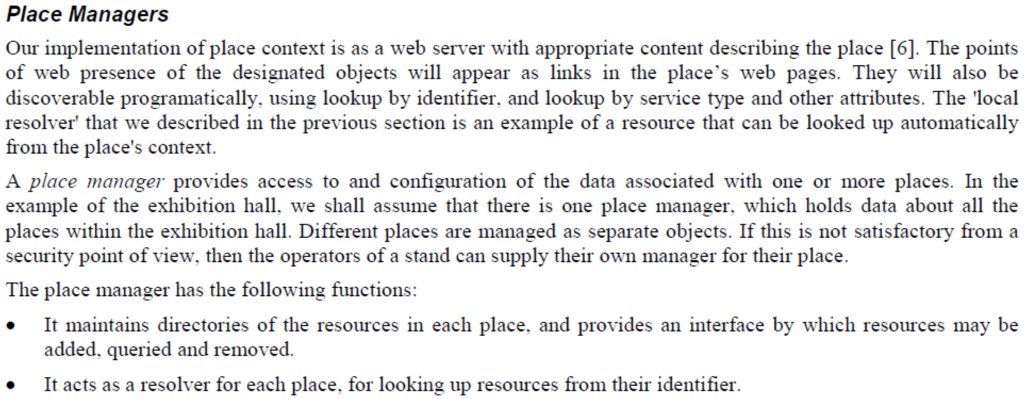 ***Place Managers*** Our implementation of place context is as a web server with appropriate content describing the place [6]. The points of web presence of the designated objects will appear as links in the place's web pages. They will also be discoverable programatically, using lookup by identifier, and lookup by service type and other attributes. The 'local resolver' that we described in the previous section is an example of a resource that can be looked up automatically from the place's context. A *place manager* provides access to and configuration of the data associated with one or more places. In the example of the exhibition hall, we shall assume that there is one place manager, which holds data about all the places within the exhibition hall. Different places are managed as separate objects. If this is not satisfactory from a security point of view, then the operators of a stand can supply their own manager for their place. The place manager has the following functions: • It maintains directories of the resources in each place, and provides an interface by which resources may be added, queried and removed. • It acts as a resolver for each place, for looking up resources from their identifier. Cooltown Developer's Network, People, Place, and Things: web presence for the real world (HP000025–34) (HP000025–34) at 20 – 21 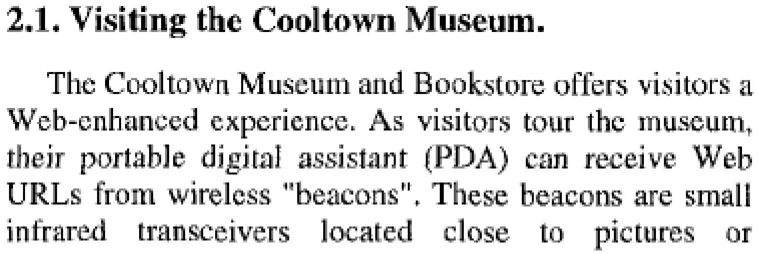 **2.1. Visiting the Cooltown Museum.** The Cooltown Museum and Bookstore offers visitors a Web-enhanced experience. As visitors tour the museum, their portable digital assistant (PDA) can receive Web URLs from wireless "beacons". These beacons are small infrared transceivers located close to pictures or

Element	'036 Claim Recitation	Exemplary Citations in Reference
		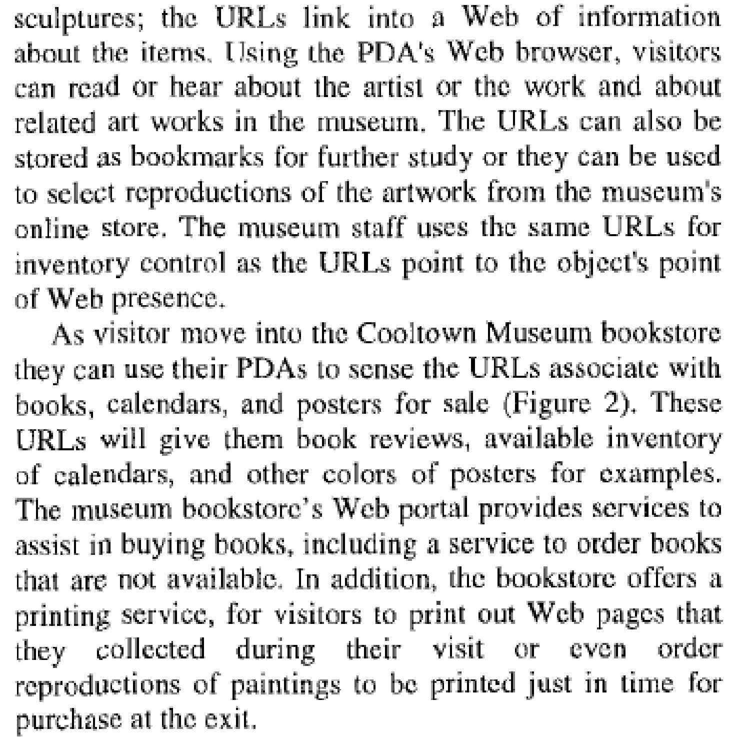 sculptures; the URLs link into a Web of information about the items. Using the PDA's Web browser, visitors can read or hear about the artist or the work and about related art works in the museum. The URLs can also be stored as bookmarks for further study or they can be used to select reproductions of the artwork from the museum's online store. The museum staff uses the same URLs for inventory control as the URLs point to the object's point of Web presence. As visitor move into the Cooltown Museum bookstore they can use their PDAs to sense the URLs associate with books, calendars, and posters for sale (Figure 2). These URLs will give them book reviews, available inventory of calendars, and other colors of posters for examples. The museum bookstore's Web portal provides services to assist in buying books, including a service to order books that are not available. In addition, the bookstore offers a printing service, for visitors to print out Web pages that they collected during their visit or even order reproductions of paintings to be printed just in time for purchase at the exit. Cooltown Developer's Network, People, Place, and Things: web presence for the real world (HP000025–34) (HP000025–34) at 23

Element	'036 Claim Recitation	Exemplary Citations in Reference
		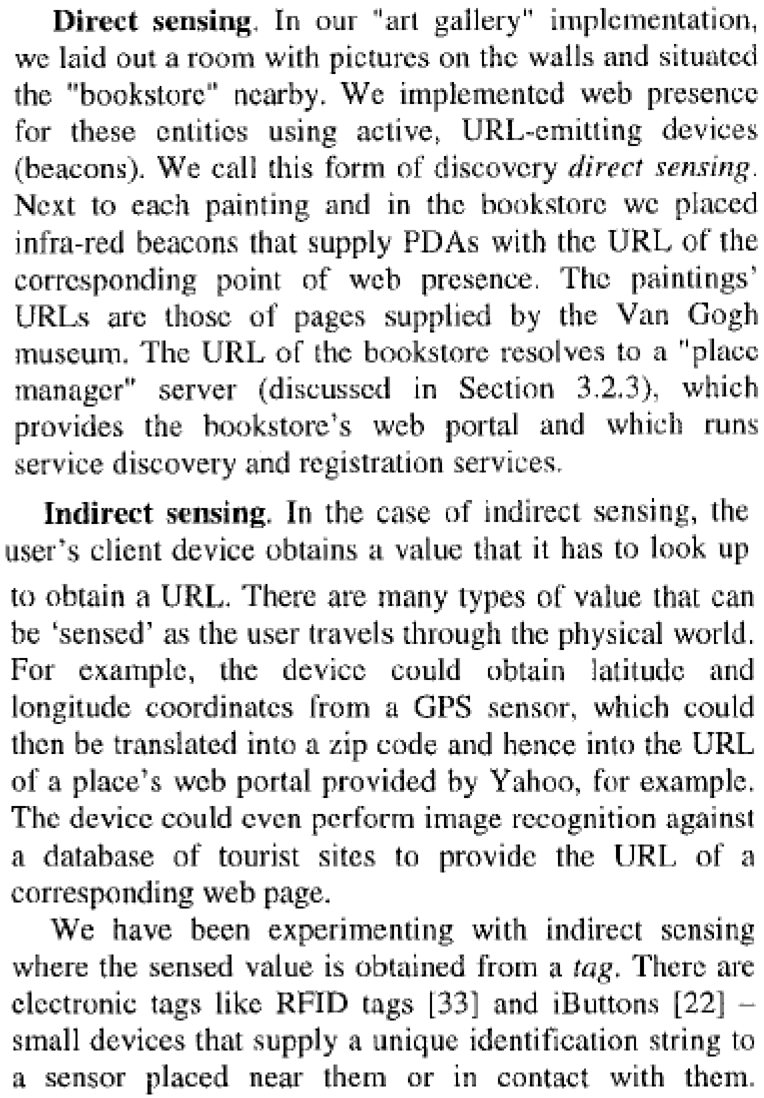 **Direct sensing**. In our "art gallery" implementation, we laid out a room with pictures on the walls and situated the "bookstore" nearby. We implemented web presence for these entities using active, URL-emitting devices (beacons). We call this form of discovery *direct sensing*. Next to each painting and in the bookstore we placed infra-red beacons that supply PDAs with the URL of the corresponding point of web presence. The paintings' URLs are those of pages supplied by the Van Gogh museum. The URL of the bookstore resolves to a "place manager" server (discussed in Section 3.2.3), which provides the bookstore's web portal and which runs service discovery and registration services. **Indirect sensing**. In the case of indirect sensing, the user's client device obtains a value that it has to look up to obtain a URL. There are many types of value that can be 'sensed' as the user travels through the physical world. For example, the device could obtain latitude and longitude coordinates from a GPS sensor, which could then be translated into a zip code and hence into the URL of a place's web portal provided by Yahoo, for example. The device could even perform image recognition against a database of tourist sites to provide the URL of a corresponding web page. We have been experimenting with indirect sensing where the sensed value is obtained from a *tag*. There are electronic tags like RFID tags [33] and iButtons [22] – small devices that supply a unique identification string to a sensor placed near them or in contact with them.

Element	'036 Claim Recitation	Exemplary Citations in Reference
		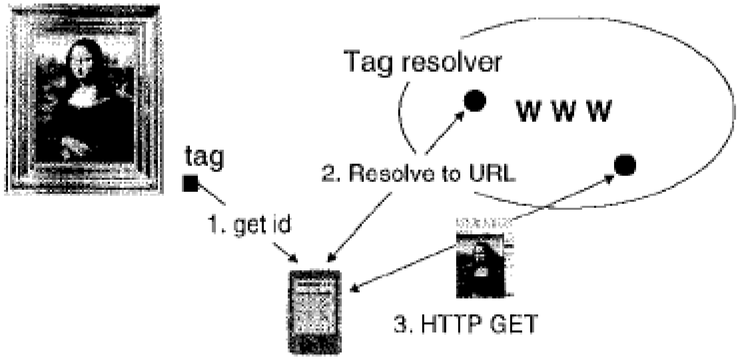 **Figure 4. Indirect sensing: the web presence for a painting.** Barcodes can be used instead of an electronic tag. Figure 4 shows a different art gallery demonstrator in which a PDA client senses a painting's tag. A correlated web page is resolved from the tag identifier and the client presents it to the user. We have produced a prototype implementation for sensing tagged things and displaying a corresponding context-dependent web page. The sensing device knows the URL of a *resolver*. When the client senses a tag, it sends the identifier to the resolver. If the resolver has an entry for the identifier, it sends back the corresponding URL. The user may choose a resolver that reflects his or her personal requirements, or those of a group of users with interests or tasks in common. For example, visitors to an art gallery could be offered resolvers that provide the URLs of pages in their native language, or pages provided by the art school at which they study. The Cooltown User Experience (HP000088 – 093) at 3

Element	'036 Claim Recitation	Exemplary Citations in Reference
		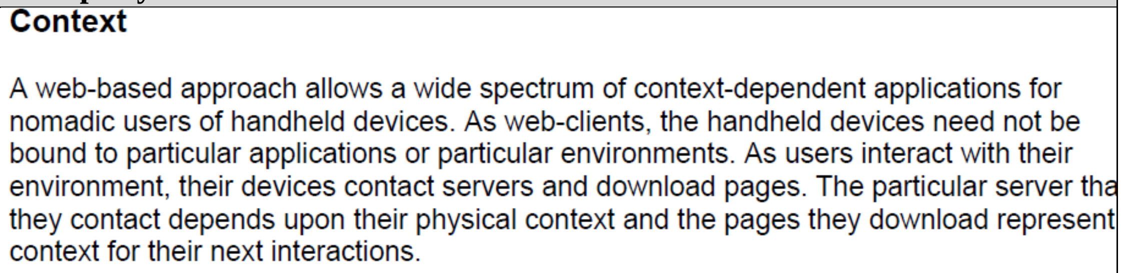 **Context** A web-based approach allows a wide spectrum of context-dependent applications for nomadic users of handheld devices. As web-clients, the handheld devices need not be bound to particular applications or particular environments. As users interact with their environment, their devices contact servers and download pages. The particular server tha they contact depends upon their physical context and the pages they download represent context for their next interactions. To the extent it is found that CoolTown does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that CoolTown does not anticipate this claim, CoolTown renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine CoolTown with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1C.i]	obtains content information related to the known target; and	CoolTown discloses [a content service coupled with the identification platform] that obtains content information related to the known target. A Web-Based Nomadic Computing System (HP000074 – 87) at Fig. 2

Element	'036 Claim Recitation	Exemplary Citations in Reference
		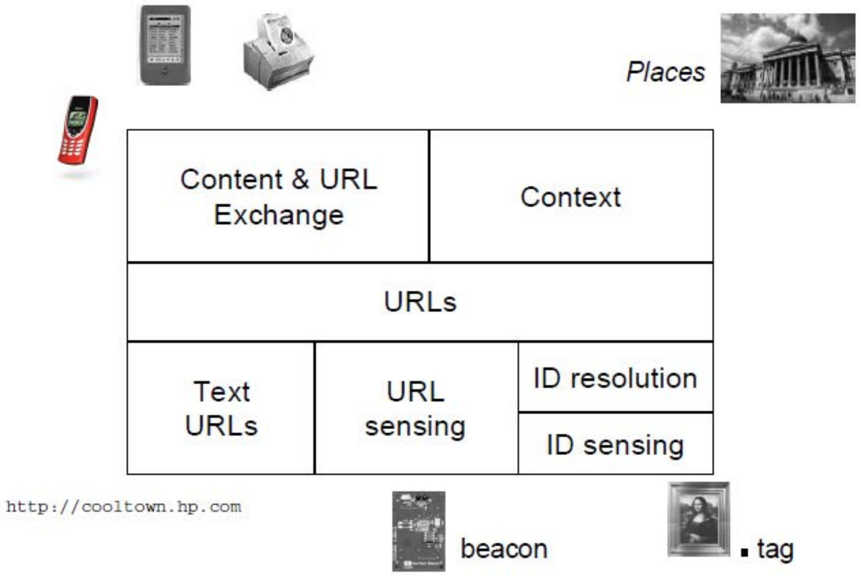 Figure 2. Component relationships. Sources of URLs, including keyboards, beacons, and resolved tags point to context or content. A Web-Based Nomadic Computing System (HP000074 – 87) at 3 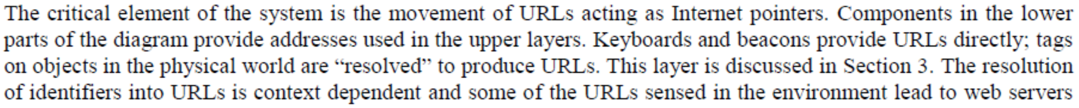 The critical element of the system is the movement of URLs acting as Internet pointers. Components in the lower parts of the diagram provide addresses used in the upper layers. Keyboards and beacons provide URLs directly; tags on objects in the physical world are "resolved" to produce URLs. This layer is discussed in Section 3. The resolution of identifiers into URLs is context dependent and some of the URLs sensed in the environment lead to web servers A Web-Based Nomadic Computing System (HP000074 – 87) at 4

Element	'036 Claim Recitation	Exemplary Citations in Reference
		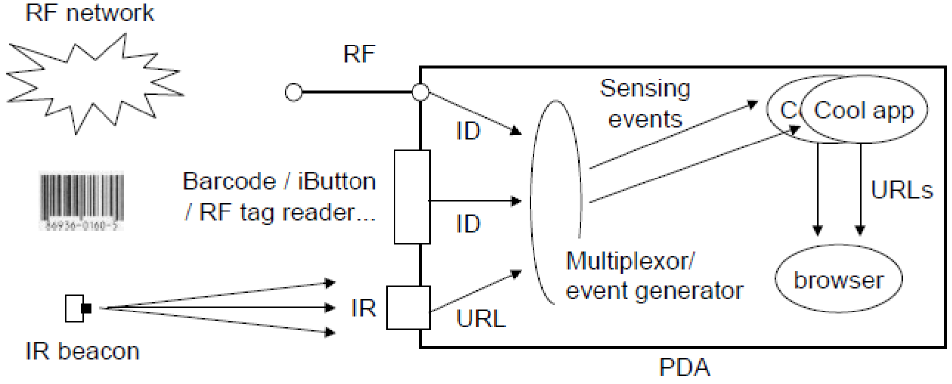 Figure 3. A sketch of a PDA equipped for sensing URLs and identifiers. A Web-Based Nomadic Computing System (HP000074 – 87) at 4–5 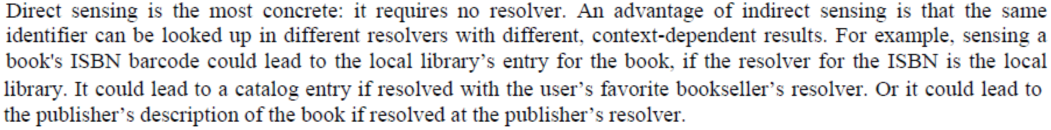 Direct sensing is the most concrete: it requires no resolver. An advantage of indirect sensing is that the same identifier can be looked up in different resolvers with different, context-dependent results. For example, sensing a book's ISBN barcode could lead to the local library's entry for the book, if the resolver for the ISBN is the local library. It could lead to a catalog entry if resolved with the user's favorite bookseller's resolver. Or it could lead to the publisher's description of the book if resolved at the publisher's resolver. A Web-Based Nomadic Computing System (HP000074 – 87) at 5 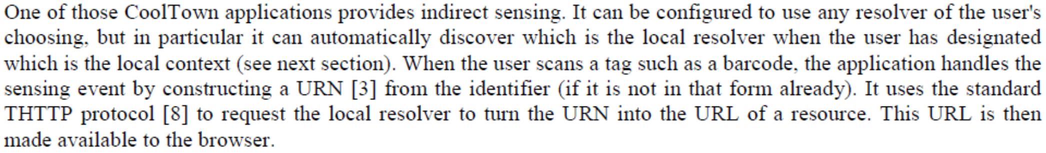 One of those CoolTown applications provides indirect sensing. It can be configured to use any resolver of the user's choosing, but in particular it can automatically discover which is the local resolver when the user has designated which is the local context (see next section). When the user scans a tag such as a barcode, the application handles the sensing event by constructing a URN [3] from the identifier (if it is not in that form already). It uses the standard THTTP protocol [8] to request the local resolver to turn the URN into the URL of a resource. This URL is then made available to the browser. A Web-Based Nomadic Computing System (HP000074 – 87) at 6 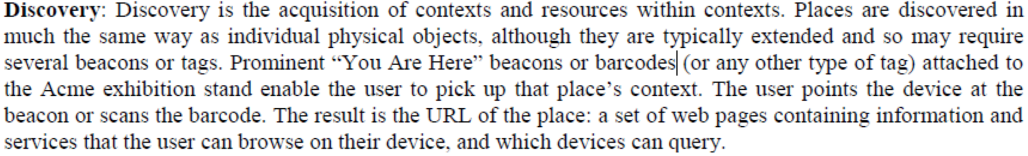 **Discovery**: Discovery is the acquisition of contexts and resources within contexts. Places are discovered in much the same way as individual physical objects, although they are typically extended and so may require several beacons or tags. Prominent "You Are Here" beacons or barcodes (or any other type of tag) attached to the Acme exhibition stand enable the user to pick up that place's context. The user points the device at the beacon or scans the barcode. The result is the URL of the place: a set of web pages containing information and services that the user can browse on their device, and which devices can query. A Web-Based Nomadic Computing System (HP000074 – 87) at 6

Element	'036 Claim Recitation	Exemplary Citations in Reference
		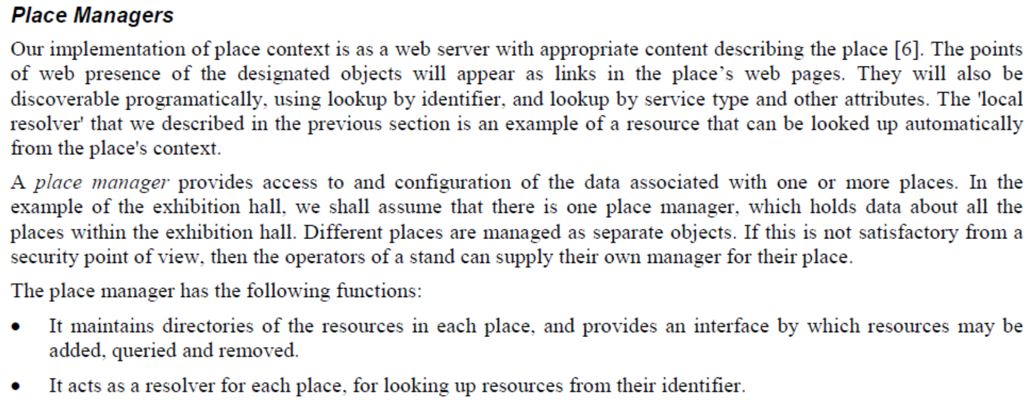 ***Place Managers*** Our implementation of place context is as a web server with appropriate content describing the place [6]. The points of web presence of the designated objects will appear as links in the place's web pages. They will also be discoverable programatically, using lookup by identifier, and lookup by service type and other attributes. The 'local resolver' that we described in the previous section is an example of a resource that can be looked up automatically from the place's context. A *place manager* provides access to and configuration of the data associated with one or more places. In the example of the exhibition hall, we shall assume that there is one place manager, which holds data about all the places within the exhibition hall. Different places are managed as separate objects. If this is not satisfactory from a security point of view, then the operators of a stand can supply their own manager for their place. The place manager has the following functions: • It maintains directories of the resources in each place, and provides an interface by which resources may be added, queried and removed. • It acts as a resolver for each place, for looking up resources from their identifier. Cooltown Developer's Network, People, Place, and Things: web presence for the real world (HP000025–34) (HP000025–34) at 20 – 21 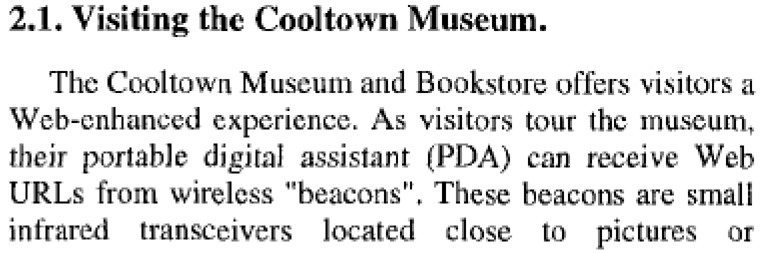 **2.1. Visiting the Cooltown Museum.** The Cooltown Museum and Bookstore offers visitors a Web-enhanced experience. As visitors tour the museum, their portable digital assistant (PDA) can receive Web URLs from wireless "beacons". These beacons are small infrared transceivers located close to pictures or

Element	'036 Claim Recitation	Exemplary Citations in Reference
		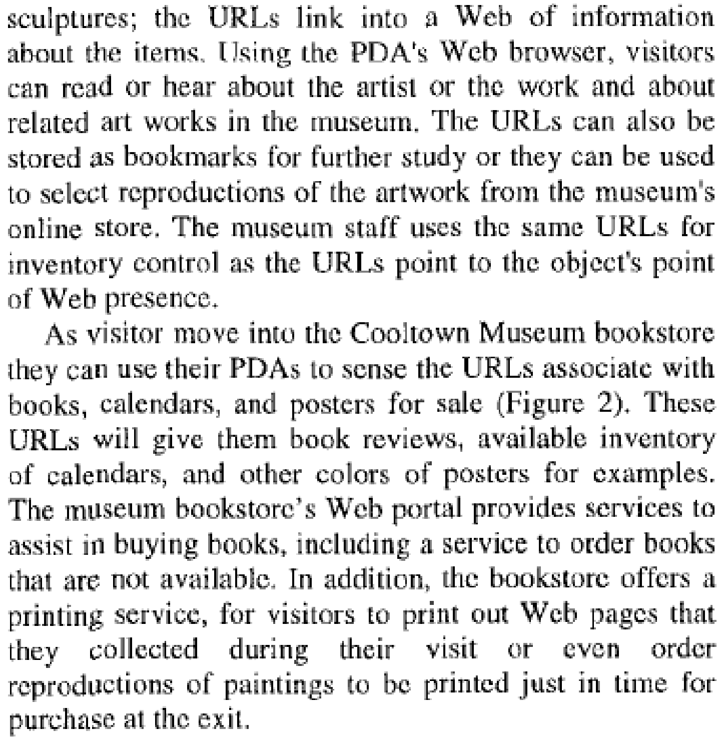 sculptures; the URLs link into a Web of information about the items. Using the PDA's Web browser, visitors can read or hear about the artist or the work and about related art works in the museum. The URLs can also be stored as bookmarks for further study or they can be used to select reproductions of the artwork from the museum's online store. The museum staff uses the same URLs for inventory control as the URLs point to the object's point of Web presence. As visitor move into the Cooltown Museum bookstore they can use their PDAs to sense the URLs associate with books, calendars, and posters for sale (Figure 2). These URLs will give them book reviews, available inventory of calendars, and other colors of posters for examples. The museum bookstore's Web portal provides services to assist in buying books, including a service to order books that are not available. In addition, the bookstore offers a printing service, for visitors to print out Web pages that they collected during their visit or even order reproductions of paintings to be printed just in time for purchase at the exit. Cooltown Developer's Network, People, Place, and Things: web presence for the real world (HP000025–34) (HP000025–34) at 23

Element	'036 Claim Recitation	Exemplary Citations in Reference
		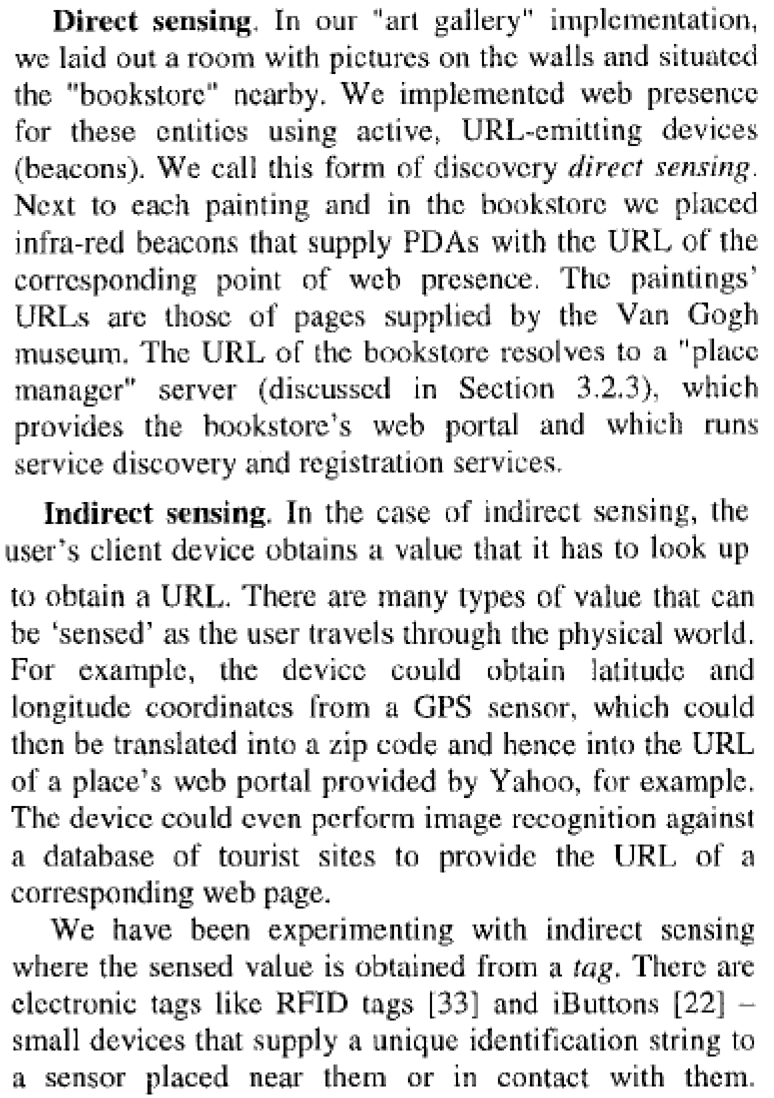 **Direct sensing**. In our "art gallery" implementation, we laid out a room with pictures on the walls and situated the "bookstore" nearby. We implemented web presence for these entities using active, URL-emitting devices (beacons). We call this form of discovery *direct sensing*. Next to each painting and in the bookstore we placed infra-red beacons that supply PDAs with the URL of the corresponding point of web presence. The paintings' URLs are those of pages supplied by the Van Gogh museum. The URL of the bookstore resolves to a "place manager" server (discussed in Section 3.2.3), which provides the bookstore's web portal and which runs service discovery and registration services. **Indirect sensing**. In the case of indirect sensing, the user's client device obtains a value that it has to look up to obtain a URL. There are many types of value that can be 'sensed' as the user travels through the physical world. For example, the device could obtain latitude and longitude coordinates from a GPS sensor, which could then be translated into a zip code and hence into the URL of a place's web portal provided by Yahoo, for example. The device could even perform image recognition against a database of tourist sites to provide the URL of a corresponding web page. We have been experimenting with indirect sensing where the sensed value is obtained from a *tag*. There are electronic tags like RFID tags [33] and iButtons [22] – small devices that supply a unique identification string to a sensor placed near them or in contact with them.

Element	'036 Claim Recitation	Exemplary Citations in Reference
		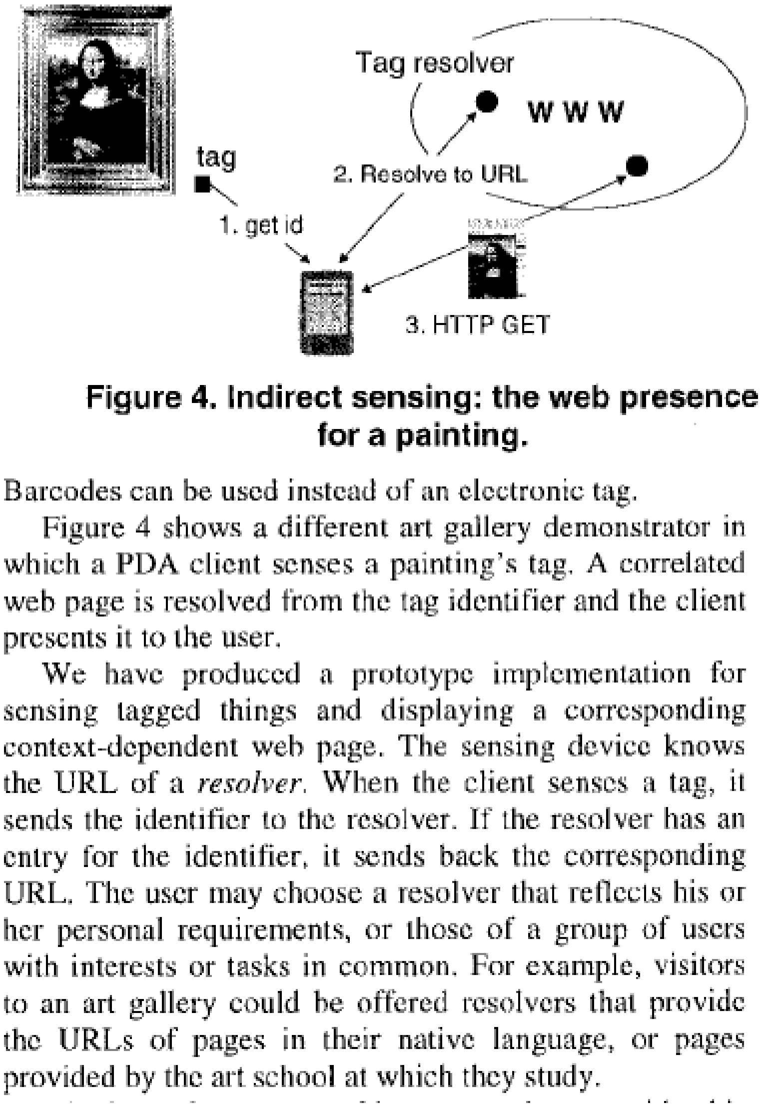 **Figure 4. Indirect sensing: the web presence for a painting.** Barcodes can be used instead of an electronic tag. Figure 4 shows a different art gallery demonstrator in which a PDA client senses a painting's tag. A correlated web page is resolved from the tag identifier and the client presents it to the user. We have produced a prototype implementation for sensing tagged things and displaying a corresponding context-dependent web page. The sensing device knows the URL of a *resolver*. When the client senses a tag, it sends the identifier to the resolver. If the resolver has an entry for the identifier, it sends back the corresponding URL. The user may choose a resolver that reflects his or her personal requirements, or those of a group of users with interests or tasks in common. For example, visitors to an art gallery could be offered resolvers that provide the URLs of pages in their native language, or pages provided by the art school at which they study.

Element	'036 Claim Recitation	Exemplary Citations in Reference
		To the extent it is found that CoolTown does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that CoolTown does not anticipate this claim, CoolTown renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine CoolTown with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1C.ii]	sends the content information to at least one of the identification platform and the mobile device.	CoolTown discloses [a content service coupled with the identification platform] that sends the content information to at least one of the identification platform and the mobile device. A Web-Based Nomadic Computing System (HP000074 – 87) at Fig. 2

Element	**'036 Claim Recitation**	**Exemplary Citations in Reference**
		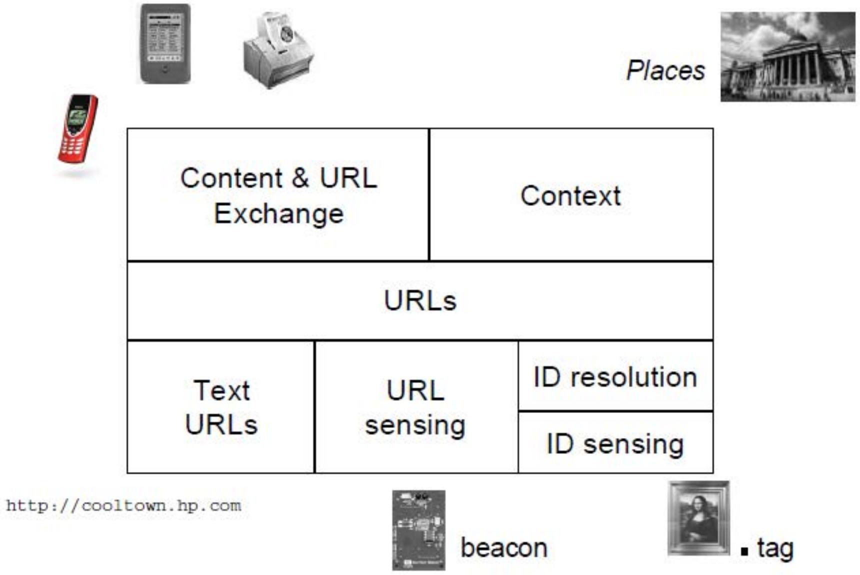 Figure 2. Component relationships. Sources of URLs, including keyboards, beacons, and resolved tags point to context or content. A Web-Based Nomadic Computing System (HP000074 – 87) at 4

Element	'036 Claim Recitation	Exemplary Citations in Reference
		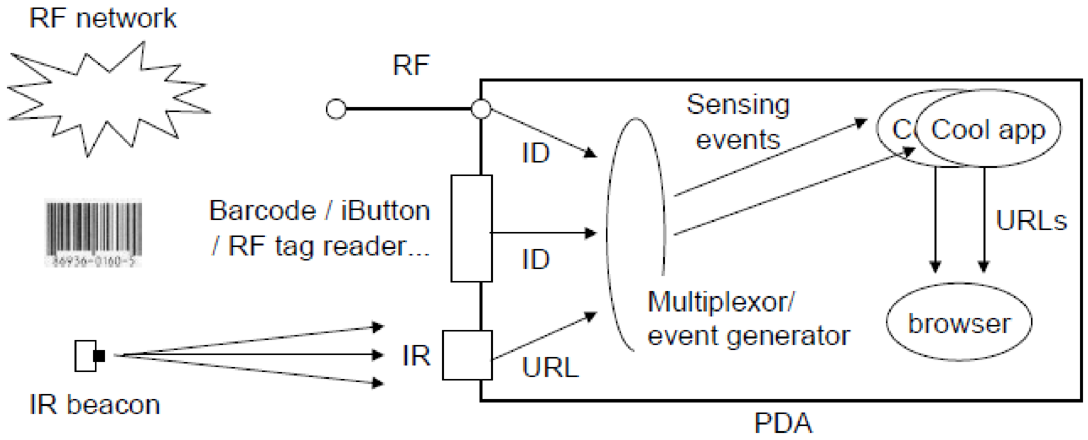 Figure 3. A sketch of a PDA equipped for sensing URLs and identifiers. A Web-Based Nomadic Computing System (HP000074 – 87) at 5 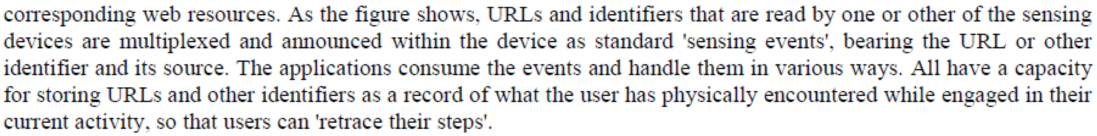 Figure 3 sketches a PDA with several sensing mechanisms (IR, barcode and RF) similar to prototypes we have developed in the CoolTown project. The device runs a web browser, and one or more CoolTown applications. The latter are simple applications that obtain URLs and present them as links, using the standard browser to display the corresponding web resources. As the figure shows, URLs and identifiers that are read by one or other of the sensing devices are multiplexed and announced within the device as standard 'sensing events', bearing the URL or other identifier and its source. The applications consume the events and handle them in various ways. All have a capacity for storing URLs and other identifiers as a record of what the user has physically encountered while engaged in their current activity, so that users can 'retrace their steps'. A Web-Based Nomadic Computing System (HP000074 – 87) at 5 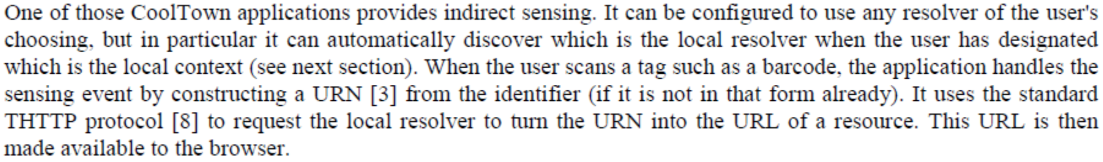 One of those CoolTown applications provides indirect sensing. It can be configured to use any resolver of the user's choosing, but in particular it can automatically discover which is the local resolver when the user has designated which is the local context (see next section). When the user scans a tag such as a barcode, the application handles the sensing event by constructing a URN [3] from the identifier (if it is not in that form already). It uses the standard THTTP protocol [8] to request the local resolver to turn the URN into the URL of a resource. This URL is then made available to the browser. A Web-Based Nomadic Computing System (HP000074 – 87) at 6

Element	'036 Claim Recitation	Exemplary Citations in Reference
		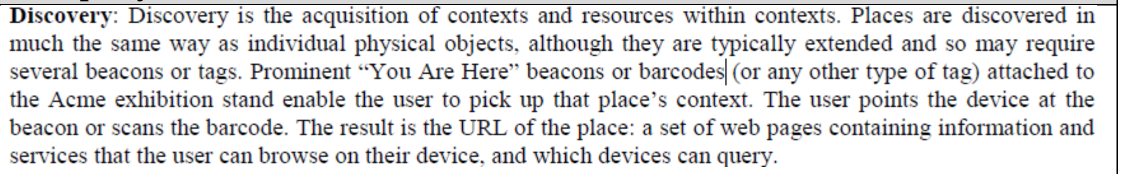 **Discovery**: Discovery is the acquisition of contexts and resources within contexts. Places are discovered in much the same way as individual physical objects, although they are typically extended and so may require several beacons or tags. Prominent "You Are Here" beacons or barcodes (or any other type of tag) attached to the Acme exhibition stand enable the user to pick up that place's context. The user points the device at the beacon or scans the barcode. The result is the URL of the place: a set of web pages containing information and services that the user can browse on their device, and which devices can query. Cooltown Developer's Network, People, Place, and Things: web presence for the real world (HP000025–34) at 5 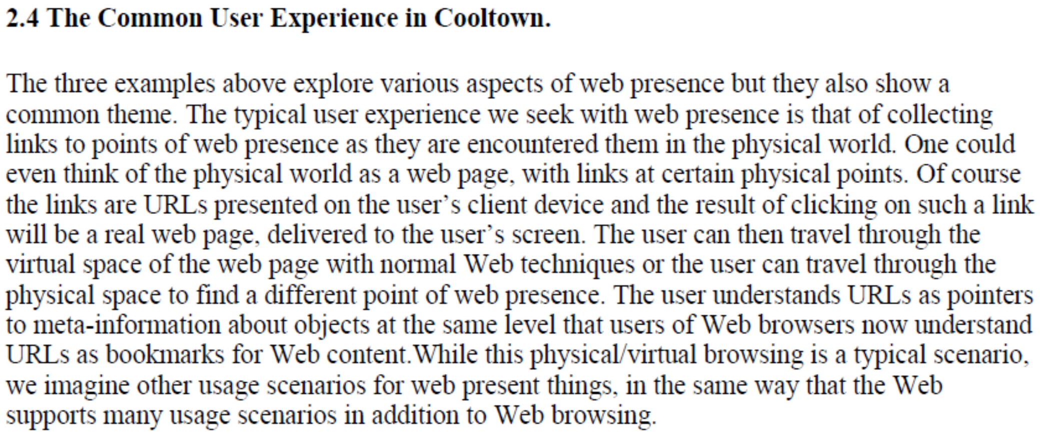 **2.4 The Common User Experience in Cooltown.** The three examples above explore various aspects of web presence but they also show a common theme. The typical user experience we seek with web presence is that of collecting links to points of web presence as they are encountered them in the physical world. One could even think of the physical world as a web page, with links at certain physical points. Of course the links are URLs presented on the user's client device and the result of clicking on such a link will be a real web page, delivered to the user's screen. The user can then travel through the virtual space of the web page with normal Web techniques or the user can travel through the physical space to find a different point of web presence. The user understands URLs as pointers to meta-information about objects at the same level that users of Web browsers now understand URLs as bookmarks for Web content.While this physical/virtual browsing is a typical scenario, we imagine other usage scenarios for web present things, in the same way that the Web supports many usage scenarios in addition to Web browsing. Cooltown Developer's Network, People, Place, and Things: web presence for the real world (HP000025–34) (HP000025–34) at 23

Element	'036 Claim Recitation	Exemplary Citations in Reference
		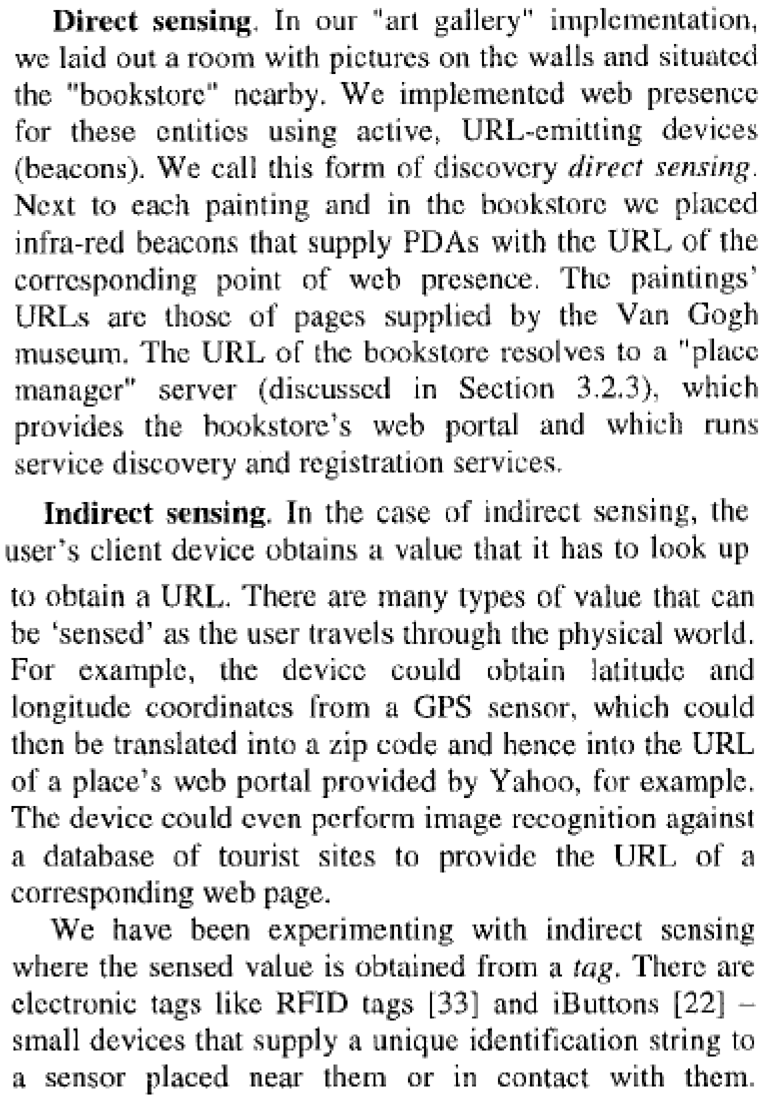 **Direct sensing**. In our "art gallery" implementation, we laid out a room with pictures on the walls and situated the "bookstore" nearby. We implemented web presence for these entities using active, URL-emitting devices (beacons). We call this form of discovery *direct sensing*. Next to each painting and in the bookstore we placed infra-red beacons that supply PDAs with the URL of the corresponding point of web presence. The paintings' URLs are those of pages supplied by the Van Gogh museum. The URL of the bookstore resolves to a "place manager" server (discussed in Section 3.2.3), which provides the bookstore's web portal and which runs service discovery and registration services. **Indirect sensing**. In the case of indirect sensing, the user's client device obtains a value that it has to look up to obtain a URL. There are many types of value that can be 'sensed' as the user travels through the physical world. For example, the device could obtain latitude and longitude coordinates from a GPS sensor, which could then be translated into a zip code and hence into the URL of a place's web portal provided by Yahoo, for example. The device could even perform image recognition against a database of tourist sites to provide the URL of a corresponding web page. We have been experimenting with indirect sensing where the sensed value is obtained from a *tag*. There are electronic tags like RFID tags [33] and iButtons [22] – small devices that supply a unique identification string to a sensor placed near them or in contact with them.

Element	'036 Claim Recitation	Exemplary Citations in Reference
		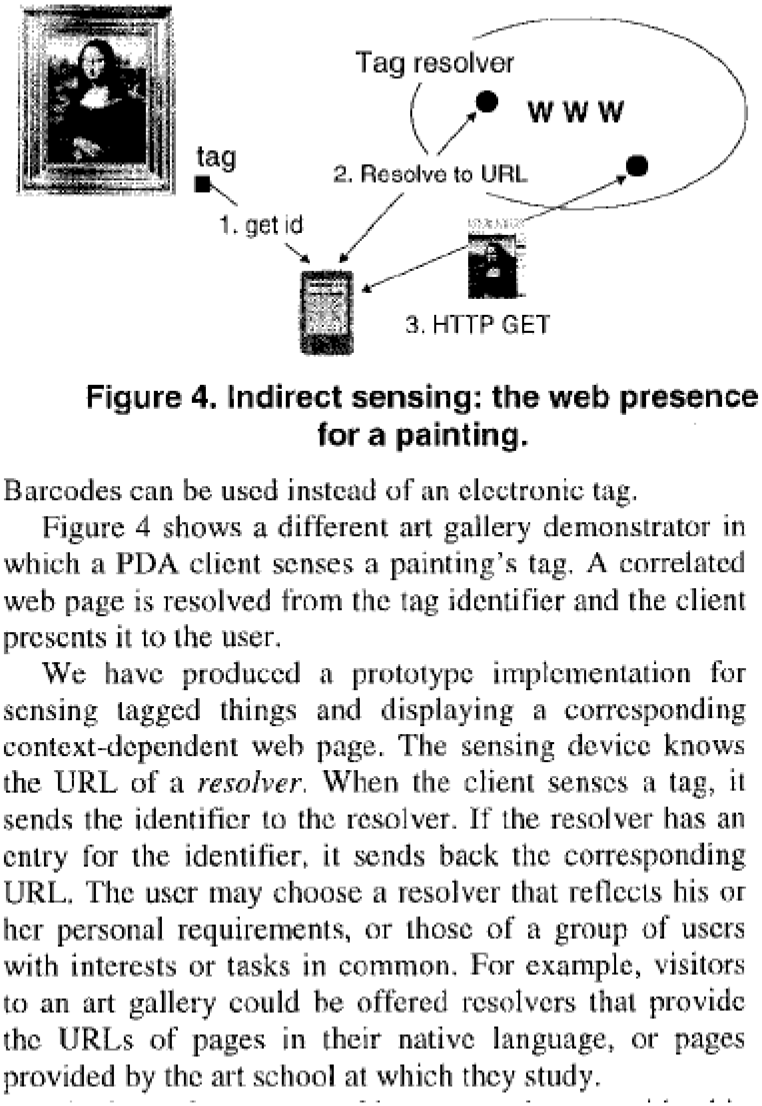 **Figure 4. Indirect sensing: the web presence for a painting.** Barcodes can be used instead of an electronic tag. Figure 4 shows a different art gallery demonstrator in which a PDA client senses a painting's tag. A correlated web page is resolved from the tag identifier and the client presents it to the user. We have produced a prototype implementation for sensing tagged things and displaying a corresponding context-dependent web page. The sensing device knows the URL of a *resolver*. When the client senses a tag, it sends the identifier to the resolver. If the resolver has an entry for the identifier, it sends back the corresponding URL. The user may choose a resolver that reflects his or her personal requirements, or those of a group of users with interests or tasks in common. For example, visitors to an art gallery could be offered resolvers that provide the URLs of pages in their native language, or pages provided by the art school at which they study. Cooltown Developer's Network, People, Place, and Things: web presence for the real world (HP000025–34) at 21

Element	’036 Claim Recitation	Exemplary Citations in Reference
		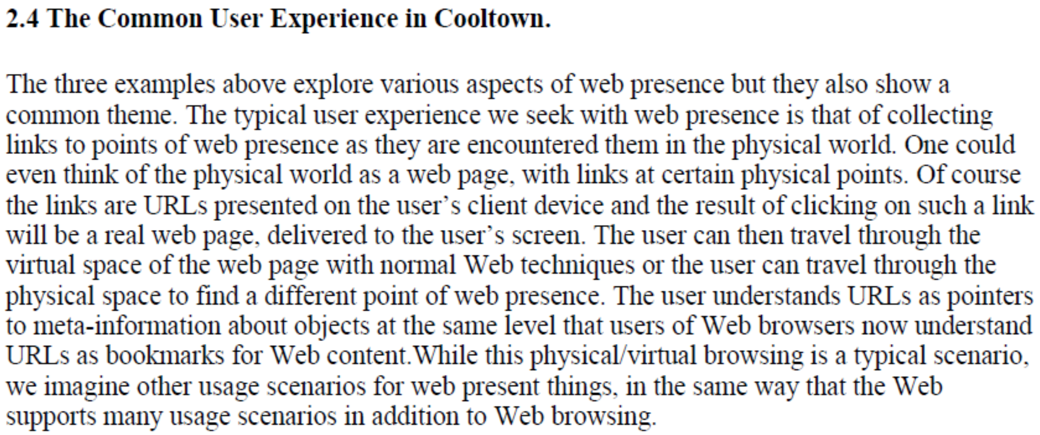 **2.4 The Common User Experience in Cooltown.** The three examples above explore various aspects of web presence but they also show a common theme. The typical user experience we seek with web presence is that of collecting links to points of web presence as they are encountered them in the physical world. One could even think of the physical world as a web page, with links at certain physical points. Of course the links are URLs presented on the user’s client device and the result of clicking on such a link will be a real web page, delivered to the user’s screen. The user can then travel through the virtual space of the web page with normal Web techniques or the user can travel through the physical space to find a different point of web presence. The user understands URLs as pointers to meta-information about objects at the same level that users of Web browsers now understand URLs as bookmarks for Web content.While this physical/virtual browsing is a typical scenario, we imagine other usage scenarios for web present things, in the same way that the Web supports many usage scenarios in addition to Web browsing. Cooltown Developer’s Network, People, Place, and Things: web presence for the real world (HP000025–34) at 9

<table>
<tr><th>Element</th><th>'036 Claim Recitation</th><th>Exemplary Citations in Reference</th></tr>
<tr><td></td><td></td><td>
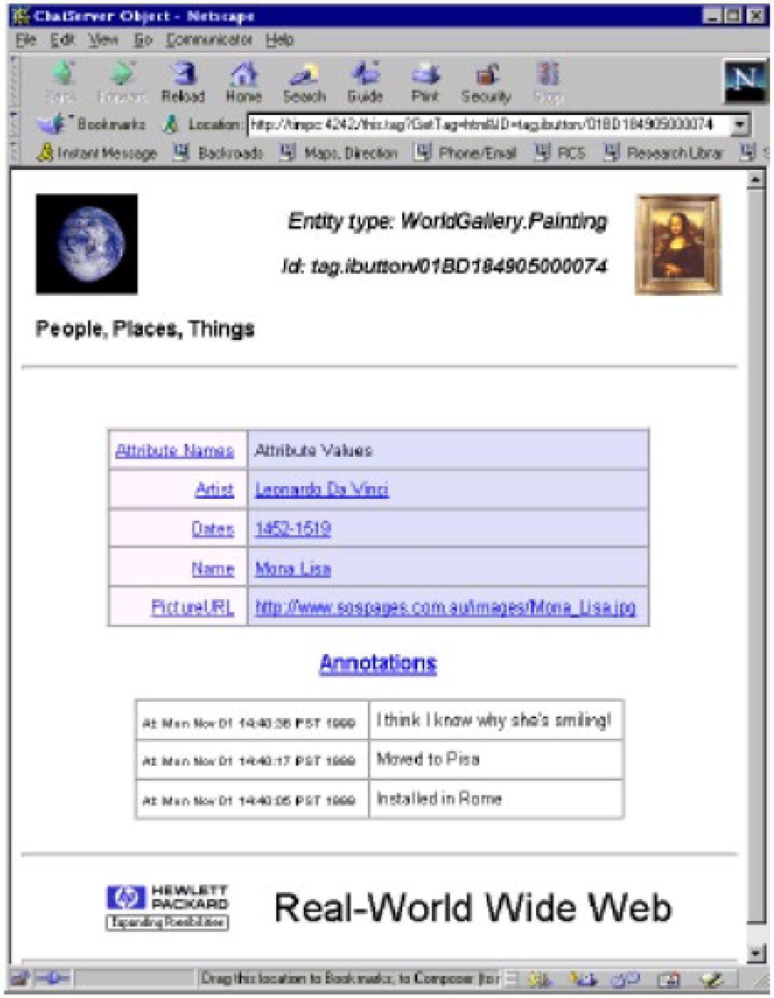

Location: http://hmpc:4242/this.tag?GetTag=htmlID=tag.ibutton/01BD184905000074

Entity type: WorldGallery.Painting

Id: tag.ibutton/01BD184905000074

People, Places, Things

<table>
<tr><td>Attribute Names</td><td>Attribute Values</td></tr>
<tr><td>Artist</td><td>Leonardo Da Vinci</td></tr>
<tr><td>Dates</td><td>1452-1519</td></tr>
<tr><td>Name</td><td>Mona Lisa</td></tr>
<tr><td>PictureURL</td><td>http://www.sospages.com.au/images/Mona_Lisa.jpg</td></tr>
</table>

Annotations

<table>
<tr><td>At Mon Nov 01 14:40:38 PST 1999</td><td>I think I know why she's smiling!</td></tr>
<tr><td>At Mon Nov 01 14:40:17 PST 1999</td><td>Moved to Pisa</td></tr>
<tr><td>At Mon Nov 01 14:40:05 PST 1999</td><td>Installed in Rome</td></tr>
</table>

Real-World Wide Web

Figure 5. A Web presence for *Mona Lisa*.
</td></tr>
</table>

Element	'036 Claim Recitation	Exemplary Citations in Reference
		To the extent it is found that CoolTown does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that CoolTown does not anticipate this claim, CoolTown renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine CoolTown with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

Exhibit E-3

Bank of America Corporation ("BofA" or "Defendant") provides this chart subject to all reservations, objections, statements, and disclaimers set forth herein and in BofA's Final Invalidity Contentions Cover Pleading, as well as any amendment, supplement, or modification thereof, which are incorporated herein by reference in their entirety.

On information and belief, the application that matured into Rhoads was filed on May 15, 2000, claiming the priority date of May 19, 1999. Rhoads published at least as of September 20, 2005, and is available as prior art at least under 35 U.S.C. § 102(e).

As shown in the chart below, Rhoads anticipates asserted claims 18, 27, 29, and 31 of the '252 Patent. To the extent Rhoads is found not to expressly disclose certain limitations in the asserted claims, such limitations are inherent. To the extent Rhoads is found not to anticipate any asserted claims or claim elements of the '252 Patent, the reference nevertheless renders those claims or claim elements obvious under 35 U.S.C. § 103, either alone or in combination with other art identified herein or in the cover pleading of BofA's Final Invalidity Contentions, for the reasons identified in the cover pleading. These Final Invalidity Contentions are not an admission by BofA that the accused products, including any current or past versions of the accused products, are covered by, or infringe the asserted claims, but are based instead on the recognition that if the claims are interpreted to be broad enough to encompass the accused products, the claims must also be construed to have that same scope when considering whether they are invalid.

The following chart is based on the Court's claim construction rulings (Dkt. Nos. 145, 153, 236), and to the extent not contradictory, Plaintiffs' apparent interpretations of the claims in Plaintiff's "Preliminary" Final Infringement Contentions. BofA notes that in many instances, Plaintiff's "Preliminary" Final Infringement Contentions ignore and/or contradict the Court's claim construction rulings. BofA does not accept the assumptions concerning the scope and meaning implicit in Plaintiff's Preliminary Final Infringement Contentions, to the extent those assumptions are discernible, and reserves the right to challenge Plaintiff's proposed (or implied) constructions. Furthermore, Plaintiffs' "Preliminary" Final Contentions fail to comply with the level of specificity required by SPR 2.1.2 and 2.1.3, and include improper reservation of rights to supplement with "additional ways in which BoA Accused Products infringe." Any alleged deficient discovery is of Plaintiffs own making. Most if not all of the deficiencies in Plaintiffs' Preliminary Final Contentions have been documented by Defendants through multiple letters since at least May 18, 2021 and up to Bank of America's most recent letter of August 19, 2023. Plaintiffs' shifting infringement contentions, deficiencies in its infringement contentions, and changing claim scope have prejudiced BofA from being able to reasonably prepare its defenses, including its invalidity contentions. BofA also reserves the right to revise and supplement these charts if and when Plaintiff is permitted to provide revised or supplemental Infringement Contentions. Citations given in the chart below are merely representative of the respective elements to provide notice to Plaintiffs' of BofA's invalidity theories and are not meant to be exhaustive.

1. Claim 18

Element	'252 Claim Recitation	Exemplary Citations in Reference
[1Pre]	A method for retrieving information from image processing, the method comprising:	Rhoads at Abstract • "A cell phone is equipped with a 2D optical sensor, enabling a variety of applications. For example, such a phone may also be provided with a digital watermark decoder, permitting decoding of steganographic data on imaged objects. Movement of a phone may be inferred by sensing movement of an imaged pattern across the optical sensor's field of view, allowing use of the phone as a gestural input device through which a user can signal instructions to a computer-based process. A variety of other arrangements by which electronic devices can interact with the physical world are also detailed, e.g., involving sensing and responding to digital watermarks, bar codes, RFIDs, etc." Rhoads at 5:54-56. • "Basically, the technology detailed in this disclosure may be regarded as enhanced systems by which users can interact with computer-based devices." Rhoads at 9:35–41 • "The Steganographic decoding process may entail three Steps. In the first, the object is located. In the Second, the objects orientation is discerned. In the third, the Bedoop data is extracted from the image data corresponding to the Bedoop object." Rhoads at Fig. 2 (and accompanying description)

Element	'252 Claim Recitation	Exemplary Citations in Reference
		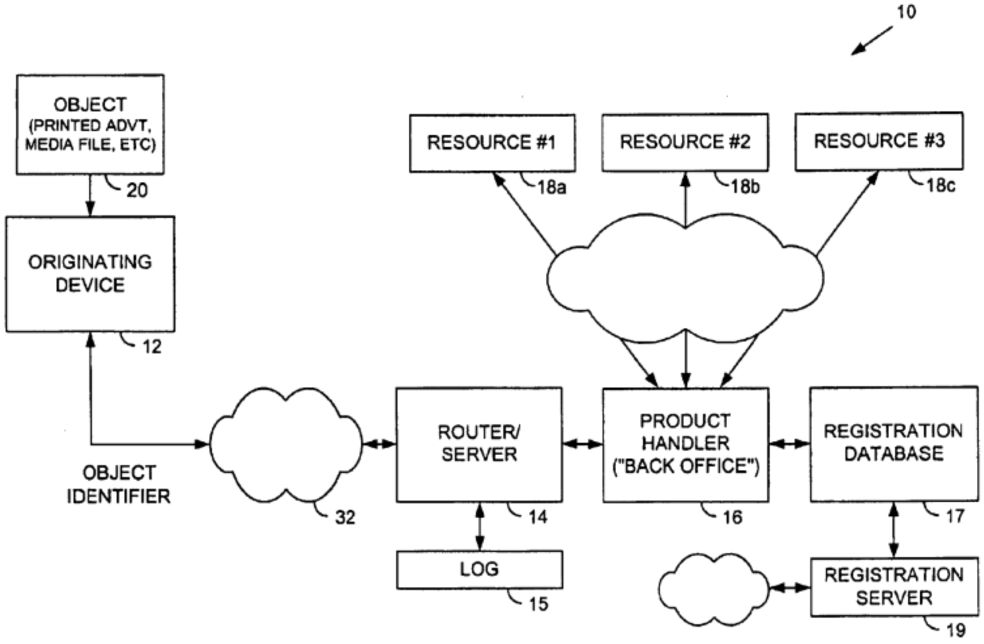 **FIG. 2** Rhoads at 49:15-18 • "Referring to FIG. 2, a system 10 according to the exemplary embodiment includes an originating device 12, a router/server 14, a product handler 16, a registration database 17, and one or more remote resources 18." Rhoads at 50:21-26 • "As more particularly detailed below, the handler 16 provides a response in accordance with the particular watermark payload. The response may be provided directly by the product handler to the device 12, or the handler may respond by communicating with a remote resource 18 (which may be, e.g., a data repository or service provider)."

Element	'252 Claim Recitation	Exemplary Citations in Reference
[1A]	providing a mobile device having a camera the mobile device configured to capture an image and configured to transmit data relating to the image to an image processing platform;	Rhoads discloses providing a mobile device having a camera the mobile device configured to capture an image and configured to transmit data relating to the image to an image processing platform. Rhoads at Fig. 2 (and accompanying description) 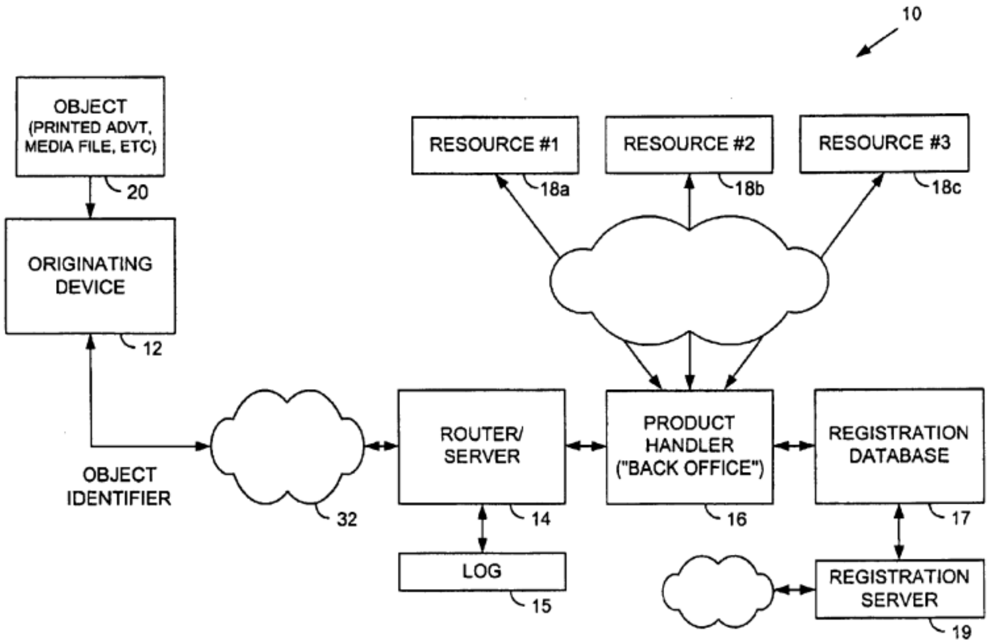 **FIG. 2** Rhoads at 3:41–51 • "According to another aspect, the invention includes a promotional method comprising: (a) presenting an object within the field of view of an optical sensor device, the object being selected from the list consisting of a retail product, packaging for a retail product, or printed advertising; (b) acquiring optical data corresponding to the object; (c) decoding plural-bit digital data from the optical data; (d) submitting at least some of said decoded data to a

Element	'252 Claim Recitation	Exemplary Citations in Reference
		remote computer; and (e) determining at the remote computer whether a prize should be awarded in response to submission of said decoded data." Rhoads at 7:24-42 • "Device 12 interacts with an object 20. The object can be electronic or not. Electronic objects can include computer files, representations of audio, video, or still imagery (e.g., files or in streaming form), etc. Non-electronic objects can include physical objects such as newspapers, magazine pages, posters, product packaging, event tickets, credit cards, paper currency, etc. Non-electronic objects can also include sounds produced by loudspeakers. When used with nonelectronic objects, device 12 (FIG. 2) typically includes some form of sensor or transducer 22 to produce electronic signals or data corresponding to the object. Examples include CCD- or CMOS-based optical sensors (either as part of still- or video cameras, flatbed scanners, mice, or otherwise), microphones, barcode scanners, RF ID sensors, mag stripe readers, etc. In such cases, the sensor 22 may be coupled to associated interface electronics 24, which in turn may be coupled to device driver software 26, which in turn may be coupled to one or more application programs 28." Rhoads at Fig. 11

Element	'252 Claim Recitation	Exemplary Citations in Reference
		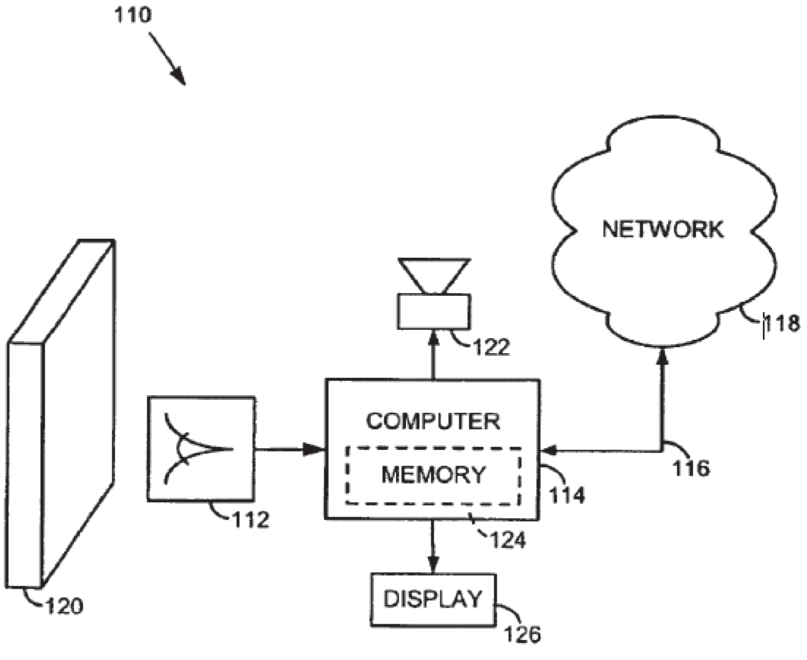 FIG. 11 Rhoads at 7:13-29 • "Referring to FIG. 11, a basic embodiment 110 of the present technology includes an optical sensor 112, a computer 114, and a network connection 116 to the internet 118. The illustrated optical sensor 112 is a digital camera having a resolution of 320 by 200 pixels (color or black and white) that stares out, grabbing frames of image data five times per second and storing same in one or more frame buffers. These frames of image data are analyzed

Element	'252 Claim Recitation	Exemplary Citations in Reference
		by a computer 114 for the presence of Bedoop data. (Essentially, Bedoop data is any form of plural-bit data encoding recognized by the system 110—data which, in many embodiments, initiates some action.) Once detected, the system responds in accordance with the detected Bedoop data (e.g., by initiating some local action, or by communication with a remote computer, such as over the internet, via an online service such as AOL, or using point-to-point dial-up communications, as with a bulletin board system)." Rhoads 8:66 – 9:12 • "In still other embodiments, the camera (or other sensor) can be equipped with one or more auxiliary fixed-focus lenses that can be selectively used, depending on the particular application. Some such embodiments have a first fixed focused lens that always overlies the sensor, with which one or more auxiliary lenses can be optically cascaded (e.g., by hinge or slide arrangements). Such arrangements are desirable, e.g., when a camera is not a dedicated Bedoop sensor but also performs other imaging tasks. When the camera is to be used for Bedoop, the auxiliary lens is positioned (e.g., flipped into) place, changing the focal length of the first lens (which may by unsuitably long for Bedoop purposes, such as infinity) to an appropriate Bedoop imaging range (such as one foot)." Rhoads at 17:46–54 • "The building can be provided with an image sensor (such as a video camera or the like), an associated Bedoop. detection system, and the proximity detector. When a user wearing the badge approaches, the proximity detector signals the camera to capture image data. The Bedoop detection system identifies the badge photograph (e.g., by clues as are described in the prior applications, or without such aids), captures optical data, and decodes same to extract the steganographically-embedded data hidden therein."

Element	'252 Claim Recitation	Exemplary Citations in Reference
		Rhoads at 46:47-53 • "The provision of image Sensors in cell phones enables other functionality. One is the capture of still or video imagery. Such image data can be compressed (typically by lossy processes such as MPEG, JPEG, or the like, implemented with dedicated hardware CODECs) and transmitted with the audio data. The Screens on Such phones can likewise be used for display of incoming image or Video data." To the extent it is found that Rhoads does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Rhoads does not anticipate this claim, Rhoads renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Rhoads with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1B]	configuring the image processing platform to receive the data relating to the image and to conduct image processing, including:	Rhoads discloses configuring the image processing platform to receive the data relating to the image and to conduct image processing. Rhoads at Fig. 2 (including accompanying description)

Element	’252 Claim Recitation	Exemplary Citations in Reference
		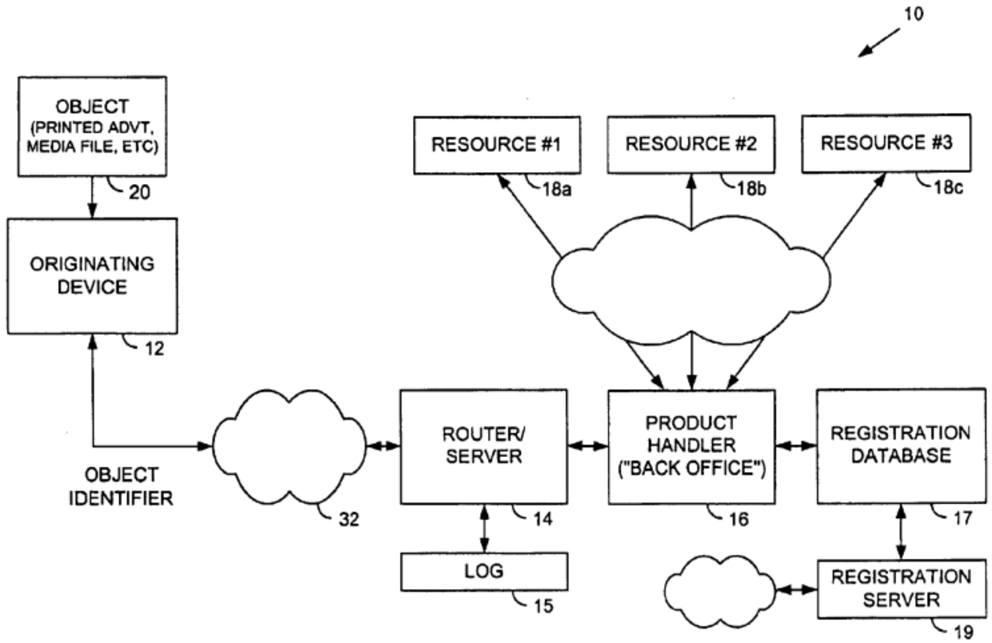 FIG. 2 Rhoads at 3:52–60 • “According to another aspect, the invention includes a method of interacting with a magazine using a computer, the computer including an internet web browser, the method including: (a) providing a peripheral device having a sensor; (b) positioning the peripheral device adjacent a first advertisement in the magazine to direct the web browser to a first internet address; and (c) positioning the peripheral device adjacent a second advertisement in the magazine to direct the web browser to a second internet address.” Rhoads at 49:15–35

Element	'252 Claim Recitation	Exemplary Citations in Reference
		- "Referring to FIG. 2, a system 10 according to the exemplary embodiment includes an originating device 12, a router/server 14, a product handler 16, a registration database 17, and one or more remote resources 18. The originating device 12 can take many different forms, e.g., a cell phone, a personal digital assistant (e.g., a Palm Pilot), a personal computer, a barcode scanning system, etc. For expository convenience, the embodiment is described with reference to a personal computer for device 12. Device 12 interacts with an object 20. The object can be electronic or not. Electronic objects can include computer files, representations of audio, video, or still imagery (e.g., files or in streaming form), etc. Non-electronic objects can include physical objects such as newspapers, magazine pages, posters, product packaging, event tickets, credit cards, paper currency, etc. Non-electronic objects can also include sounds produced by loudspeakers. When used with nonelectronic objects, device 12 (FIG. 2) typically includes some form of sensor or transducer 22 to produce electronic signals or data corresponding to the object…." Rhoads at 10:22–30 "Some embodiments filter the image data at some point in the process to aid in ultimate Bedoop data extraction. One use of such filtering is to mitigate image data artifacts due to the particular optical sensor. For example, CCD arrays have regularly-spaced sensors that sample the optical image at uniformly spaced discrete points. This discrete sampling effects a transformation of the image data, leading to certain image artifacts. An appropriately configured filter can reduce some of these artifacts." Rhoads at 19:61 – 20:15 - "The coffee cup is an example of a non-planar object. Another is soft drink

Element	'252 Claim Recitation	Exemplary Citations in Reference
		cans. Special issues can arise when encoding and decoding markings on such objects. For example, when sensing such an object with a camera or the like, part of the image will be out of focus due to differing distances from the camera to different parts of the can surface. While parts of an image sensed from a non-planar object, such as a can, may be out of focus, they still convey useful image data. The out of focus areas are just blurred—as if filtered by a low pass filter. But to make use of this information, a further complication must first be addressed: warping. When viewed from a camera, the planar artwork with which the can is wrapped becomes warped. Portions of the can nearest the camera appear at a nominal full scale, while areas successively further around the can curvature (as viewed from the camera) appear progressively more and more spatially compressed. Regardless of the watermarking technology being employed, the physical warp of the can's surface is likewise manifested as a warping of the encoded watermark data." Rhoads at 38:1 – 19 • "Post-It® Notes Pads of Post-It® notes, or other pads of paper, can be marked by the manufacturer (either by texturing, watermarked tinting, ink-jet spattering, etc.) to convey steganographic data (e.g., Bedoop data). When such a note is presented to a Bedoop system, the system may launch an application that stores a snapshot of the note. More particularly, the application may mask the note-portion of the image data from the other image data, virtually re-map it to a rectangular format of standardized pixel dimensions, JPEG-compress the resulting image, and store it in a particular computer subdirectory with a name indicating the date of image acquisition, together with the color and/or size of the note. (These latter two data may be

Element	'252 Claim Recitation	Exemplary Citations in Reference
		indicated by data included in the Bedoop payload.) If the color of the note is indicated by digital data (e.g, in the file name), then the image itself may be stored in grey-scale. When later recalled for display, the white image background can be flooded with color in accordance with the digital color data." Rhoads at 42:1–12 • "The optical data collected by the sensor can be processed within the peripheral's processor to extract the steganographically encoded binary Bedoop data therefrom. Or this processing burden can be undertaken by the associated computer system, with the peripheral simply processing and formatting the raw sensor data into sequential frames of image data to be output to that system. While scanning peripherals of the type described above are typically wired to an associated host system, wireless links (e.g., radio, infrared, ultrasonic, etc.) can of course be used, freeing the user from the constraint imposed by the cable." Rhoads at 45:19–55 • "Bedoop technology can be integrated into portable telecommunication terminals, such as cell phones (manufactured, e.g., by Motorola, Nokia, Qualcomm, and others). Such a phone can be equipped with a 1D or 2D image sensor, the output of which is applied to Bedoop decoding circuitry within the phone. This decoding circuitry can be the phone's main CPU, or can be a processing circuit dedicated to Bedoop functionality. (In this as in other embodiments, the decoding can be effected by dedicated hardware, by decoding software executing on a general purpose CPU, etc.) Cell phones are already equipped with numerous features that make them

Element	'252 Claim Recitation	Exemplary Citations in Reference
		well suited for Bedoop operation. One is that cell phones typically include an LCD or similar screen for display of textual or graphic information, and additionally include buttons or other controls for selecting among menu options presented on the screen (e.g., by moving a cursor). Moreover, cell phones naturally include both audio input and output devices (i.e., microphone and speaker). Still further, the protocol by which cell phones transmit data includes data identifying the phone, so that such data need not be separately encoded. And finally, cell phones obviously provide ready links to remote computer systems. Collectively, these capabilities rival those of the most fully-equipped desktop computer system. Thus, essentially all of the applications detailed elsewhere in this specification can be implemented using cell phone Bedoop systems As with the other Bedoop systems, when Bedoop data is sensed, the phone can respond to the data locally, or it can forward Same over the cellular network to a remote System (or computer network) for handling." Rhoads at 46:10–12. • "Bedoop decoding function within the phone, in other embodiments the image data can be transmitted from the phone and decoded at a remote location." Rhoads at 46:47 – 53 • "The provision of image sensors in cell phones enables other functionality. One is the capture of still or video imagery. Such image data can be compressed (typically by lossy processes such as MPEG, JPEG, or the like, implemented with dedicated hardware CODECs) and transmitted with the audio data. The screens on such phones can likewise be used for display of incoming image or video data."

Element	'252 Claim Recitation	Exemplary Citations in Reference
		Rhoads at Fig. 3 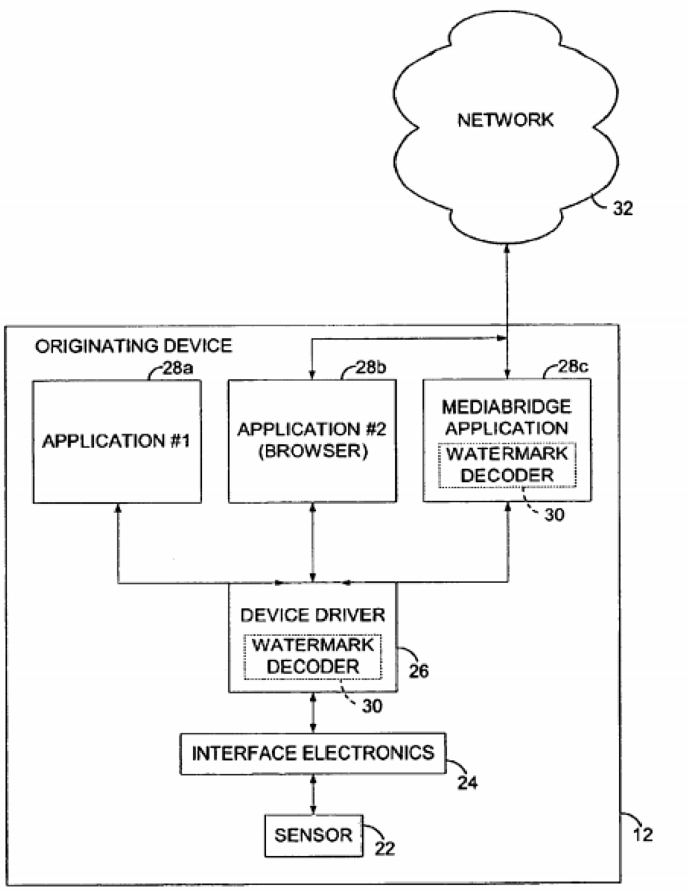 FIG. 3 Rhoads at 45:19–55

Element	'252 Claim Recitation	Exemplary Citations in Reference
		- "Bedoop technology can be integrated into portable telecommunication terminals, such as cell phones (manufactured, e.g., by Motorola, Nokia, Qualcomm, and others). Such a phone can be equipped with a 1D or 2D image sensor, the output of which is applied to Bedoop decoding circuitry within the phone. This decoding circuitry can be the phone's main CPU, or can be a processing circuit dedicated to Bedoop functionality. (In this as in other embodiments, the decoding can be effected by dedicated hardware, by decoding software executing on a general purpose CPU, etc.) Cell phones are already equipped with numerous features that make them well suited for Bedoop operation. One is that cell phones typically include an LCD or similar screen for display of textual or graphic information, and additionally include buttons or other controls for selecting among menu options presented on the screen (e.g., by moving a cursor). Moreover, cell phones naturally include both audio input and output devices (i.e., microphone and speaker). Still further, the protocol by which cell phones transmit data includes data identifying the phone, so that such data need not be separately encoded. And finally, cell phones obviously provide ready links to remote computer systems. Collectively, these capabilities rival those of the most fully-equipped desktop computer system. Thus, essentially all of the applications detailed elsewhere in this specification can be implemented using cell phone Bedoop systems As with the other Bedoop systems, when Bedoop data is sensed, the phone can respond to the data locally, or it can forward Same over the cellular network to a remote System (or computer network) for handling." Rhoads at 46:10–12. - "Bedoop decoding function within the phone, in other embodiments the image data can be transmitted from the phone and decoded at a remote

Element	'252 Claim Recitation	Exemplary Citations in Reference
		location." Rhoads at 46:47 – 53 "The provision of image sensors in cell phones enables other functionality. One is the capture of still or video imagery. Such image data can be compressed (typically by lossy processes such as MPEG, JPEG, or the like, implemented with dedicated hardware CODECs) and transmitted with the audio data. The screens on such phones can likewise be used for display of incoming image or video data." Rhoads at 49:48-65 • "The detailed embodiment contemplates that the object 20 is a magazine advertisement encoded with a steganographic watermark conveying a plural bit object identifier. The watermark is hidden in the advertisement's image in a manner indiscernable to human observers, but detectable by computer analysis. That analysis is performed by a watermark detector 30. Watermark detector 30 can be implemented in various different locations in the system of FIG. 1. Typically, the detector is implemented in the originating device 12, e.g., in the driver software 26, or in application software 28 c that serves to link to external resources based on detected watermarks. But it may be implemented elsewhere, e.g., in hardware in the interface electronics 24, in an operating system associated with the device, or outside device 12 altogether. Some systems may have plural watermark detectors, implemented at different locations throughout the system." Rhoads at Fig. 12 (and accompanying description)

Element	'252 Claim Recitation	Exemplary Citations in Reference
		FIG. 12 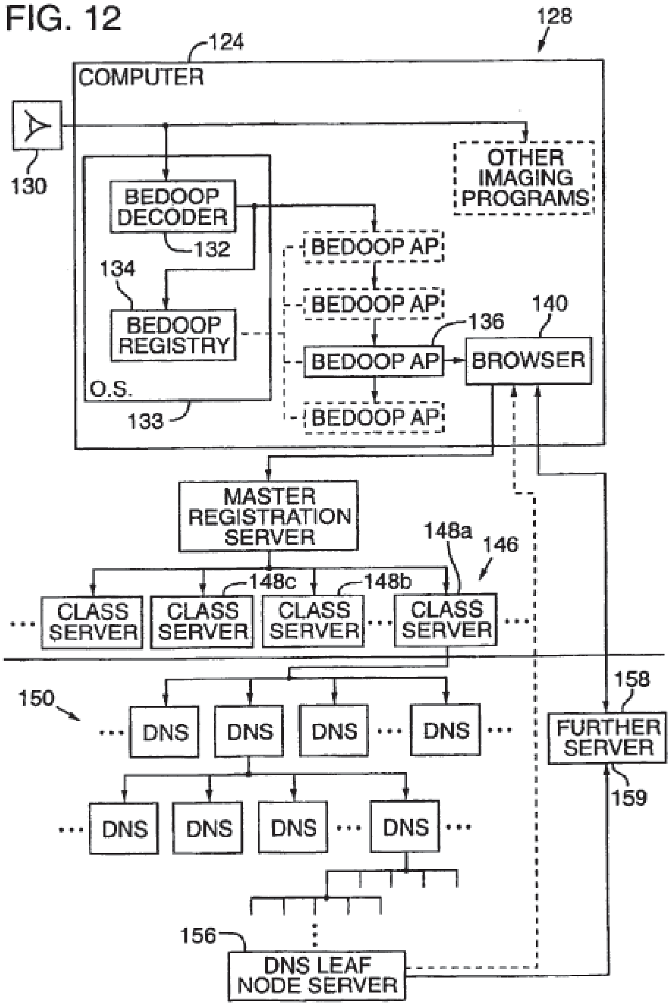 Rhoads at 12:23 – 33 • "Sometimes the computer system may encounter a Bedoop object for which it does not have a registered application program. In such case, a default

Element	'252 Claim Recitation	Exemplary Citations in Reference
		Bedoop application can be invoked. This default application can, e.g., establish an internet link to a remote server computer (or a network of such computers), and can transmit the Bedoop data (or a part of the Bedoop data) to that remote computer. The remote server can undertake the response itself, it can instruct the originating computer how to respond appropriately, or it can undertake some combination of these two responses. (Such arrangements are further considered below.)" Rhoads at 16:66 – 17:10 • "Inside the building the candidate may encounter a security guard. Seeing an unfamiliar person, the guard may visually compare the photo on the badge with the candidate's face. Additionally, the guard can present the badge to a portable Bedoop device, or to one of many Bedoop systems scattered through the building (e g., at every telephone). The Bedoop system extracts the Bedoop data from the card (i.e., from the DMV photograph), interrogates the DMV's DNS server with this Bedoop data, and receives in reply the name of the person depicted in the photograph. (If the Bedoop system is a telephone, the name may be displayed on a small LCD display commonly provided on telephones.)" To the extent it is found that Rhoads does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Rhoads does not anticipate this claim, Rhoads renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Rhoads with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar

Element	'252 Claim Recitation	Exemplary Citations in Reference
		known problems, according to known methods which would yield predictable results.
[1B.i]	operating on the data relating to the image to determine if the image contains one or more recognizable symbols; and	Rhoads discloses operating on the data relating to the image to determine if the image contains one or more recognizable symbols. Rhoads at Abstract • "A cell phone is equipped with a 2D optical sensor, enabling a variety of applications. For example, such a phone may also be provided with a digital watermark decoder, permitting decoding of steganographic data on imaged objects. Movement of a phone may be inferred by sensing movement of an imaged pattern across the optical sensor's field of view, allowing use of the phone as a gestural input device through which a user can signal instructions to a computer-based process. A variety of other arrangements by which electronic devices can interact with the physical world are also detailed, e.g., involving sensing and responding to digital watermarks, bar codes, RFIDs, etc." Rhoads at 4:19–26 • "According to another aspect, the invention comprises a wireless telephony handset including a microphone, a modulator, and an RF amplifier, the device serving to receive audio and transmit an RF signal conveying audio modulation, the handset further including an optical sensor producing optical data, a lens for imaging an object onto the sensor, and a decoder for decoding plural bit identifier data conveyed by a barcode or a digital watermark on the object." Rhoads at 7:2–11 • "Although the following specification focuses on applications employing digital watermarking, certain of such applications can alternatively employ other data encoding techniques, including 1D and 2D barcodes, magnetic ink

Element	'252 Claim Recitation	Exemplary Citations in Reference
		character recognition (MICR), optical character recognition (OCR), optical mark recognition (OMR), radio frequency identification (RF/ID), UV/IR identification technologies, data glyphs, organic transistors, magnetic stripe, etc., depending on the particular application requirements." Rhoads at 9:35–41 • "The Steganographic decoding process may entail three Steps. In the first, the object is located. In the Second, the objects orientation is discerned. In the third, the Bedoop data is extracted from the image data corresponding to the Bedoop object." Rhoads at 9:35–10:10 • "The Steganographic decoding process may entail three Steps. In the first, the object is located. In the Second, the objects orientation is discerned. In the third, the Bedoop data is extracted from the image data corresponding to the Bedoop object. The first step, object location, can be assisted by various clues. One is the placement of the object; typically the center of the image field will be a point on the object. The surrounding data can then be analyzed to try and discern the object's boundaries. Another location technique is slight movement. Although the user will typically try to hold the object still, there will usually be some jitter of the Bedoop object within the image frame (e.g., a few pixels back and forth). Background visual clutter, in contrast, will typically be stationary. Such movement may thus be sensed and used to identify the Bedoop object from within the image data. Still another object-location clue is object shape. Many Bedoop objects are rectangular in shape (or trapezoidal as viewed by the camera). Straight edge

Element	'252 Claim Recitation	Exemplary Citations in Reference
		boundaries can thus be used to define an area of likely Bedoop data. Color is a further object identification clue that may be useful in some contexts. Yet another object location clue is spatial frequency. In imaging systems with well defined focal zones, undesired visual clutter may be at focal distances that results in blurring. The Bedoop object, in contrast, will be in focus and may be characterized by fine detail. Analyzing the image data for the high frequencies associated with fine detail can be used to distinguish the intended object from others. Characteristic markings on the object (as discussed below in connection with determining object orientation) can also be sensed and used in locating the object. Once the Bedoop object has been located within the image data, masking can be applied (if desired) to eliminate image data not corresponding to the intended object." Rhoads at 44:66 – 45:11 • "Just as the scanner 210 can perform watermark decoding, it can similarly perform barcode decoding. Conventional pattern-recognition algorithms can be applied by CPU 212 to either the raw or final scan data to identify barcode patterns. Once the pattern is identified, decoding is straightforward by applying known barcode alphabets. As in the watermark case, barcode decoding can be performed autonomously, or in response to user/auxiliary device command. Similarly, the decoded barcode data can be provided to the auxiliary device whenever detected, or in response to an auxiliary device query. The 2D data can likewise be used to augment the information provided to the CPU so as to aid in the barcode detection/decoding process."

Element	'252 Claim Recitation	Exemplary Citations in Reference
		Rhoads at 49:3-14: - "Detection is the complementary operation to embedding, i.e., discerning a digital identifier from an object. Response refers to the action taken based on the discerned identifier. The middle two steps—embedding and detection—can employ any of myriad well-known technologies, including 1D and 2D barcodes, magnetic ink character recognition (MICR), optical character recognition (OCR), optical mark recognition (OMR), radio frequency identification (RF/ID), data glyphs, organic transistors, magnetic stripe, metadata, file header information, UV/IR identifiers, and other machine-readable indicia and techniques for associating plural-bit digital data with an electronic or physical object. The detailed embodiment employs watermarking technology, although this is illustrative only." Rhoads at 49:48-65: - "The detailed embodiment contemplates that the object 20 is a magazine advertisement encoded with a steganographic watermark conveying a plural bit object identifier. The watermark is hidden in the advertisement's image in a manner indiscernable to human observers, but detectable by computer analysis. That analysis is performed by a watermark detector 30. Watermark detector 30 can be implemented in various different locations in the system of FIG. 1. Typically, the detector is implemented in the originating device 12, e.g., in the driver software 26, or in application software 28 c that serves to link to external resources based on detected watermarks. But it may be implemented elsewhere, e.g., in hardware in the interface electronics 24, in an operating system associated with the device, or outside device 12 altogether. Some systems may have plural watermark detectors, implemented at different locations throughout the system."

Element	'252 Claim Recitation	Exemplary Citations in Reference
		Rhoads at 91:3–6 • "As indicated above, while steganographic encoding of the digital data is used in the preferred embodiments, visible forms of digital encoding—such as bar codes—can naturally be employed where aesthetic considerations permit." To the extent it is found that Rhoads does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Rhoads does not anticipate this claim, Rhoads renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Rhoads with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1B.ii]	decoding the recognizable symbols to extract symbol information by analyzing the recognizable symbols according to type;	Rhoads discloses decoding the recognizable symbols to extract symbol information by analyzing the recognizable symbols according to type. Rhoads at 4:18 – 26 • "According to another aspect, the invention comprises a wireless telephony handset including a microphone, a modulator, and an RF amplifier, the device serving to receive audio and transmit an RF signal conveying audio modulation, the handset further including an optical sensor producing optical data, a lens for imaging an object onto the sensor, and a decoder for decoding plural bit identifier data conveyed by a barcode or a digital watermark on the object." Rhoads at 7:2–11

Element	'252 Claim Recitation	Exemplary Citations in Reference
		• "Although the following specification focuses on applications employing digital watermarking, certain of such applications can alternatively employ other data encoding techniques, including 1D and 2D barcodes, magnetic ink character recognition (MICR), optical character recognition (OCR), optical mark recognition (OMR), radio frequency identification (RF/ID), UV/IR identification technologies, data glyphs, organic transistors, magnetic stripe, etc., depending on the particular application requirements." Rhoads at 9:35–10:21 • "The Steganographic decoding process may entail three Steps. In the first, the object is located. In the Second, the objects orientation is discerned. In the third, the Bedoop data is extracted from the image data corresponding to the Bedoop object. The first step, object location, can be assisted by various clues. One is the placement of the object; typically the center of the image field will be a point on the object. The surrounding data can then be analyzed to try and discern the object's boundaries. Another location technique is slight movement. Although the user will typically try to hold the object still, there will usually be some jitter of the Bedoop object within the image frame (e.g., a few pixels back and forth). Background visual clutter, in contrast, will typically be stationary. Such movement may thus be sensed and used to identify the Bedoop object from within the image data. Still another object-location clue is object shape. Many Bedoop objects are rectangular in shape (or trapezoidal as viewed by the camera). Straight edge boundaries can thus be used to define an area of likely Bedoop data. Color is a further object identification clue that may be useful in some

Element	'252 Claim Recitation	Exemplary Citations in Reference
		contexts. Yet another object location clue is spatial frequency. In imaging systems with well defined focal zones, undesired visual clutter may be at focal distances that results in blurring. The Bedoop object, in contrast, will be in focus and may be characterized by fine detail. Analyzing the image data for the high frequencies associated with fine detail can be used to distinguish the intended object from others. Characteristic markings on the object (as discussed below in connection with determining object orientation) can also be sensed and used in locating the object. Once the Bedoop object has been located within the image data, masking can be applied (if desired) to eliminate image data not corresponding to the intended object. Once the Bedoop object has been located within the image data, masking can be applied (if desired) to eliminate image data not corresponding to the intended object. The next step in the decoding process, determining orientation of the Bedoop data, can likewise be discerned by reference to visual clues. For example, some objects include subliminal graticule data, or other calibration data, steganographically encoded with the Bedoop data to aid in determining orientation. Others can employ overt markings, either placed for that sole purpose (e.g. reference lines or fiducials), or serving another purpose as well (e.g. lines of text), to discern orientation. Edge-detection algorithms can also be employed to deduce the orientation of the object by reference to its edges."

Element	**'252 Claim Recitation**	**Exemplary Citations in Reference**

		Rhoads at 11:19–22
		- "Details of particular encoding and decoding techniques can be found in U.S. Pat. No. 5,862,260 and application Ser. No. 09/503,881, the disclosures of which are incorporated by reference."
		- The 260 patent discloses:
		  o "Step one of FIG. 25 is the standard "scanning" of a personal cash card 950 within the optical window 960. This can be performed using linear optical sensors which scan the window, or via a two dimensional optical detector array such as a CCD. The resulting scan is digitized into a grey scale image and stored in an image frame memory buffer such as a "framegrabber," as is now common in the designs of optical imaging systems. Once the card is scanned, a first image processing step would probably be locating the four fiducial center points, 954, and using these four points to guide all further image processing operations (i.e. the four fiducials "register" the corresponding patterns and barcodes on the personal cash card). Next, the barcode ID number would be extracted using common barcode reading image processing methods. Generally, the user's account number would be determined in this step." '260 patent at 58:45–60
		  o "It will be appreciated that the information embedded in the photograph need not be visually hidden or steganographically embedded. For example, the photograph incorporated into the identification card may be a composite of an image of the individual and one-, or two-dimensional bar codes. The bar code information would be subject to conventional optical scanning techniques (including internal cross checks) so that the information derived from the code may be compared, for example, to the information printed on the identification document." '260 patent at 61:62–62:4.
		  o "Generally speaking, three classes of data can be steganographically

Element	'252 Claim Recitation	Exemplary Citations in Reference
		embedded in an object: a number (e.g. a serial or identification number, encoded in binary), an alphanumeric message (e.g. a human readable name or telephone number, encoded in ASCII or a reduced bit code), or computer instructions (e.g. JAVA or HTML instructions). The embedded URLs and the like detailed above begin to explore this third class, but a more detailed exposition of the possibilities may be helpful." '260 patent at 64:59–67. Rhoads at 22:9–27 • "To aid in identification, the Bedoop objects (e.g., badges) can be colored a distinctive color, permitting the system to more easily identify candidate objects from other items within the optical capture devices' field of view. Or the object can be provided with a retro-reflective coating, and the elevator can be equipped with one or more illumination sources of known spectral or temporal quality (e.g., constant infra red, or constant illumination with a single- or multi-line spectrum, or a pulsed light source of known periodicity; LEDs or semiconductor lasers, each with an associated diffuser, can be used for each the foregoing and can be paired with the image capture devices). Other such tell-tale clues can likewise be used to aid in object location. In all such cases, the optical capture device can sense the tell-tale clue(s) using a wide field of view sensor. The device can then be physically or electronically steered, and/or zoomed, to acquire a higher resolution image of the digitally-encoded object suitable for decoding." Rhoads at 26:13 – 45 • "Gestural Decoding Module There are various ways in which the Bedoop system's decoding of gestural input can be effected. In some Bedoop systems, this functionality is provided as part of the Bedoop applications. Generally, however, the applications must be provided with the raw frame data in order to discern

Element	'252 Claim Recitation	Exemplary Citations in Reference
		the gestural movements. Since this functionality is typically utilized by many Bedoop applications, it is generally preferable to provide a single set of gestural interpretation software functions (commonly at the operating system level) to analyze the frame data, and make available gestural output data in standardized form to all Bedoop applications. In one such system, a gestural decoding module tracks the encoded object within the series of image data frames, and outputs various parameters characterizing the object's position and manipulation over time. Two of these parameters indicate the X-Y position of the object within current frame of image data. The module can identify a reference point (or several) on the object, and output two corresponding position data (X and Y). The first represents the horizontal offset of the reference point from the center of the image frame, represented as a percentage of frame width. A two's complement representation, or other representation capable of expressing both positive and negative values, can be used so that this parameter has a positive value if the reference point is right of center-frame, and has a negative value if the reference point is left of center frame. The second parameter, Y, similarly characterizes the position of the reference point above or below center-frame (with above-being represented by a positive value). Each of these two parameters can be expressed as a seven-bit byte. A new pair of X, Y parameters is output from the gestural decoding module each time a new frame of image data is processed." Rhoads at 44:66 – 45:11 • "Just as the scanner 210 can perform watermark decoding, it can similarly perform barcode decoding. Conventional pattern-recognition algorithms can be applied by CPU 212 to either the raw or final scan data to identify barcode patterns. Once the pattern is identified, decoding is straightforward by applying known barcode alphabets. As in the watermark case, barcode decoding can be performed autonomously, or in response to user/auxiliary

Element	'252 Claim Recitation	Exemplary Citations in Reference
		device command. Similarly, the decoded barcode data can be provided to the auxiliary device whenever detected, or in response to an auxiliary device query. The 2D data can likewise be used to augment the information provided to the CPU so as to aid in the barcode detection/decoding process." Rhoads at 49:3-14: • "Detection is the complementary operation to embedding, i.e., discerning a digital identifier from an object. Response refers to the action taken based on the discerned identifier. The middle two steps—embedding and detection—can employ any of myriad well-known technologies, including 1D and 2D barcodes, magnetic ink character recognition (MICR), optical character recognition (OCR), optical mark recognition (OMR), radio frequency identification (RF/ID), data glyphs, organic transistors, magnetic stripe, metadata, file header information, UV/IR identifiers, and other machine-readable indicia and techniques for associating plural-bit digital data with an electronic or physical object. The detailed embodiment employs watermarking technology, although this is illustrative only." Rhoads at 49:48-65 • "The detailed embodiment contemplates that the object 20 is a magazine advertisement encoded with a steganographic watermark conveying a plural bit object identifier. The watermark is hidden in the advertisement's image in a manner indiscernable to human observers, but detectable by computer analysis. That analysis is performed by a watermark detector 30. Watermark detector 30 can be implemented in various different locations in the system of FIG. 1. Typically, the detector is implemented in the originating device 12, e.g., in the driver software 26, or in application software 28 c that serves to link to external resources based on detected

Element	'252 Claim Recitation	Exemplary Citations in Reference
		watermarks. But it may be implemented elsewhere, e.g., in hardware in the interface electronics 24, in an operating system associated with the device, or outside device 12 altogether. Some systems may have plural watermark detectors, implemented at different locations throughout the system." 49:48-65. To the extent it is found that Rhoads does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Rhoads does not anticipate this claim, Rhoads renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Rhoads with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1C]	providing access to a distal server configured to use a database to identify pertinent information associated with the recognizable symbols based on the symbol information; and	Rhoads discloses providing access to a distal server configured to use a database to identify pertinent information associated with the recognizable symbols based on the symbol information. Rhoads at Fig. 12 (and accompanying description)

Element	'252 Claim Recitation	Exemplary Citations in Reference
		FIG. 12 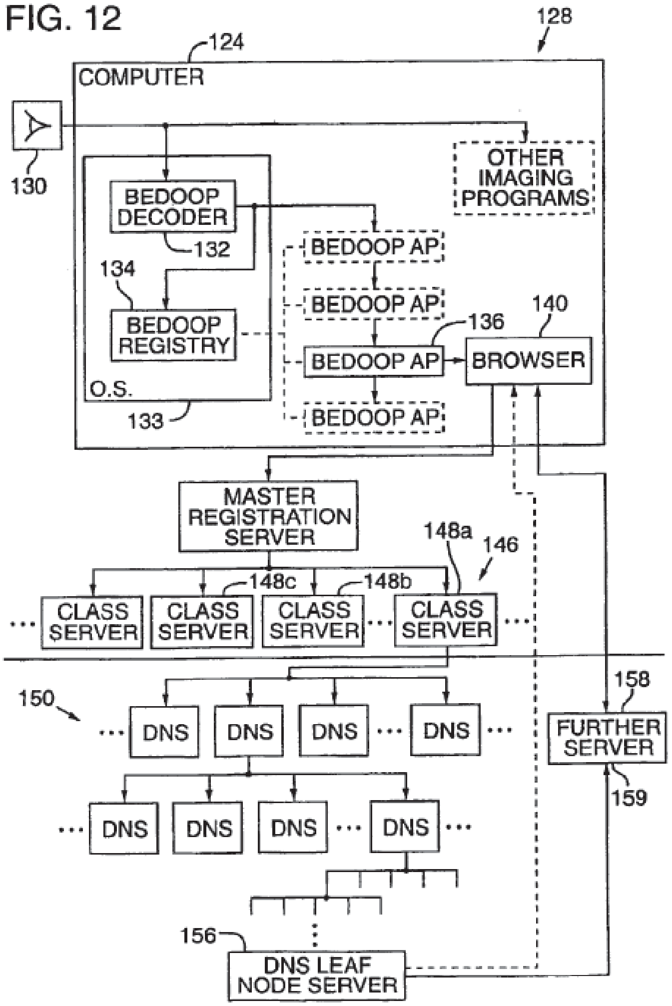 Rhoads at 12:23 – 33 • "Sometimes the computer system may encounter a Bedoop object for which it does not have a registered application program. In such case, a default

Element	'252 Claim Recitation	Exemplary Citations in Reference
		Bedoop application can be invoked. This default application can, e.g., establish an internet link to a remote server computer (or a network of such computers), and can transmit the Bedoop data (or a part of the Bedoop data) to that remote computer. The remote server can undertake the response itself, it can instruct the originating computer how to respond appropriately, or it can undertake some combination of these two responses. (Such arrangements are further considered below.)" Rhoads at 16:66 – 17:10 • "Inside the building the candidate may encounter a security guard. Seeing an unfamiliar person, the guard may visually compare the photo on the badge with the candidate's face. Additionally, the guard can present the badge to a portable Bedoop device, or to one of many Bedoop systems scattered through the building (e g., at every telephone). The Bedoop system extracts the Bedoop data from the card (i.e., from the DMV photograph), interrogates the DMV's DNS server with this Bedoop data, and receives in reply the name of the person depicted in the photograph. (If the Bedoop system is a telephone, the name may be displayed on a small LCD display commonly provided on telephones.)" Rhoads at 22:28–33 • "Magazine (and newspaper) pages can be steganographically encoded with Bedoop data to provide another "paper as portal" experience. As with the earlier-described office document case, the encoded data yields an address to a computer location (e.g., a web page) having the same, or related, content." Rhoads at 45:19–55 • "Bedoop technology can be integrated into portable telecommunication terminals, such as cell phones (manufactured, e.g., by Motorola, Nokia,

Element	'252 Claim Recitation	Exemplary Citations in Reference
		Qualcomm, and others). Such a phone can be equipped with a 1D or 2D image sensor, the output of which is applied to Bedoop decoding circuitry within the phone. This decoding circuitry can be the phone's main CPU, or can be a processing circuit dedicated to Bedoop functionality. (In this as in other embodiments, the decoding can be effected by dedicated hardware, by decoding software executing on a general purpose CPU, etc.) Cell phones are already equipped with numerous features that make them well suited for Bedoop operation. One is that cell phones typically include an LCD or similar screen for display of textual or graphic information, and additionally include buttons or other controls for selecting among menu options presented on the screen (e.g., by moving a cursor). Moreover, cell phones naturally include both audio input and output devices (i.e., microphone and speaker). Still further, the protocol by which cell phones transmit data includes data identifying the phone, so that such data need not be separately encoded. And finally, cell phones obviously provide ready links to remote computer systems. Collectively, these capabilities rival those of the most fully-equipped desktop computer system. Thus, essentially all of the applications detailed elsewhere in this specification can be implemented using cell phone Bedoop systems As with the other Bedoop systems, when Bedoop data is sensed, the phone can respond to the data locally, or it can forward Same over the cellular network to a remote System (or computer network) for handling." Rhoads at 46:10–12. • "Bedoop decoding function within the phone, in other embodiments the image data can be transmitted from the phone and decoded at a remote location." Rhoads at 48:32–53

Element	'252 Claim Recitation	Exemplary Citations in Reference
		• "As noted, the following disclosure focuses on one particular application—a system for linking print media to electronic content It should be reiterated, however, that the technology is not so limited, and may more generally be viewed as a system for linking any object (physical or electronic) to a corresponding networked or local resource. In accordance with an exemplary application of the below-detailed technology, digital watermarking is employed to convey a plural bit identifier within print media, such as magazine advertisements or articles, direct mail coupons or catalogs, bank- or credit-cards, and business cards. This identifier is read by software in a user's computing device and forwarded to a remote database. The remote database identifies a URL corresponding to the identifier, and provides the URL back to the user's computer—permitting a browser on the user's computer to display the URL-identified web page. That web page can provide additional information or services—more timely and/or more extensive than that provided by the print material. By such arrangement, more efficient internet navigation and access is provided to consumers, and more effective means for linking readers to e-commerce points of sale is provided to advertisers." Rhoads at 49:48-65 • "The detailed embodiment contemplates that the object 20 is a magazine advertisement encoded with a steganographic watermark conveying a plural bit object identifier. The watermark is hidden in the advertisement's image in a manner indiscernable to human observers, but detectable by computer analysis. That analysis is performed by a watermark detector 30. Watermark detector 30 can be implemented in various different locations in the system of FIG. 1. Typically, the detector is implemented in the originating device 12, e.g., in the driver software 26, or in application software 28 c that serves to link to external resources based on detected

Element	'252 Claim Recitation	Exemplary Citations in Reference
		watermarks. But it may be implemented elsewhere, e.g., in hardware in the interface electronics 24, in an operating system associated with the device, or outside device 12 altogether. Some systems may have plural watermark detectors, implemented at different locations throughout the system." Rhoads at 50:21-48 • "As more particularly detailed below, the handler 16 provides a response in accordance with the particular watermark payload. The response may be provided directly by the product handler to the device 12, or the handler may respond by communicating with a remote resource 18 (which may be, e.g., a data repository or service provider). In the former case, the handler 16 may identify a URL corresponding to the watermark (using the database 17), and return the URL to the application 28 c. Application 28 c can then pass the URL to a web browser 28 b in the device 12, and initiate a link to the internet site identified by the URL. Or the handler may have some locally stored data (e.g., audio or video, or software updates) and send it to the device 12 in response to the watermark. In the latter case, the handler 16 does not respond directly to the device 12. Instead, the handler responds by communicating with a remote resource 18. The communication can be as simple as logging receipt of the watermark message in a remote repository. Or it can be to authenticate device 12 (or a user thereof) to a remote resource in anticipation of a further transaction (e.g., the communication can form part of an on-line licensing or digital rights management transaction). Or the communication can request the remote resource to provide data or a service back to device 12 or to another destination (e.g., to initiate an FTP file transfer, or to request that a song selection identified by the watermark be downloaded to a user's personal music library, or to update software installed on device 12). In still other cases, hybrids of the two foregoing cases can be employed, e.g., handler 16 can send some data back to device 12, while also communicating with a remote resource 18.

Element	'252 Claim Recitation	Exemplary Citations in Reference
		In some cases, the response returned to the device 12 by handler 6 (or a remote resource 18) can serve to trigger some further action by the device 12. For example, the response returned to device 12 can include a WindowsMedia audio file, together with a request that the device 12 launch the WindowsMedia player installed on the device. (The launching of a browser pointed to a URL is another example of such triggering.)" Rhoads at 51:18-37 • "When shown a watermarked image, the application 28 c analyzes the image and extracts the embedded watermark payload (more particularly detailed below) from the image. The application sends some or all of this information in a message format to the router 14. The router 14 decodes the received message, looking for vendor and product information. Based on this information, it passes the message to a corresponding product handler 16. The product handler receives the message and attempts to match the detected watermark serial number to a registered watermark serial number earlier stored in the database 17. If a match is found, the product handler performs the desired action. As noted, typical actions include returning a URL for web redirection, serving up an HTML page for initial user navigation, initiating software downloads, etc. If a match is not found, the product handler returns an error code and message to the application 28 c. If a match is found, but the corresponding action is unavailable, incomplete, inactive or invalid, the product handler returns an error code and message to the calling To the extent it is found that Rhoads does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Rhoads does not anticipate this claim, Rhoads renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other

Element	'252 Claim Recitation	Exemplary Citations in Reference
		prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Rhoads with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1D]	allowing the mobile device to receive the pertinent information over a network.	Rhoads discloses allowing the mobile device to receive the pertinent information over a network. Rhoads at Fig. 2

Element	'252 Claim Recitation	Exemplary Citations in Reference
		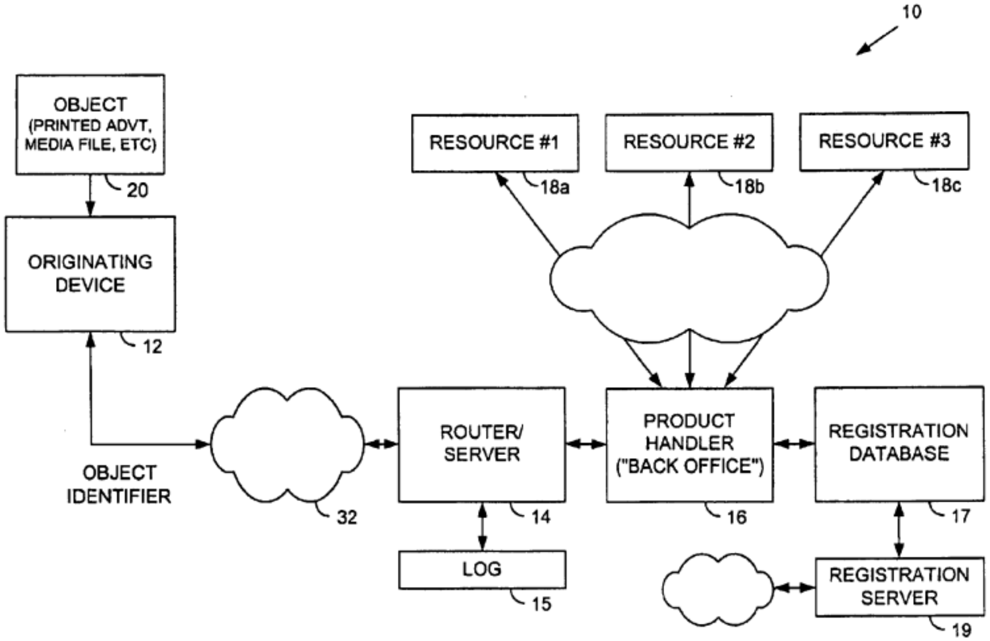 FIG. 2 Rhoads at Fig. 12 (and accompanying description)

Element	'252 Claim Recitation	Exemplary Citations in Reference
		FIG. 12 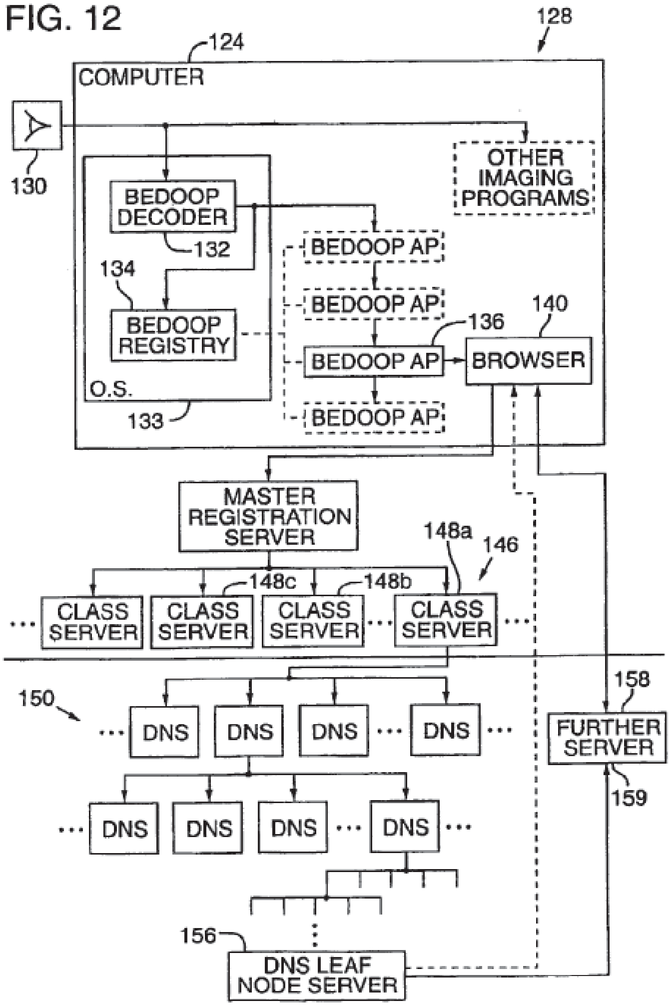 Rhoads at 12:23 – 33 • "Sometimes the computer system may encounter a Bedoop object for which it does not have a registered application program. In such case, a default

Element	'252 Claim Recitation	Exemplary Citations in Reference
		Bedoop application can be invoked. This default application can, e.g., establish an internet link to a remote server computer (or a network of such computers), and can transmit the Bedoop data (or a part of the Bedoop data) to that remote computer. The remote server can undertake the response itself, it can instruct the originating computer how to respond appropriately, or it can undertake some combination of these two responses. (Such arrangements are further considered below.)" Rhoads at 16:66 – 17:10 • "Inside the building the candidate may encounter a security guard. Seeing an unfamiliar person, the guard may visually compare the photo on the badge with the candidate's face. Additionally, the guard can present the badge to a portable Bedoop device, or to one of many Bedoop systems scattered through the building (e g., at every telephone). The Bedoop system extracts the Bedoop data from the card (i.e., from the DMV photograph), interrogates the DMV's DNS server with this Bedoop data, and receives in reply the name of the person depicted in the photograph. (If the Bedoop system is a telephone, the name may be displayed on a small LCD display commonly provided on telephones.)" Rhoads at 49:19-47 • "Device 12 interacts with an object 20. The object can be electronic or not. Electronic objects can include computer files, representations of audio, video, or still imagery (e.g., files or in streaming form), etc. Non-electronic objects can include physical objects such as newspapers, magazine pages, posters, product packaging, event tickets, credit cards, paper currency, etc. Non-electronic objects can also include sounds produced by loudspeakers. When used with nonelectronic objects, device 12 (FIG. 2) typically includes some form of sensor or transducer 22 to produce electronic signals or data corresponding to the object. Examples include CCD- or CMOS-based

Element	'252 Claim Recitation	Exemplary Citations in Reference
		optical sensors (either as part of still- or video cameras, flatbed scanners, mice, or otherwise), microphones, barcode scanners, RF ID sensors, mag stripe readers, etc. In such cases, the sensor 22 may be coupled to associated interface electronics 24, which in turn may be coupled to device driver software 26, which in turn may be coupled to one or more application programs 28. Device driver software 26 serves as a software interface, communicating at a relatively high level with the application programs 28 (e.g., through API instructions whose content and format are standardized to facilitate application programming), and at a relatively low level with the interface electronics 24." Rhoads at 49:48-65 • "The detailed embodiment contemplates that the object 20 is a magazine advertisement encoded with a steganographic watermark conveying a plural bit object identifier. The watermark is hidden in the advertisement's image in a manner indiscernable to human observers, but detectable by computer analysis. That analysis is performed by a watermark detector 30. Watermark detector 30 can be implemented in various different locations in the system of FIG. 1. Typically, the detector is implemented in the originating device 12, e.g., in the driver software 26, or in application software 28 c that serves to link to external resources based on detected watermarks. But it may be implemented elsewhere, e.g., in hardware in the interface electronics 24, in an operating system associated with the device, or outside device 12 altogether. Some systems may have plural watermark detectors, implemented at different locations throughout the system." Rhoads at 50:21–51:3 • "As more particularly detailed below, the handler 16 provides a response in accordance with the particular watermark payload. The response may be

Element	'252 Claim Recitation	Exemplary Citations in Reference
		provided directly by the product handler to the device 12, or the handler may respond by communicating with a remote resource 18 (which may be, e.g., a data repository or service provider). In the former case, the handler 16 may identify a URL corresponding to the watermark (using the database 17), and return the URL to the application 28 c. Application 28 c can then pass the URL to a web browser 28 b in the device 12, and initiate a link to the internet site identified by the URL. Or the handler may have some locally stored data (e.g., audio or video, or software updates) and send it to the device 12 in response to the watermark. In the latter case, the handler 16 does not respond directly to the device 12. Instead, the handler responds by communicating with a remote resource 18. The communication can be as simple as logging receipt of the watermark message in a remote repository. Or it can be to authenticate device 12 (or a user thereof) to a remote resource in anticipation of a further transaction (e.g., the communication can form part of an on-line licensing or digital rights management transaction). Or the communication can request the remote resource to provide data or a service back to device 12 or to another destination (e.g., to initiate an FTP file transfer, or to request that a song selection identified by the watermark be downloaded to a user's personal music library, or to update software installed on device 12). In still other cases, hybrids of the two foregoing cases can be employed, e.g., handler 16 can send some data back to device 12, while also communicating with a remote resource 18. In some cases, the response returned to the device 12 by handler 6 (or a remote resource 18) can serve to trigger some further action by the device 12. For example, the response returned to device 12 can include a WindowsMedia audio file, together with a request that the device 12 launch the WindowsMedia player installed on the device. (The launching of a

Element	'252 Claim Recitation	Exemplary Citations in Reference
		browser pointed to a URL is another example of such triggering.) The illustrated product handler 16 comprises essentially the same hardware elements as the router 14, e.g., CPU, memory, etc. Although FIG. 2 shows just one product handler, several product handlers can be included in the system—either co-located of geographically distributed. Different handlers can be dedicated to different functions (e.g., serving URLs, serving music, etc.) or to different watermark sources (e.g., one responds to watermarks found in audio, another responds to watermarks found in print advertising, etc.)." Rhoads at Fig. 12 (and accompanying description)

Element	'252 Claim Recitation	Exemplary Citations in Reference
		FIG. 12 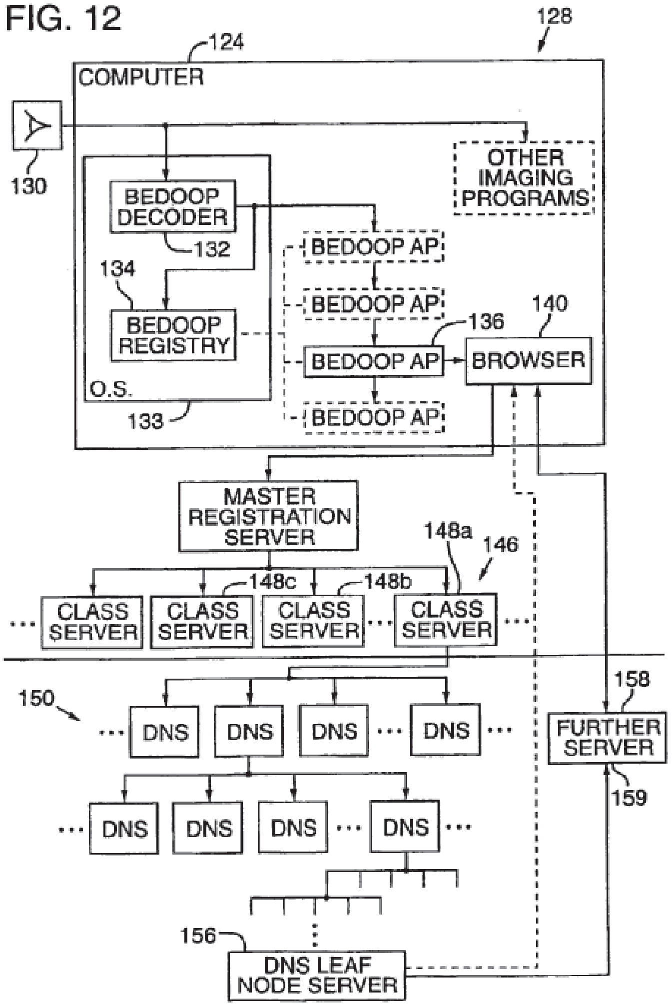

Element	'252 Claim Recitation	Exemplary Citations in Reference
		To the extent it is found that Rhoads does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Rhoads does not anticipate this claim, Rhoads renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Rhoads with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

2. Claim 26

Claim	'252 Claim Recitation	Exemplary Citations in Reference
26	The method of claim 18, wherein the data relating to the image comprises the image.	Rhoads discloses the data relating to the image comprises the image. Rhoads at Fig. 2 (and accompanying description)

		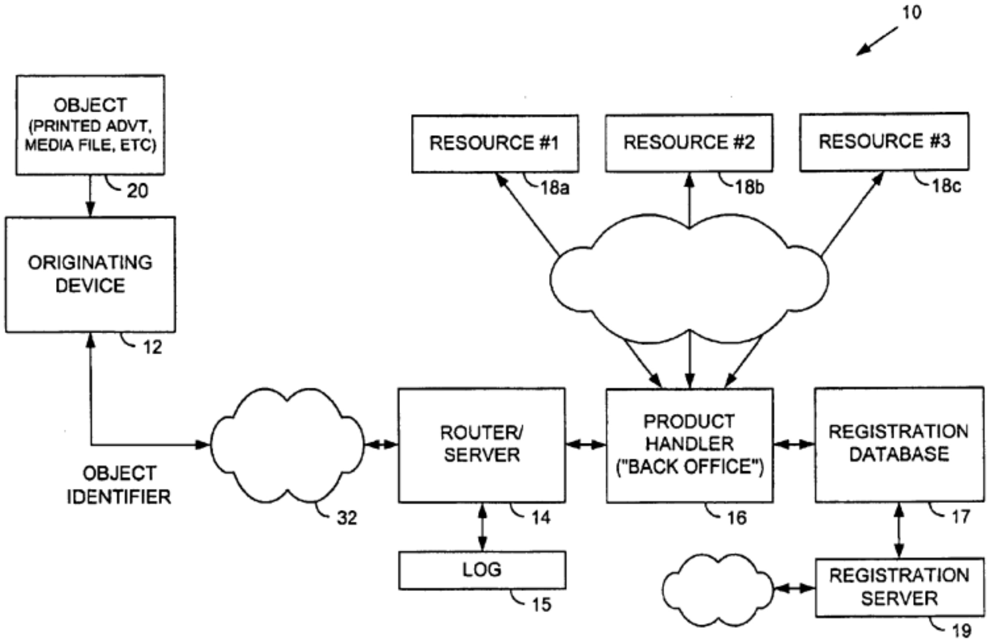 FIG. 2 Rhoads at Fig. 3 (and accompanying description)

		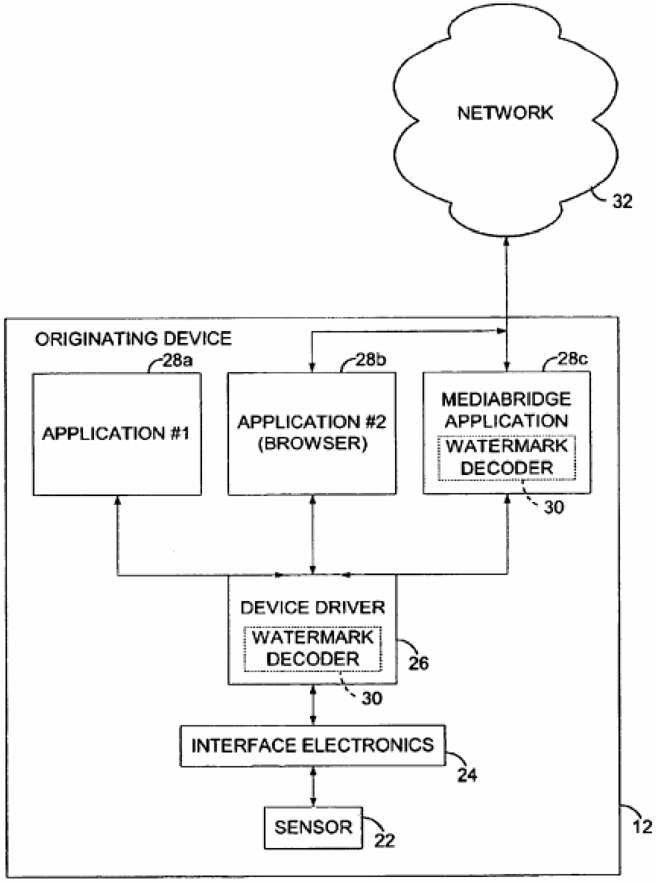 **FIG. 3** Rhoads at 12:23 – 33 • "Sometimes the computer system may encounter a Bedoop object for which it does not have a registered application program. In such case, a default Bedoop application can be invoked. This default application can,

		e.g., establish an internet link to a remote server computer (or a network of such computers), and can transmit the Bedoop data (or a part of the Bedoop data) to that remote computer. The remote server can undertake the response itself, it can instruct the originating computer how to respond appropriately, or it can undertake some combination of these two responses. (Such arrangements are further considered below.)" Rhoads at 16:66 – 17:10 • "Inside the building the candidate may encounter a security guard. Seeing an unfamiliar person, the guard may visually compare the photo on the badge with the candidate's face. Additionally, the guard can present the badge to a portable Bedoop device, or to one of many Bedoop systems scattered through the building (e g., at every telephone). The Bedoop system extracts the Bedoop data from the card (i.e., from the DMV photograph), interrogates the DMV's DNS server with this Bedoop data, and receives in reply the name of the person depicted in the photograph. (If the Bedoop system is a telephone, the name may be displayed on a small LCD display commonly provided on telephones.)" To the extent it is found that Rhoads does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Rhoads does not anticipate this claim, Rhoads renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Rhoads with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results."

		Rhoads at 49:19-47 • "Device **12** interacts with an object **20**. The object can be electronic or not. Electronic objects can include computer files, representations of audio, video, or still imagery (e.g., files or in streaming form), etc. Non-electronic objects can include physical objects such as newspapers, magazine pages, posters, product packaging, event tickets, credit cards, paper currency, etc. Non-electronic objects can also include sounds produced by loudspeakers. When used with nonelectronic objects, device **12** (FIG. 2) typically includes some form of sensor or transducer **22** to produce electronic signals or data corresponding to the object. Examples include CCD- or CMOS-based optical sensors (either as part of still- or video cameras, flatbed scanners, mice, or otherwise), microphones, barcode scanners, RF ID sensors, mag stripe readers, etc. In such cases, the sensor **22** may be coupled to associated interface electronics **24**, which in turn may be coupled to device driver software **26**, which in turn may be coupled to one or more application programs **28**. Device driver software **26** serves as a software interface, communicating at a relatively high level with the application programs **28** (e.g., through API instructions whose content and format are standardized to facilitate application programming), and at a relatively low level with the interface electronics **24**." Rhoads at 49:48-65 • "The detailed embodiment contemplates that the object 20 is a magazine advertisement encoded with a steganographic watermark conveying a plural bit object identifier. The watermark is hidden in the advertisement's image in a manner indiscernable to human observers, but detectable by computer analysis. That analysis is performed by a watermark detector 30. Watermark detector 30 can be implemented in various different locations

		in the system of FIG. 1. Typically, the detector is implemented in the originating device 12, e.g., in the driver software 26, or in application software 28 c that serves to link to external resources based on detected watermarks. But it may be implemented elsewhere, e.g., in hardware in the interface electronics 24, in an operating system associated with the device, or outside device 12 altogether. Some systems may have plural watermark detectors, implemented at different locations throughout the system." To the extent it is found that Rhoads does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Rhoads does not anticipate this claim, Rhoads renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Rhoads with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

3. Claim 27

Claim	'252 Claim Recitation	Exemplary Citations in Reference
27	The method of claim 26, wherein image processing further comprises compressing the image to form the data relating to the image.	Rhoads discloses image processing further comprises compressing the image to form the data relating to the image. Rhoads at 38:1-19 • "Post-It® Notes

		Pads of Post-It® notes, or other pads of paper, can be marked by the manufacturer (either by texturing, watermarked tinting, ink-jet spattering, etc.) to convey steganographic data (e.g., Bedoop data). When such a note is presented to a Bedoop system, the system may launch an application that stores a snapshot of the note. More particularly, the application may mask the note-portion of the image data from the other image data, virtually re-map it to a rectangular format of standardized pixel dimensions, JPEG-compress the resulting image, and store it in a particular computer subdirectory with a name indicating the date of image acquisition, together with the color and/or size of the note. (These latter two data may be indicated by data included in the Bedoop payload.) If the color of the note is indicated by digital data (e.g, in the file name), then the image itself may be stored in grey-scale." Rhoads at 46:47 – 53 - "The provision of image sensors in cell phones enables other functionality. One is the capture of still or video imagery. Such image data can be compressed (typically by lossy processes such as MPEG, JPEG, or the like, implemented with dedicated hardware CODECs) and transmitted with the audio data. The screens on such phones can likewise be used for display of incoming image or video data." To the extent it is found that Rhoads does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Rhoads does not anticipate this claim, Rhoads renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Rhoads with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to

		address similar known problems, according to known methods which would yield predictable results.

4. Claim 29

Claim	'252 Claim Recitation	Exemplary Citations in Reference
29	The method of claim 18, wherein image processing further includes converting at least a portion of the data relating to the image into a grayscale image.	Rhoads discloses image processing further includes converting at least a portion of the data relating to the image into a grayscale image. Rhoads at 38:1-17 • "Post-It® Notes Pads of Post-It® notes, or other pads of paper, can be marked by the manufacturer (either by texturing, watermarked tinting, ink-jet spattering, etc.) to convey steganographic data (e.g., Bedoop data). When such a note is presented to a Bedoop system, the system may launch an application that stores a snapshot of the note. More particularly, the application may mask the note-portion of the image data from the other image data, virtually re-map it to a rectangular format of standardized pixel dimensions, JPEG-compress the resulting image, and store it in a particular computer subdirectory with a name indicating the date of image acquisition, together with the color and/or size of the note. (These latter two data may be indicated by data included in the Bedoop payload.) If the color of the note is indicated by digital data (e.g, in the file name), then the image itself may be stored in grey-scale." Rhoads at 50:21-48 • "As more particularly detailed below, the handler 16 provides a response in accordance with the particular watermark payload. The response may be provided directly by the product handler to the device 12, or the

		handler may respond by communicating with a remote resource 18 (which may be, e.g., a data repository or service provider). In the former case, the handler 16 may identify a URL corresponding to the watermark (using the database 17), and return the URL to the application 28 c. Application 28 c can then pass the URL to a web browser 28 b in the device 12, and initiate a link to the internet site identified by the URL. Or the handler may have some locally stored data (e.g., audio or video, or software updates) and send it to the device 12 in response to the watermark. In the latter case, the handler 16 does not respond directly to the device 12. Instead, the handler responds by communicating with a remote resource 18. The communication can be as simple as logging receipt of the watermark message in a remote repository. Or it can be to authenticate device 12 (or a user thereof) to a remote resource in anticipation of a further transaction (e.g., the communication can form part of an on-line licensing or digital rights management transaction). Or the communication can request the remote resource to provide data or a service back to device 12 or to another destination (e.g., to initiate an FTP file transfer, or to request that a song selection identified by the watermark be downloaded to a user's personal music library, or to update software installed on device 12). In still other cases, hybrids of the two foregoing cases can be employed, e.g., handler 16 can send some data back to device 12, while also communicating with a remote resource 18. In some cases, the response returned to the device 12 by handler 6 (or a remote resource 18) can serve to trigger some further action by the device 12. For example, the response returned to device 12 can include a WindowsMedia audio file, together with a request that the device 12 launch the WindowsMedia player installed on the device. (The launching of a browser pointed to a URL is another example of such triggering.)" To the extent it is found that Rhoads does not expressly disclose this limitation,

		the limitation is inherent. And to the extent it is found that Rhoads does not anticipate this claim, Rhoads renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Rhoads with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

5. Claim 31

Claim	'252 Claim Recitation	Exemplary Citations in Reference
31	The method of claim 18, further comprising the distal server using database to identify the pertinent information based on image characteristics of the image derived by the image processing platform based on the data relating to the image.	Rhoads discloses the distal server using database to identify the pertinent information based on image characteristics of the image derived by the image processing platform based on the data relating to the image. Rhoads at Fig. 2 (and accompanying description)

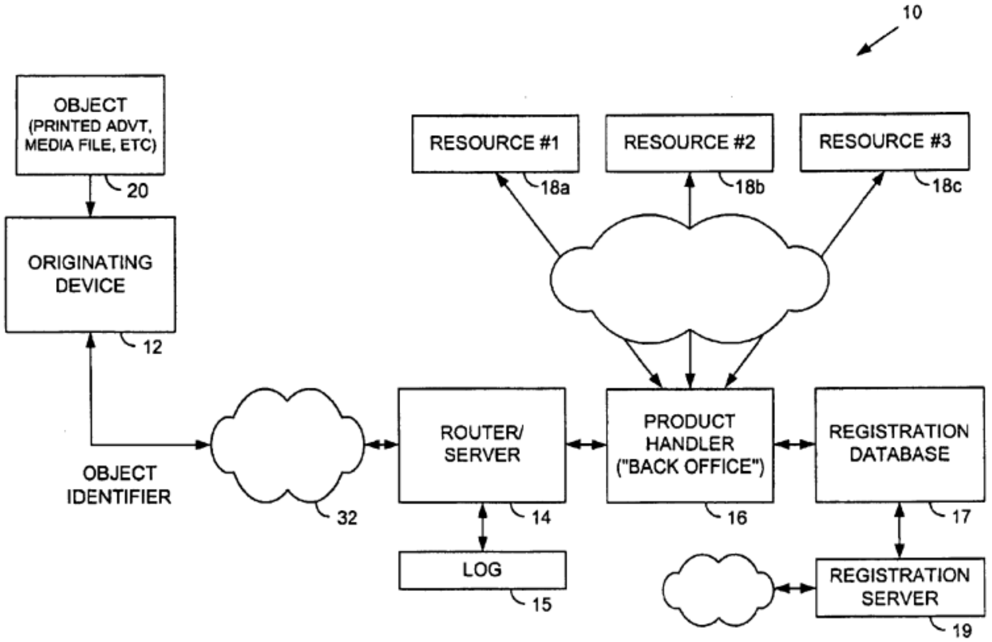

FIG. 2

Rhoads at 7:13-29

- "Referring to FIG. 11, a basic embodiment 110 of the present technology includes **an optical sensor 112**, a computer 114, and a network connection 116 to the internet 118. The illustrated optical sensor 112 is a digital camera having a resolution of 320 by 200 pixels (color or black and white) that stares out, grabbing frames of image data five times per second and storing same in one or more frame buffers. These frames of image data are analyzed by a computer 114 for the presence of Bedoop data. (Essentially, Bedoop data is any form of plural-bit data encoding recognized by the system 110—data which, in many embodiments, initiates some action.) Once detected, the system responds in accordance with the detected Bedoop data (e.g., by initiating some local action, or by

		communication with a remote computer, such as over the internet, via an online service such as AOL, or using point-to-point dial-up communications, as with a bulletin board system)." Rhoads 7:30 – 46 • "Consider the milk carton example. The artwork on a milk carton can be adapted to convey Bedoop data. In the preferred embodiment, the Bedoop data is steganographically encoded (e.g., digitally watermarked) on the carton Numerous digital watermarking techniques are known—all of which convey data in a hidden form (i.e., on human inspection, it is not apparent that digitally encoded data is present). Exemplary techniques operate by slightly changing the luminance, or contours of selected points on artwork or text printed on the carton, or by splattering tiny droplets of ink on the carton in a seemingly random pattern. Each of these techniques has the effect of changing the local luminance at areas across the carton—luminance changes that can be detected by the computer 114 and decoded to extract the encoded digital data. In the case of a milk carton, the data may serve to identify the object as, e.g., a half gallon carton of Alpenrose brand skim milk." Rhoads at 9:35–10:21 • "The Steganographic decoding process may entail three Steps. In the first, the object is located. In the Second, the objects orientation is discerned. In the third, the Bedoop data is extracted from the image data corresponding to the Bedoop object. The first step, object location, can be assisted by various clues. One is the placement of the object; typically the center of the image field will be a point on the object. The surrounding data can then be analyzed to try and discern the object's boundaries. Another location technique is slight movement. Although the user will

		typically try to hold the object still, there will usually be some jitter of the Bedoop object within the image frame (e.g., a few pixels back and forth). Background visual clutter, in contrast, will typically be stationary. Such movement may thus be sensed and used to identify the Bedoop object from within the image data. Still another object-location clue is object shape. Many Bedoop objects are rectangular in shape (or trapezoidal as viewed by the camera). Straight edge boundaries can thus be used to define an area of likely Bedoop data. Color is a further object identification clue that may be useful in some contexts. Yet another object location clue is spatial frequency. In imaging systems with well defined focal zones, undesired visual clutter may be at focal distances that results in blurring. The Bedoop object, in contrast, will be in focus and may be characterized by fine detail. Analyzing the image data for the high frequencies associated with fine detail can be used to distinguish the intended object from others. Characteristic markings on the object (as discussed below in connection with determining object orientation) can also be sensed and used in locating the object. Once the Bedoop object has been located within the image data, masking can be applied (if desired) to eliminate image data not corresponding to the intended object. The next step in the decoding process, determining orientation of the Bedoop data, can likewise be discerned by reference to visual clues. For example, some objects include subliminal graticule data, or other calibration data, steganographically encoded with the Bedoop data to aid in determining orientation. Others can employ overt markings, either

		placed for that sole purpose (e.g. reference lines or fiducials), or serving another purpose as well (e.g. lines of text), to discern orientation. Edge-detection algorithms can also be employed to deduce the orientation of the object by reference to its edges." Rhoads at 16:66 – 17:10 • "Inside the building the candidate may encounter a security guard. Seeing an unfamiliar person, the guard may visually compare the photo on the badge with the candidate's face. Additionally, the guard can present the badge to a portable Bedoop device, or to one of many Bedoop systems scattered through the building (e g., at every telephone). The Bedoop system extracts the Bedoop data from the card (i.e., from the DMV photograph), interrogates the DMV's DNS server with this Bedoop data, and receives in reply the name of the person depicted in the photograph. (If the Bedoop system is a telephone, the name may be displayed on a small LCD display commonly provided on telephones.)" Rhoads at 22:9–27 • "To aid in identification, the Bedoop objects (e.g., badges) can be colored a distinctive color, permitting the system to more easily identify candidate objects from other items within the optical capture devices' field of view. Or the object can be provided with a retro-reflective coating, and the elevator can be equipped with one or more illumination sources of known spectral or temporal quality (e.g., constant infra red, or constant illumination with a single- or multi-line spectrum, or a pulsed light source of known periodicity; LEDs or semiconductor lasers, each with an associated diffuser, can be used for each the foregoing and can be paired with the image capture devices). Other such tell-tale clues can likewise be used to aid in object location. In all such cases, the optical capture device can sense the tell-tale clue(s) using a wide field of view sensor. The device can then be physically or electronically steered, and/or zoomed, to acquire

		a higher resolution image of the digitally-encoded object suitable for decoding.” Rhoads at 23:33–54 • “When a recipient of a business card holds it in front of a Bedoop sensor, the operating system on the local system launches a local Bedoop application. That local Bedoop application, in turn, establishes an external internet connection to a remote business card server. The address of that server may already be known to the local Bedoop application (e.g., having been stored from previous use), or the local Bedoop system can traverse the above-described public network of DNS servers to reach the business card server. A database on the business card name server maintains a large collection of business card data, one database record per UID. When that server receives Bedoop data from a local Bedoop system, it parses out the UID and accesses the corresponding database record. This record typically includes more information than is commonly printed on conventional business cards. Sample fields from the record may include, for example, name, title, office phone, office fax, home phone, home fax, cellular phone, email address, company name, corporate web page address, personal web page address, secretary's name, spouse's name, and birthday. This record is transmitted back to the originating Bedoop system.” Rhoads at 48:32–53 • “As noted, the following disclosure focuses on one particular application—a system for linking print media to electronic content It should be reiterated, however, that the technology is not so limited, and may more generally be viewed as a system for linking any object (physical or electronic) to a corresponding networked or local resource.

		In accordance with an exemplary application of the below-detailed technology, digital watermarking is employed to convey a plural bit identifier within print media, such as magazine advertisements or articles, direct mail coupons or catalogs, bank- or credit-cards, and business cards. This identifier is read by software in a user's computing device and forwarded to a remote database. The remote database identifies a URL corresponding to the identifier, and provides the URL back to the user's computer—permitting a browser on the user's computer to display the URL-identified web page. That web page can provide additional information or services—more timely and/or more extensive than that provided by the print material. By such arrangement, more efficient internet navigation and access is provided to consumers, and more effective means for linking readers to e-commerce points of sale is provided to advertisers." Rhoads at 50:21-48 • "As more particularly detailed below, the handler 16 provides a response in accordance with the particular watermark payload. The response may be provided directly by the product handler to the device 12, or the handler may respond by communicating with a remote resource 18 (which may be, e.g., a data repository or service provider). In the former case, the handler 16 may identify a URL corresponding to the watermark (using the database 17), and return the URL to the application 28 c. Application 28 c can then pass the URL to a web browser 28 b in the device 12, and initiate a link to the internet site identified by the URL. Or the handler may have some locally stored data (e.g., audio or video, or software updates) and send it to the device 12 in response to the watermark. In the latter case, the handler 16 does not respond directly to the device 12. Instead, the handler responds by communicating with a remote resource 18. The communication can be as simple as logging receipt of the watermark message in a remote repository. Or it can be to authenticate

		device 12 (or a user thereof) to a remote resource in anticipation of a further transaction (e.g., the communication can form part of an on-line licensing or digital rights management transaction). Or the communication can request the remote resource to provide data or a service back to device 12 or to another destination (e.g., to initiate an FTP file transfer, or to request that a song selection identified by the watermark be downloaded to a user's personal music library, or to update software installed on device 12). In still other cases, hybrids of the two foregoing cases can be employed, e.g., handler 16 can send some data back to device 12, while also communicating with a remote resource 18. In some cases, the response returned to the device 12 by handler 6 (or a remote resource 18) can serve to trigger some further action by the device 12. For example, the response returned to device 12 can include a WindowsMedia audio file, together with a request that the device 12 launch the WindowsMedia player installed on the device. (The launching of a browser pointed to a URL is another example of such triggering.)" Rhoads at 51:18-37 • "When shown a watermarked image, the application 28 c analyzes the image and extracts the embedded watermark payload (more particularly detailed below) from the image. The application sends some or all of this information in a message format to the router 14. The router 14 decodes the received message, looking for vendor and product information. Based on this information, it passes the message to a corresponding product handler 16. The product handler receives the message and attempts to match the detected watermark serial number to a registered watermark serial number earlier stored in the database 17. If a match is found, the product handler performs the desired action. As noted, typical actions include returning a

		URL for web redirection, serving up an HTML page for initial user navigation, initiating software downloads, etc. If a match is not found, the product handler returns an error code and message to the application 28 c. If a match is found, but the corresponding action is unavailable, incomplete, inactive or invalid, the product handler returns an error code and message to the calling application." Rhoads at 84:11 – 17 • "B55. A method comprising: presenting an object to an optical sensor, the optical sensor producing image data, the object being steganographically encoded with plural bits of digital data; decoding said digital data from the image data; and in response to said digital data, displaying an image of a like object." To the extent it is found that Rhoads does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Rhoads does not anticipate this claim, Rhoads renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Rhoads with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

Exhibit E-4

Bank of America Corporation ("BofA" or "Defendant") provides this chart subject to all reservations, objections, statements, and disclaimers set forth herein and in BofA's Final Invalidity Contentions Cover Pleading, as well as any amendment, supplement, or modification thereof, which are incorporated herein by reference in their entirety.

On information and belief, Mault was filed on September 7, 2001, claiming the priority date of September 7, 2000. Mault published at least as of April 25, 2002, and is available as prior art at least under 35 U.S.C. § 102(e).

As shown in the chart below, Mault anticipates asserted claims 18, 27, 29 and 31 of the '252 Patent. To the extent Mault is found not to expressly disclose certain limitations in the asserted claims, such limitations are inherent. To the extent Mault is found not to anticipate any asserted claims or claim elements of the '252 Patent, the reference nevertheless renders those claims or claim elements obvious under 35 U.S.C. § 103, either alone or in combination with other art identified herein or in the cover pleading of BofA's Final Invalidity Contentions, for the reasons identified in the cover pleading. These Final Invalidity Contentions are not an admission by BofA that the accused products, including any current or past versions of the accused products, are covered by, or infringe the asserted claims, but are based instead on the recognition that if the claims are interpreted to be broad enough to encompass the accused products, the claims must also be construed to have that same scope when considering whether they are invalid.

The following chart is based on the Court's claim construction rulings (Dkt. Nos. 145, 153, 236), and to the extent not contradictory, Plaintiffs' apparent interpretations of the claims in Plaintiffs' "Preliminary" Final Infringement Contentions. BofA notes that in many instances, Plaintiffs' "Preliminary" Final Infringement Contentions ignore and/or contradict the Court's claim construction rulings. BofA does not accept the assumptions concerning the scope and meaning implicit in Plaintiffs' Preliminary Final Infringement Contentions, to the extent those assumptions are discernible, and reserves the right to challenge Plaintiffs' proposed (or implied) constructions. Furthermore, Plaintiffs' "Preliminary" Final Contentions fail to comply with the level of specificity required by SPR 2.1.2 and 2.1.3, and include improper reservation of rights to supplement with "additional ways in which BoA Accused Products infringe." Any alleged deficient discovery is of Plaintiffs own making. Most if not all of the deficiencies in Plaintiffs' Preliminary Final Contentions have been documented by Defendants through multiple letters since at least May 18, 2021 and up to Bank of America's most recent letter of August 19, 2023. Plaintiffs' shifting infringement contentions, deficiencies in its infringement contentions, and changing claim scope have prejudiced BofA from being able to reasonably prepare its defenses, including its invalidity contentions. BofA also reserves the right to revise and supplement these charts if and when Plaintiffs are permitted to provide revised or supplemental Infringement Contentions. Citations given in the chart below are merely representative of the respective elements to provide notice to Plaintiffs of BofA's invalidity theories and are not meant to be exhaustive.

1. Claim 18

Element	'252 Claim Recitation	Exemplary Citations in Reference
[1Pre]	A method for retrieving information from image processing, the method comprising:	Mault discloses a method for retrieving information from image processing. Mault at Abstract • "A method for assisting a person create a record of food items consumed by the person comprises the capturing of food images corresponding to food items consumed, for example by creating an image using an optical image sensor; recording the food images in the memory of an electronic device; viewing the food images on the display of the electronic device, so as to identify food items consumed; selecting food item identifiers corresponding to the food images, and recording food item identifiers in a memory, So as to create a record of food items consumed." Mault at [0038] • "A person can be provided with a portable electronic device, having an image creation capability. For example, the portable electronic device may be a personal digital assistant (PDA) having an optical image sensor. The optical image sensor can be an accessory to the PDA, or can be within the housing of the PDA." To the extent it is found that Mault does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Mault does not anticipate this claim, Mault renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Mault with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success because each of these

Element	'252 Claim Recitation	Exemplary Citations in Reference
		references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1A]	providing a mobile device having a camera the mobile device configured to capture an image and configured to transmit data relating to the image to an image processing platform;	Mault discloses providing a mobile device having a camera the mobile device configured to capture an image and configured to transmit data relating to the image to an image processing platform. Mault at Fig. 5. 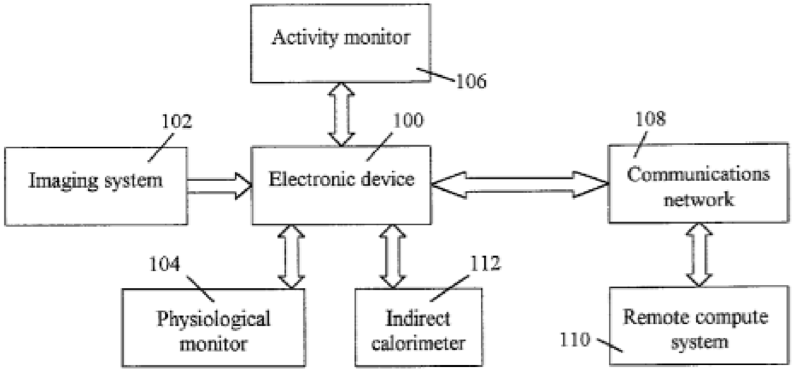 Figure 5 Mault at Figs. 12 and 13.

Element	'252 Claim Recitation	Exemplary Citations in Reference
		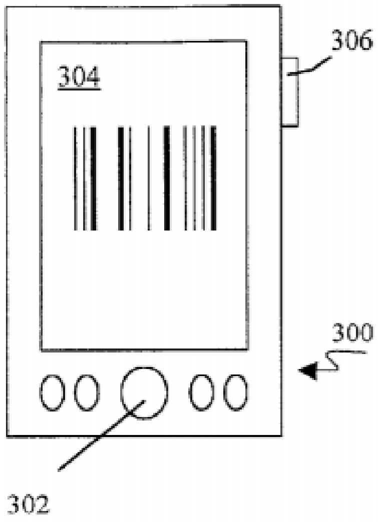 Figure 12

Element	**'252 Claim Recitation**	**Exemplary Citations in Reference**
		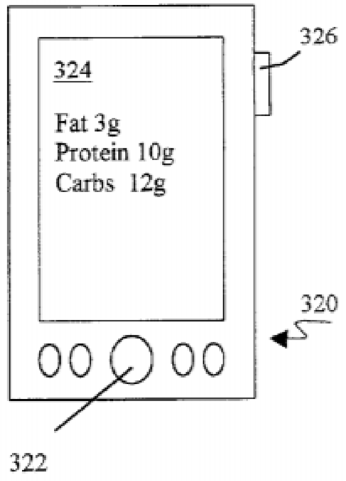 Figure 13 Mault at [0008] - "The image can be stored, for example, on the memory of an imaging device such as a camera, a memory of a device in communication with the imaging device, a memory module (such as a memory stick), a memory of a portable computing device, or on a memory associated with a remote computer system, such as one accessible over a communications network. In this context, the term image refers to an electronic file, or equivalent, containing information corresponding to an image. An example would be a conventional JPEG file." Mault at [0009] - "After recording the image, at a convenient later time the person can retrieve a stored image and view the image on a display. The display can be part of the imaging device, or other device such as a personal digital assistant, personal computer, Internet appliance, WebTV, e-book, tablet computer, pager, cell

Element	'252 Claim Recitation	Exemplary Citations in Reference
		phone, interactive TV, spectacle mounted display, visor, and the like. An electronic device having a display used to view an image may be in communication with a separate device used to create the image, for example using wireless or cable-based links." Mault at [0011] • "The imaging device, used to create an image of foods consumed, can be a digital camera, device having the functionality of a camera, such as a personal digital assistant (PDA), cell phone, electronic book, pager, calculator, or other consumer electronics device." Mault at [0013] • "An imaging device can comprise elements such as an optical imaging sensor such as a charge coupled device (CCD), a memory, a memory slot adapted to receive a memory module such as a non-volatile memory module, an electrical power source, a transceiver, a clock, a display, an electrical interface, an audible signal generator, an indicator light, a data port (such as a socket for a cable), an IR transmitter or transceiver, or other wireless transceiver." Mault at [0027] • " FIG. 5 illustrates a system allowing diet log creation;" Mault at [0038–39] • "A person can be provided with a portable electronic device, having an image creation capability. For example, the portable electronic device may be a personal digital assistant (PDA) having an optical image sensor. The optical image sensor can be an accessory to the PDA, or can be within the housing of the PDA. The portable electronic device may also be a digital camera, wireless telephone, pager, organizer, calculator (numeric calculator, diet calculator,

Element	'252 Claim Recitation	Exemplary Citations in Reference
		graphing calculator), spectacle mounted device, wrist mounted device, or other electronic device having image production capabilities. FIG. 1A shows a schematic of an electronic device which may be used in embodiments of the present invention, having a processor **10**, an image capturing system **12**, a data entry mechanism **14**, a memory **16**, a clock **18**, and display **18**." Mault at [0046] • "Other options which may be presented to the user include the entering of a food by name, for example into diet log software, or scanning the barcode of a food item. The electronic device may comprise or communicate with a barcode scanner, or the optical image sensor can be used to record an image of the barcode. The barcode can then be interpreted in terms of a UPC (universal product code) and/or product name using a database accessible to the PDA. A UPC database may be accessible over a communications network such as the Internet." Mault at [0048] • FIG. 5 is a schematic of a system comprising an electronic device 100, an imaging system 102, a physiological sensor 104, an activity sensor 106, and a communications link through communications network 108 to a remote computer system (such as a server system) 110. The electronic device can be a PDA or other device carried by the person. The imaging device can be included with the electronic device in a unitary device. The physiological sensor may detect eating, for example a blood glucose sensor can detect the rise in blood glucose after a meal, allowing the electronic device to alert the person to make a diet log entry. The activity sensor provides a signal correlated with a physical activity level of the person.

Element	'252 Claim Recitation	Exemplary Citations in Reference
		Mault at [0054] • "FIG. 9 shows a schematic of a system in which an imaging system **220**, such as a camera, transmits images over a communications network **222** to a remote computer system **224**." Mault at [0079] • "Embodiments of the present invention allow an electronic diet log to be created by any person with access to the recorded images. After recording an image using the PDA, the image may be transferred to another location using a communications network. For example, the image may be transmitted via a wireless Internet connection to another computer system." To the extent it is found that Mault does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Mault does not anticipate this claim, Mault renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Mault with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1B]	configuring the image processing platform to receive the data relating to the image and to conduct image processing, including:	Mault discloses configuring the image processing platform to receive the data relating to the image and to conduct image processing. Mault at Fig. 5

Element	'252 Claim Recitation	Exemplary Citations in Reference
		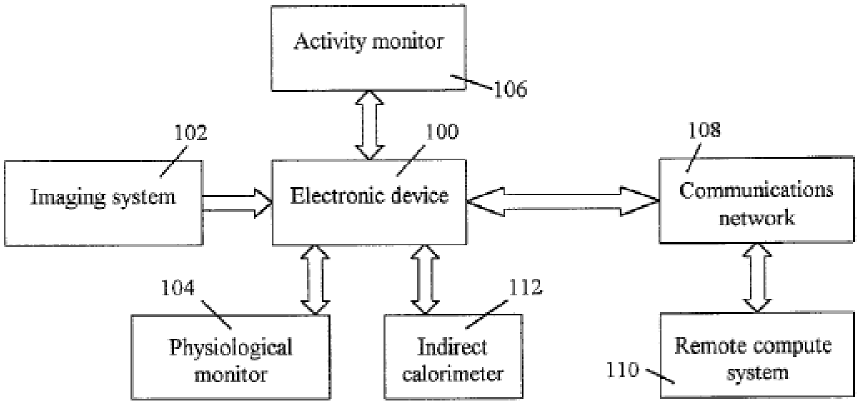 Figure 5 Mault at [0017] • "A method of creating a diet log comprises: recording an image of a label, the label corresponding to a food item consumed; analyzing the image of the label (for example using optical character recognition, image analysis, image recognition, or neural network techniques) so as to determine the identity of the food item consumed; and storing the determined identity of the food item consumed in an electronic memory, so as to create a diet log. The label can be a printed label, a barcode, an alphanumeric code, a nutritional data label, a list of ingredients, a product code, graphical designs (Such as trademarks), product names, and the like." Mault at [0018] • "A method by which a person can create a record of foods consumed comprises: recording a food image of a food consumed; transmitting the food image to a remote computer system, for example over a communications network such as the internet, receiving food identification data from the remote computer system (which may be generated by a software running on the remote computer System, a computer expert system, an image analysis

Element	'252 Claim Recitation	Exemplary Citations in Reference
		system, a human operative in communication with the remote computer system, or some combination of techniques); and recording the food identification data So as to create a record of foods consumed." Mault at [0027] - " FIG. 5 illustrates a system allowing diet log creation;" Mault at [0048] - FIG. 5 is a schematic of a system comprising an electronic device 100, an imaging system 102, a physiological sensor 104, an activity sensor 106, and a communications link through communications network 108 to a remote computer system (such as a server system) 110. The electronic device can be a PDA or other device carried by the person. The imaging device can be included with the electronic device in a unitary device. The physiological sensor may detect eating, for example a blood glucose sensor can detect the rise in blood glucose after a meal, allowing the electronic device to alert the person to make a diet log entry. The activity sensor provides a signal correlated with a physical activity level of the person. Mault at [0049] - "Images recorded with the imaging device are transmitted over the communications network to the remote computer. The remote computer system can be a server system, and may comprise a computer expert system for image analysis. Image analysis can be performed by one or more persons, a computer expert system, or other mechanism." Mault at [0051] - "FIG. 6 is a flow chart of a method for diet logging. Box **120** corresponds to recording an image of food. Box **122** corresponds to transmission of the image to a remote location, such as a remote computer system. Other data can be

Element	'252 Claim Recitation	Exemplary Citations in Reference
		transmitted with the image, for example a sales receipt, data provided by the food vendor, location the image was taken (which may be correlated with a restaurant identity), a menu image, and the weight of the food (which may be provided by a mat)." To the extent it is found that Mault does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Mault does not anticipate this claim, Mault renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Mault with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1B.i]	operating on the data relating to the image to determine if the image contains one or more recognizable symbols; and	Mault discloses operating on the data relating to the image to determine if the image contains one or more recognizable symbols. *See* limitation [1B] above. Mault at [0017] • "A method of creating a diet log comprises: recording an image of a label, the label corresponding to a food item consumed; analyzing the image of the label (for example using optical character recognition, image analysis, image recognition, or neural network techniques) so as to determine the identity of the food item consumed; and storing the determined identity of the food item consumed in an electronic memory, so as to create a diet log. The label can be a printed label, a barcode, an alphanumeric code, a nutritional data label, a list

Element	'252 Claim Recitation	Exemplary Citations in Reference
		of ingredients, a product code, graphical designs (Such as trademarks), product names, and the like." Mault at [0018]. • "A method by which a person can create a record of foods consumed comprises: recording a food image of a food consumed; transmitting the food image to a remote computer system, for example over a communications network such as the internet, receiving food identification data from the remote computer system (which may be generated by a software running on the remote computer system, a computer expert system, an image analysis system, a human operative in communication with the remote computer system, or some combination of techniques); and recording the food identification data so as to create a record of foods consumed." Mault at [0046] • "Other options which may be presented to the user include the entering of a food by name, for example into diet log software, or scanning the barcode of a food item. The electronic device may comprise or communicate with a barcode scanner, or the optical image sensor can be used to record an image of the barcode. The barcode can then be interpreted in terms of a UPC (universal product code) and/or product name using a database accessible to the PDA. A UPC database may be accessible over a communications network such as the Internet. Other options include receiving transmissions such as wireless data transmissions from food vendors or food vending devices, viewing the recorded diet log, editing the diet log, or analyzing previously recorded images." Mault at [0051] • "Box **124** corresponds to analysis of the image of the remote location. Image recognition software may be used. A trained person, at a remote location, can select food names corresponding to the images, for example in conjunction with

Element	'252 Claim Recitation	Exemplary Citations in Reference
		image recognition, providing names where image recognition has not been successful in identifying the food. A computer expert system, optical character recognition of a name on the package, barcode recognition, portion size estimate, and other food data can also be extracted from the image or images. Box **126** corresponds to determining food identity and/or nutritional content of the food within the image." Mault at [0075–77] • "Image processing techniques can be used to identify food items from images. For packaged foods, optical character recognition may be used to record and identify nutritional information and/or identify the item consumed. Computer analysis of images may be carried out on the PDA or using another computer system in communication with the PDA. For example, an image of food packaging may be used to identify the food contained. An image of the nutritional content information panel provided by the manufacturer on the package may also be recorded. Optical character recognition may be used to obtain information from the image for storage in a diet log. An image of a box of corn flakes may be recognized by a computer as such, and used to generate a corn flake serving record in a diet log. For non-packaged foods, e.g. restaurant meals, image analysis may also be used in identifying the food items. Imaging at a number of wavelengths, followed by false color image generation, can be used to help identify image components. Spectroscopic imaging can be used in computer-assisted food recognition. Image processing, image recognition, and pattern recognition algorithms are useful in analyzing stored images so as to create a diet log. Algorithms may be applied to color images, or images recorded at a number of different wavelengths (e.g. in the IR, optical, and UV) which may assist identification. Stored images can be identified by comparison with previously stored images that the person has identified, using a software learning mechanism such as a

Element	'252 Claim Recitation	Exemplary Citations in Reference
		neural network." Mault at [0080] • "Barcodes can also be used in food item identification. The PDA may also incorporate a bar-code Scanner, or alternatively, barcodes may be identified from a recorded image of the bar-code using image analysis. The barcode (e.g. universal product code) can be used to retrieve nutritional data relating to the item from a database. The database may reside on a remote computer System (remote in this context meaning a computer System not carried by the person, for example a commercial Server System, home computer, etc.). The person may scan the barcode of a box of cereal and the UPC code used to retrieve nutrition information specific to that brand of cereal from the database." Mault at [0081] • "FIG. 12 shows an image of a barcode presented on the display **304** of PDA **300** (the PDA having data entry buttons **302**, and an image sensor **306**). Image analysis software can be used to analyze the image of the barcode, so as to determine UPC, product name, nutritional data, and the like." Mault at [0082] • "FIG. 13 shows an image of nutritional information on the display **324** of a PDA **320**, having data entry buttons **322** and optical image Sensor **326**. Optical character recognition Software running on a processor of the PDA can be used to extract nutritional information from the image, and place it into memory for possible later transfer to a database." Mault at [0090] • "Colors, shapes, printed codes, and fluorescent markers of tablets can be used to assist identification, either by a person or an [sic] software image analysis program."

Element	'252 Claim Recitation	Exemplary Citations in Reference
		To the extent it is found that Mault does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Mault does not anticipate this claim, Mault renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Mault with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1B.ii]	decoding the recognizable symbols to extract symbol information by analyzing the recognizable symbols according to type;	Mault discloses decoding the recognizable symbols to extract symbol information by analyzing the recognizable symbols according to type . *See* limitation [1B.i] above. To the extent it is found that Mault does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Mault does not anticipate this claim, Mault renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Mault with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1C]	providing access to a distal server configured to use a database to	Mault discloses providing access to a distal server configured to use a database to identify pertinent information associated with the recognizable symbols based on

Element	'252 Claim Recitation	Exemplary Citations in Reference
	identify pertinent information associated with the recognizable symbols based on the symbol information; and	the symbol information. Mault at Fig. 4 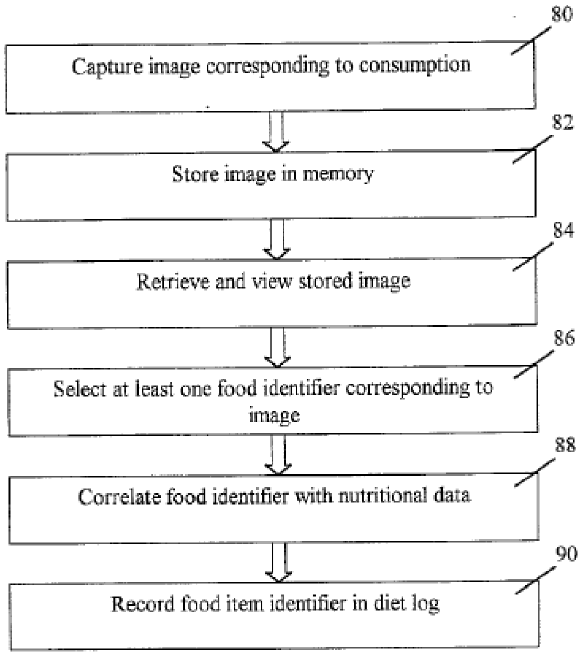 Figure 4 Mault at [0017] • "The identity of food items consumed can be correlated with nutritional data, for example using a nutritional database, and the nutritional data can stored in an electronic memory so as to create a nutritional record, for example as part of a diet log." Mault at [0047]

Element	'252 Claim Recitation	Exemplary Citations in Reference
		• "Box **88** corresponds to selecting a food item identifier, such as a food name, corresponding to the displayed image. The selection can be made from a menu displayed to the person. Box **88** corresponds to correlating a food item identifier with nutritional data. Box **90** corresponds to recording the food item identifier and corresponding nutritional data in memory, so as to create a diet log." Mault at [0059] • "For example, a person can sit down to a meal and record an image of the meal using the electronic device. The image is transmitted to a remote location where it is analyzed and converted into a caloric value. The caloric value can be transmitted back to the person, allowing the person to decide whether to eat the meal, choose another meal, or eat some portion of the meal consistent with reaching a calorie balance goal." Mault at [0082–83] • "FIG. 13 shows an image of nutritional information on the display **324** of a PDA **320**, having data entry buttons **322** and optical image sensor **326**. Optical character recognition software running on a processor of the PDA can be used to extract nutritional information from the image, and place it into memory for possible later transfer to a database. Food item identities and nutritional information can also be received from food vendors, e.g. grocery stores, on-line retailers, restaurants, vending machines, diet food retailers, etc. The PDA may prompt the person to identify the source of the food being imaged, and this information can be used to help in identification." To the extent it is found that Mault does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Mault does not anticipate this claim, Mault renders it obvious, either alone, in combination with the

Element	'252 Claim Recitation	Exemplary Citations in Reference
		knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Mault with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1D]	allowing the mobile device to receive the pertinent information over a network.	Mault discloses allowing the mobile device to receive the pertinent information over a network. Mault at [0052] • "A diet log channel, generated by remote server (or other remote computer systems, or other content provider) **176**, and received by the set top box over communications network **178**, can be selected for display on the interactive television **160**. A display of food item identifiers, such as food names in a menu format, can be displayed on the display **162** of the interactive TV **160**. An image stored in the local memory **172** can be added to the diet log channel signal, using techniques known in the art, so that the person views the diet log channel with an embedded image of the food consumed. The person then selects items corresponding to the food consumed from the displayed menu. Selection data is generated with user input **146** (for example a keyboard) and transmitted to the set top box and to the remote server. The display is changed appropriately on receipt of the selection data by the content provider and the selection is added to the diet log of the person, which is stored in a database associated with or otherwise accessible by the remote server **176**." Mault at [0051] • "FIG. 6 is a flow chart of a method for diet logging. Box **120** corresponds to

Element	'252 Claim Recitation	Exemplary Citations in Reference
		recording an image of food. Box **122** corresponds to transmission of the image to a remote location, such as a remote computer system. Other data can be transmitted with the image, for example a sales receipt, data provided by the food vendor, location the image was taken (which may be correlated with a restaurant identity), a menu image, and the weight of the food (which may be provided by a mat). Box **124** corresponds to analysis of the image of the remote location. Image recognition software may be used. A trained person, at a remote location, can select food names corresponding to the images, for example in conjunction with image recognition, providing names where image recognition has not been successful in identifying the food. A computer expert system, optical character recognition of a name on the package, barcode recognition, portion size estimate, and other food data can also be extracted from the image or images. Box **126** corresponds to determining food identity and/or nutritional content of the food within the image. Box **128** corresponds to storing the food identity and nutritional content into a database, so as to create a diet log. Box **130** corresponds to the person reviewing the diet log for accuracy at a convenient later time." Mault at [0059] • "For example, a person can sit down to a meal and record an image of the meal using the electronic device. The image is transmitted to a remote location where it is analyzed and converted into a caloric value. The caloric value can be transmitted back to the person, allowing the person to decide whether to eat the meal, choose another meal, or eat some portion of the meal consistent with reaching a calorie balance goal." Mault at [0074] • "FIG. 11 shows a PDA **280** with display **284**, buttons **282**, and imaging device **286** directed at a can of pickles **288**. The display **284** can be used to indicate the image to be captured. A button such as **282** may be pressed to

Element	'252 Claim Recitation	Exemplary Citations in Reference
		record the image. The display **184** can be used to view captured images of foods stored in the memory of the PDA, for assisting the person in creating a diet log." Mault at [0075] • "Optical character recognition may be used to obtain information from the image for storage in a diet log. An image of a box of corn flakes may be recognized by a computer as such, and used to generate a corn flake serving record in a diet log. Computer analysis may be used as a first attempt in producing a diet log, and if unsuccessful, the image can be passed to a human for analysis." To the extent it is found that Mault does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Mault does not anticipate this claim, Mault renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Mault with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

2. Claim 26

Claim	'252 Claim Recitation	Exemplary Citations in Reference
26	The method of claim 18, wherein the data relating to the image comprises the image.	Mault discloses the data relating to the image comprises the image. Mault at Fig. 5

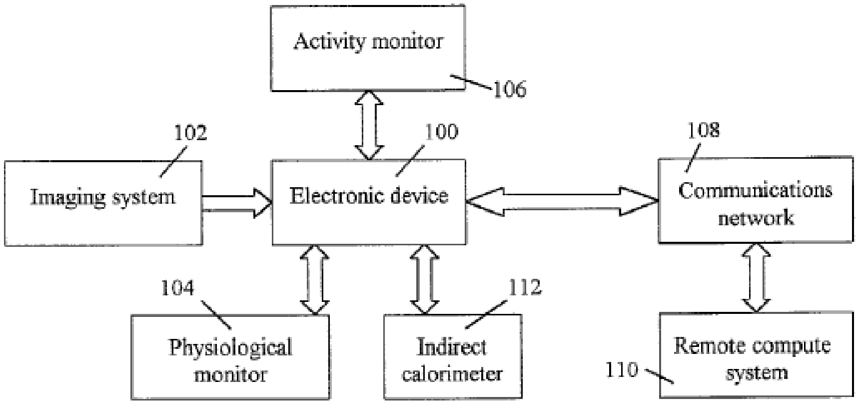

Figure 5

Mault at [0008]

- "The image can be stored, for example, on the memory of an imaging device such as a camera, a memory of a device in communication with the imaging device, a memory module (such as a memory stick), a memory of a portable computing device, or on a memory associated with a remote computer system, such as one accessible over a communications network. In this context, the term image refers to an electronic file, or equivalent, containing information corresponding to an image. An example would be a conventional JPEG file."

Mault at [0009]

- "After recording the image, at a convenient later time the person can retrieve a stored image and view the image on a display. The display can be part of the imaging device, or other device such as a personal digital assistant, personal computer, Internet appliance, WebTV, e-book, tablet computer, pager, cell phone, interactive TV, spectacle mounted display, visor, and the like. An electronic device having a display used to view an image may be in communication with a separate device used to create the image, for example

		using wireless or cable-based links.” Mault at [0011] • “The imaging device, used to create an image of foods consumed, can be a digital camera, device having the functionality of a camera, such as a personal digital assistant (PDA), cell phone, electronic book, pager, calculator, or other consumer electronics device. The imaging device can also be an ornamental component, clothing item, or other functional or decorative item having image creation capabilities, Such as spectacles, buttons, a pen, cutlery item, key ring, keys, belt, cigarette-like object, ring, body ornament, or jewelry.” Mault at [0017] • “A method of creating a diet log comprises: recording an image of a label, the label corresponding to a food item consumed; analyzing the image of the label (for example using optical character recognition, image analysis, image recognition, or neural network techniques) so as to determine the identity of the food item consumed; and storing the determined identity of the food item consumed in an electronic memory, so as to create a diet log. The label can be a printed label, a barcode, an alphanumeric code, a nutritional data label, a list of ingredients, a product code, graphical designs (Such as trademarks), product names, and the like.” Mault at [0018] • “A method by which a person can create a record of foods consumed comprises: recording a food image of a food consumed; transmitting the food image to a remote computer system, for example over a communications network Such as the internet, receiving food identification data from the remote computer System (which may be generated by a software running on the remote computer System, a computer expert System, an image analysis system, a human operative in communication with the

		remote computer system, or Some combination of techniques); and recording the food identification data So as to create a record of foods consumed." Mault at [0038–39] • "A person can be provided with a portable electronic device, having an image creation capability. For example, the portable electronic device may be a personal digital assistant (PDA) having an optical image sensor. The optical image sensor can be an accessory to the PDA, or can be within the housing of the PDA. The portable electronic device may also be a digital camera, wireless telephone, pager, organizer, calculator (numeric calculator, diet calculator, graphing calculator), spectacle mounted device, wrist mounted device, or other electronic device having image production capabilities. FIG. 1A shows a schematic of an electronic device which may be used in embodiments of the present invention, having a processor **10**, an image capturing system **12**, a data entry mechanism **14**, a memory **16**, a clock **18**, and display **18**." Mault at [0046] • "Other options which may be presented to the user include the entering of a food by name, for example into diet log software, or scanning the barcode of a food item. The electronic device may comprise or communicate with a barcode Scanner, or the optical image Sensor can be used to record an image of the barcode. The barcode can then be interpreted in terms of a UPC (universal product code) and/or product name using a database accessible to the PDA. A UPC database may be accessible over a communications network such as the Internet." Mault at [0049] • "Images recorded with the imaging device are transmitted over the communications network to the remote computer. The remote computer

		system can be a server system, and may comprise a computer expert system for image analysis. Image analysis can be performed by one or more persons, a computer expert system, or other mechanism." Mault at [0051] • "Box 124 corresponds to analysis of the image of the remote location. Image recognition software may be used. A trained person, at a remote location, can select food names corresponding to the images, for example in conjunction with image recognition, providing names where image recognition has not been successful in identifying the food. A computer expert System, optical character recognition of a name on the package, barcode recognition, portion size estimate, and other food data can also be extracted from the image or images." Mault at [0054] • "FIG. 9 shows a schematic of a system in which an imaging system **220**, such as a camera, transmits images over a communications network **222** to a remote computer system **224**." Mault at [0073] • "Fig. 11 shows a PDA 280 with display 284, buttons 282, and imaging device 286 directed at a can of pickles 288. The display 284 can be used to indicate the image to be captured. A button such as 282 may be pressed to record the image. The display 184 can be used to view captured images of foods stored in the memory of the PDA, for assisting the person in creating a diet log." Mault at [0080] • "Barcodes can also be used in food item identification. The PDA may also incorporate a bar-code scanner, or alternatively, barcodes may be identified from a recorded image of the bar-code using image analysis. The barcode

		(e.g. universal product code) can be used to retrieve nutritional data relating to the item from a database. The database may reside on a remote computer System (remote in this context meaning a computer system not carried by the person, for example a commercial server System, home computer, etc.). The person may scan the barcode of a box of cereal and the UPC code used to retrieve nutrition information specific to that brand of cereal from the database." Mault at [0081] • "FIG. 12 shows an image of a barcode presented on the display 304 of PDA 300 (the PDA having data entry buttons 302, and an image sensor 306). Image analysis software can be used to analyze the image of the barcode, So as to determine UPC, product name, nutritional data, and the like. Alternatively, a barcode Scanner can be used to read the barcode, or an accessory device which in communication with the PDA may be used as a barcode scanner." Mault at Fig. 12

		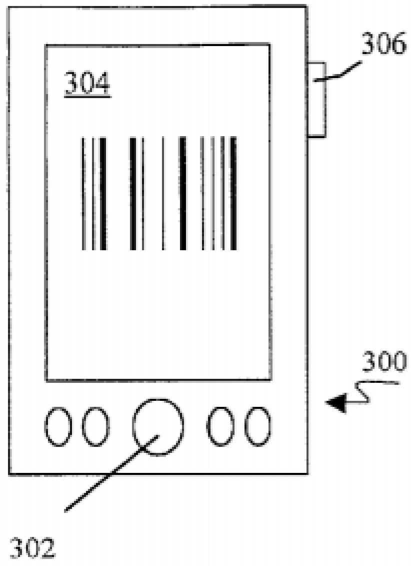 Figure 12 Mault at [0082] • "FIG. 13 shows an image of nutritional information on the display 324 of a PDA320, having data entry buttons 322 and optical image Sensor 326. Optical character recognition Software running on a processor of the PDA can be used to extract nutritional information from the image, and place it into memory for possible later transfer to a database." Mault at Fig. 13

<table>
<tr><td></td><td></td><td>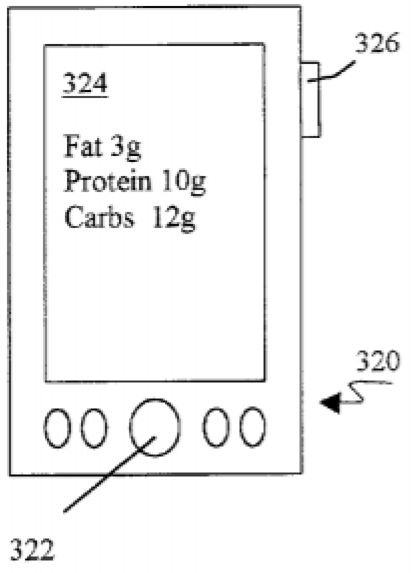
Figure 13
To the extent it is found that Mault does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Mault does not anticipate this claim, Mault renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Preliminary Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Mault with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.</td></tr>
</table>

3. Claim 27

Claim	'252 Claim Recitation	Exemplary Citations in Reference
27	The method of claim 26, wherein image processing further comprises compressing the image to form the data relating to the image.	Mault discloses image processing further comprises compressing the image to form the data relating to the image. Mault at [0008] • "The image can be stored, for example, on the memory of an imaging device such as a camera, a memory of a device in communication with the imaging device, a memory module (such as a memory stick), a memory of a portable computing device, or on a memory associated with a remote computer system, such as one accessible over a communications network. In this context, the term image refers to an electronic file, or equivalent, containing information corresponding to an image. An example would be a conventional JPEG file." To the extent it is found that Mault does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Mault does not anticipate this claim, Mault renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Mault with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

4. Claim 29

Claim	'252 Claim Recitation	Exemplary Citations in Reference
29	The method of claim 18, wherein image processing further includes converting at least a portion of the	To the extent it is found that Mault does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Mault does not anticipate this claim, Mault renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with

	data relating to the image into a grayscale image.	other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Mault with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

5. Claim 31

Claim	'252 Claim Recitation	Exemplary Citations in Reference
31	The method of claim 18, further comprising the distal server using database to identify the pertinent information based on image characteristics of the image derived by the image processing platform based on the data relating to the image.	Mault discloses the distal server using database to identify the pertinent information based on image characteristics of the image derived by the image processing platform based on the data relating to the image. Mault at [0017] • "A method of creating a diet log comprises: recording an image of a label, the label corresponding to a food item consumed; analyzing the image of the label (for example using optical character recognition, image analysis, image recognition, or neural network techniques) so as to determine the identity of the food item consumed; and storing the determined identity of the food item consumed in an electronic memory, so as to create a diet log. The label can be a printed label, a barcode, an alphanumeric code, a nutritional data label, a list of ingredients, a product code, graphical designs (Such as trademarks), product names, and the like." Mault at [0018] • "A method by which a person can create a record of foods consumed comprises: recording a food image of a food consumed; transmitting the food image to a remote computer system, for example over a communications

		network such as the internet, receiving food identification data from the remote computer System (which may be generated by a software running on the remote computer System, a computer expert system, an image analysis system, a human operative in communication with the remote computer system, or some combination of techniques); and recording the food identification data So as to create a record of foods consumed." Mault at [0046] • "Other options which may be presented to the user include the entering of a food by name, for example into diet log software, or scanning the barcode of a food item. The electronic device may comprise or communicate with a barcode scanner, or the optical image sensor can be used to record an image of the barcode. The barcode can then be interpreted in terms of a UPC (universal product code) and/or product name using a database accessible to the PDA. A UPC database may be accessible over a communications network such as the Internet. Other options include receiving transmissions such as wireless data transmissions from food vendors or food vending devices, viewing the recorded diet log, editing the diet log, or analyzing previously recorded images." Mault at [0051] • "Box 124 corresponds to analysis of the image of the remote location. Image recognition software may be used. A trained person, at a remote location, can select food names corresponding to the images, for example in conjunction with image recognition, providing names where image recognition has not been successful in identifying the food. A computer expert System, optical character recognition of a name on the package, barcode recognition, portion size estimate, and other food data can also be extracted from the image or images." Mault at [0075–76]

		• "Image processing techniques can be used to identify food items from images. For packaged foods, optical character recognition may be used to record and identify nutritional information and/or identify the item consumed. Computer analysis of images may be carried out on the PDA or using another computer system in communication with the PDA. For example, an image of food packaging may be used to identify the food contained. An image of the nutritional content information panel provided by the manufacturer on the package may also be recorded. Optical character recognition may be used to obtain information from the image for storage in a diet log. An image of a box of corn flakes may be recognized by a computer as such, and used to generate a corn flake serving record in a diet log. For non-packaged foods, e.g. restaurant meals, image analysis may also be used in identifying the food items. Imaging at a number of wavelengths, followed by false color image generation, can be used to help identify image components. Spectroscopic imaging can be used in computer-assisted food recognition." Mault at [0077] • "Image processing, image recognition, and pattern recognition algorithms are useful in analyzing stored images so as to create a diet log. Algorithms may be applied to color images, or images recorded at a number of different wavelengths (e.g. in the IR, optical, and UV) which may assist identification. Stored images can be identified by comparison with previously stored images that the person has identified, using a software learning mechanism such as a neural network." Mault at [0080] • "Barcodes can also be used in food item identification. The PDA may also incorporate a bar-code Scanner, or alternatively, barcodes may be identified from a recorded image of the bar-code using image analysis. The barcode

		(e.g. universal product code) can be used to retrieve nutritional data relating to the item from a database. The database may reside on a remote computer System (remote in this context meaning a computer System not carried by the person, for example a commercial Server System, home computer, etc.). The person may scan the barcode of a box of cereal and the UPC code used to retrieve nutrition information specific to that brand of cereal from the database." Mault at [0081] • "FIG. 12 shows an image of a barcode presented on the display **304** of PDA **300** (the PDA having data entry buttons **302**, and an image sensor **306**). Image analysis software can be used to analyze the image of the barcode, so as to determine UPC, product name, nutritional data, and the like." Mault at [0082] • "FIG. 13 shows an image of nutritional information on the display **324** of a PDA **320**, having data entry buttons **322** and optical image Sensor **326**. Optical character recognition Software running on a processor of the PDA can be used to extract nutritional information from the image, and place it into memory for possible later transfer to a database." Mault at [0090] • "Colors, shapes, printed codes, and fluorescent markers of tablets can be used to assist identification, either by a person or an [sic] software image analysis program." To the extent it is found that Mault does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Mault does not anticipate this claim, Mault renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final

		Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Mault with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

Exhibit E-6

Bank of America Corporation ("BofA" or "Defendant") provides this chart subject to all reservations, objections, statements, and disclaimers set forth herein and in BofA's Final Invalidity Contentions Cover Pleading, as well as any amendment, supplement, or modification thereof, which are incorporated herein by reference in their entirety.

On information and belief, the application that matured into Harris was filed on July 18, 2000. Harris published at least as of December 23, 2003, and is available as prior art at least under 35 U.S.C. § 102(e).

As shown in the chart below, Harris anticipates asserted claims 18, 27, 29, and 31 of the '252 Patent. To the extent Harris is found not to expressly disclose certain limitations in the asserted claims, such limitations are inherent. To the extent Harris is found not to anticipate any asserted claims or claim elements of the '252 Patent, the reference nevertheless renders those claims or claim elements obvious under 35 U.S.C. § 103, either alone or in combination with other art identified herein or in the cover pleading of BofA's Final Invalidity Contentions, for the reasons identified in the cover pleading. These Final Invalidity Contentions are not an admission by BofA that the accused products, including any current or past versions of the accused products, are covered by, or infringe the asserted claims, but are based instead on the recognition that if the claims are interpreted to be broad enough to encompass the accused products, the claims must also be construed to have that same scope when considering whether they are invalid.

The following chart is based on the Court's claim construction rulings (Dkt. Nos. 145, 153, 236), and to the extent not contradictory, Plaintiffs' apparent interpretations of the claims in Plaintiffs' "Preliminary" Final Infringement Contentions. BofA notes that in many instances, Plaintiffs' "Preliminary" Final Infringement Contentions ignore and/or contradict the Court's claim construction rulings. BofA does not accept the assumptions concerning the scope and meaning implicit in Plaintiffs' Preliminary Final Infringement Contentions, to the extent those assumptions are discernible, and reserves the right to challenge Plaintiffs' proposed (or implied) constructions. Furthermore, Plaintiffs' "Preliminary" Final Contentions fail to comply with the level of specificity required by SPR 2.1.2 and 2.1.3, and include improper reservation of rights to supplement with "additional ways in which BoA Accused Products infringe." Any alleged deficient discovery is of Plaintiffs own making. Most if not all of the deficiencies in Plaintiffs' Preliminary Final Contentions have been documented by Defendants through multiple letters since at least May 18, 2021 and up to Bank of America's most recent letter of August 19, 2023. Plaintiffs' shifting infringement contentions, deficiencies in its infringement contentions, and changing claim scope have prejudiced BofA from being able to reasonably prepare its defenses, including its invalidity contentions. BofA also reserves the right to revise and supplement these charts if and when Plaintiffs are permitted to provide revised or supplemental Infringement Contentions. Citations given in the chart below are merely representative of the respective elements to provide notice to Plaintiffs of BofA's invalidity theories and are not meant to be exhaustive.

1. Claim 18

Element	'252 Claim Recitation	Exemplary Citations in Reference
[18Pre]	A method for retrieving information from image processing, the method comprising:	
[18A]	providing a mobile device having a camera the mobile device configured to capture an image and configured to transmit data relating to the image an image processing platform;	Harris at FIG. 1A and FIG. 1B 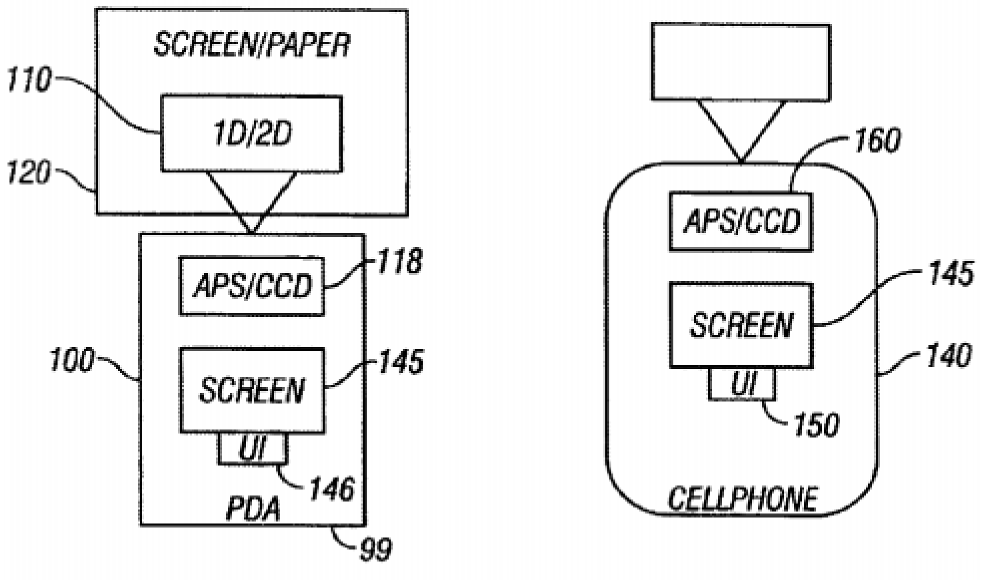 FIG. 1A FIG. 1B Harris at 2: 3-19 • "The bar codes can be imaged/scanned in a number of different ways. One

Element	'252 Claim Recitation	Exemplary Citations in Reference
		embodiment discloses using a camera to input and decode these bar codes. The embodiments are shown in FIGS. 1A and 1B. FIG. 1A describes using **a personal digital assistant 100** as the input device. One preferred input device of this type is the Palm V(TM) type hand held computer. A bar code scanner can be used, such as the commercially available Symbol Technology SPT 1700. Alternatively, a camera add-on unit can be added to the Palm V and used as described herein. The PDA includes a screen 145 and user interface 146 all within the same housing 99. FIG. 1B shows the client being **a cellular telephone** which also may include a screen 145 and user interface 150. The cellular telephone can be associated with either a dedicated bar code scanner or an image sensor 160 of the type used to obtain photographs for video telephony." Harris at FIG. 4

Element	'252 Claim Recitation	Exemplary Citations in Reference
		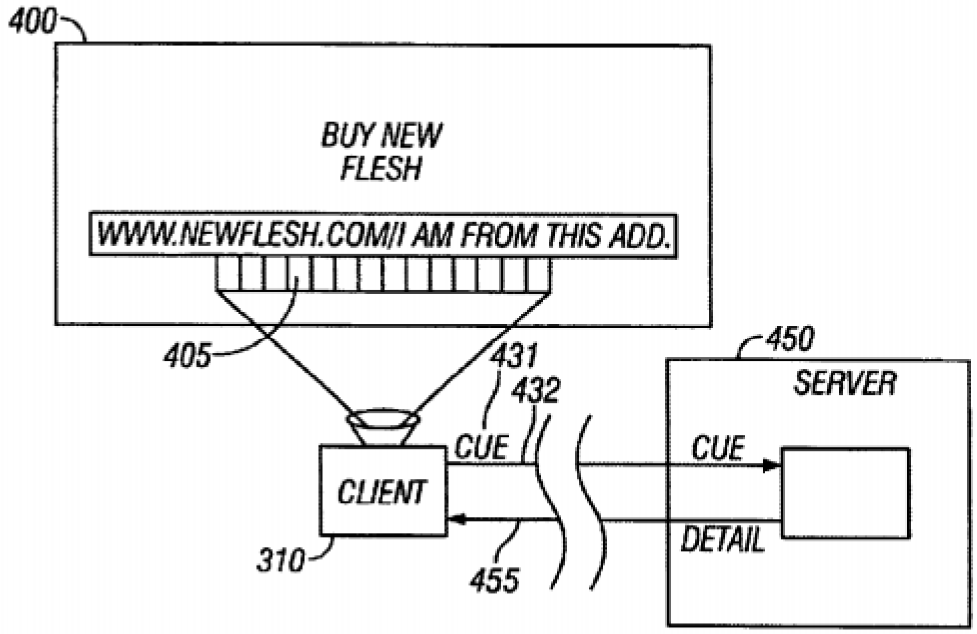 FIG. 4 Harris at 4:15-30 • "The advertisement 400 includes a bar code 405 therein. The bar code 405 is associated with the advertisement, and includes some additional information about the advertisement. For example, the bar code may include the web site address of the company preparing the advertisement, or appointment information about the advertisement, or a "vcf file" or the like which is an importable file with address information about the company or author sponsoring the advertisement. In operation, the user brings one of the clients, either the cellular telephone 310 or PDA 315, into range of the bar code 405. The client reads the bar

Element	'252 Claim Recitation	Exemplary Citations in Reference
		code and decodes it as noted above. The decoded information can represent ASCII information, a compressed file such as a zip file, G code information, or any other compressed or non-compressed information." To the extent it is found that Harris does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Harris does not anticipate this claim, Harris renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Harris with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[18B]	configuring the image processing platform to receive the data relating to the image and to conduct image processing, including:	See element 18A. To the extent it is found that Harris does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Harris does not anticipate this claim, Harris renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Harris with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results

Element	'252 Claim Recitation	Exemplary Citations in Reference
[18B.i]	operating on the data relating to the image to determine if the image contains one or more recognizable symbols; and	Harris discloses operating on the data relating to the image to determine if the image contains one or more recognizable symbols. Harris at 4:15-30 • "The advertisement 400 includes a bar code 405 therein. The bar code 405 is associated with the advertisement, and includes some additional information about the advertisement. For example, the bar code may include the web site address of the company preparing the advertisement, or appointment information about the advertisement, or a "vcf file" or the like which is an importable file with address information about the company or author sponsoring the advertisement. In operation, the user brings one of the clients, either the cellular telephone 310 or PDA 315, into range of the bar code 405. The client reads the bar code and decodes it as noted above. The decoded information can represent ASCII information, a compressed file such as a zip file, G code information, or any other compressed or non-compressed information." To the extent it is found that Harris does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Harris does not anticipate this claim, Harris renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Harris with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

Element	'252 Claim Recitation	Exemplary Citations in Reference
[18B.ii]	decoding the recognizable symbols to extract symbol information by analyzing the recognizable symbols according to type;	Harris discloses decoding the recognizable symbols to extract symbol information by analyzing the recognizable symbols according to type. Harris at FIG. 4 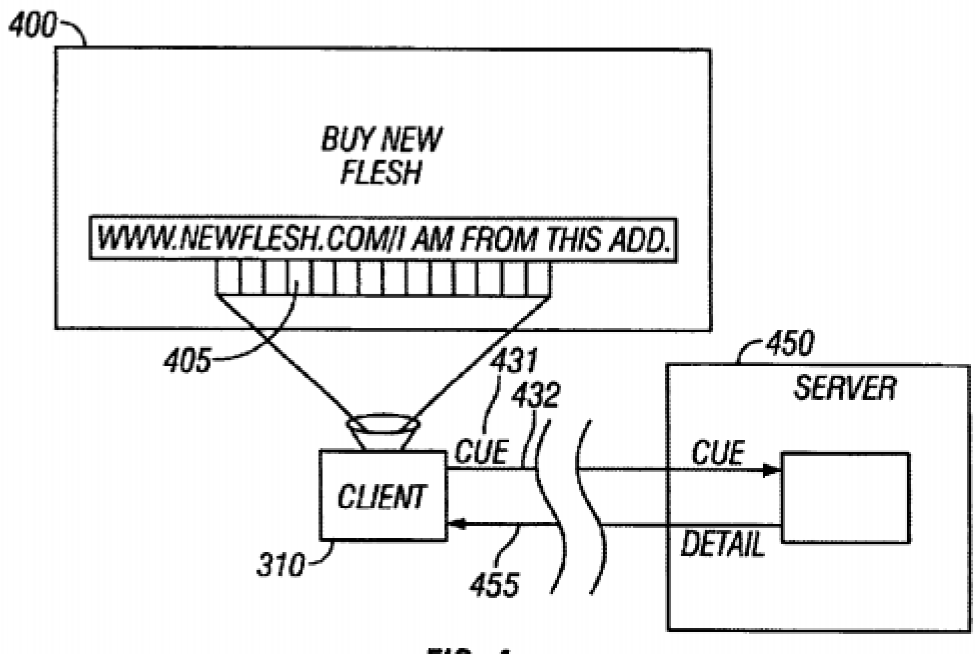 FIG. 4 Harris at 4:15-30 • "The advertisement 400 includes a bar code 405 therein. The bar code 405 is associated with the advertisement, and includes some additional information about the advertisement. For example, the bar code may include the web site address of the company preparing the advertisement, or appointment information about the advertisement, or a "vcf file" or the

Element	'252 Claim Recitation	Exemplary Citations in Reference
		like which is an importable file with address information about the company or author sponsoring the advertisement. In operation, the user brings one of the clients, either the cellular telephone 310 or PDA 315, into range of the bar code 405. The client reads the bar code and decodes it as noted above. The decoded information can represent ASCII information, a compressed file such as a zip file, G code information, or any other compressed or non-compressed information." To the extent it is found that Harris does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Harris does not anticipate this claim, Harris renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Harris with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[18C]	providing access to a distal server configured to use a database to identify pertinent information associated with the recognizable symbols based on the symbol information; and	Harris discloses providing access to a distal server configured to use a database to identify pertinent information associated with the recognizable symbols based on the symbol information. Harris at FIG. 4

Element	'252 Claim Recitation	Exemplary Citations in Reference
		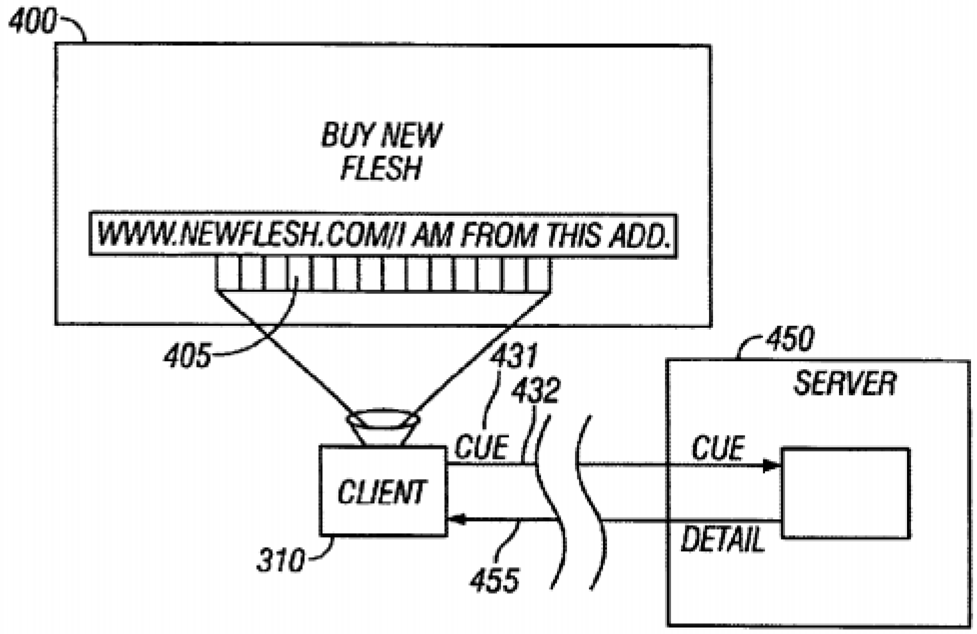 FIG. 4 Harris at 4:15-30 • "The advertisement 400 includes a bar code 405 therein. The bar code 405 is associated with the advertisement, and includes some additional information about the advertisement. For example, the bar code may include the web site address of the company preparing the advertisement, or appointment information about the advertisement, or a "vcf file" or the like which is an importable file with address information about the company or author sponsoring the advertisement. In operation, the user brings one of the clients, either the cellular telephone 310 or PDA 315, into range of the bar code 405. The client reads the bar code and decodes it as noted above. The decoded information can represent

Element	'252 Claim Recitation	Exemplary Citations in Reference
		ASCII information, a compressed file such as a zip file, G code information, or any other compressed or non-compressed information." Harris at 4:41-48 • "In this system, the information can represent a cue, e.g. a number, that is associated with a special function. The cue is used during a hot sync to access a publicly available information network, e.g. by direct connection or by the Internet. The cue 431 is sent over the Internet 432 to **the server 450**, and addresses more detailed information in a memory of the server. The server returns that information as 455, and the client receives the more detailed information." Harris at 4:55-58 • "During hot sync, the database returns the full text of the detailed information, e.g., "visit the website at http://www.pdascan.com/˜more to get a free gift." Any desired length or size of information can be returned." To the extent it is found that Harris does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Harris does not anticipate this claim, Harris renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Harris with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

Element	'252 Claim Recitation	Exemplary Citations in Reference
[18D]	allowing the mobile device to receive the pertinent information over a network.	Harris discloses allowing the mobile device to receive the pertinent information over a network. Harris at FIG. 4 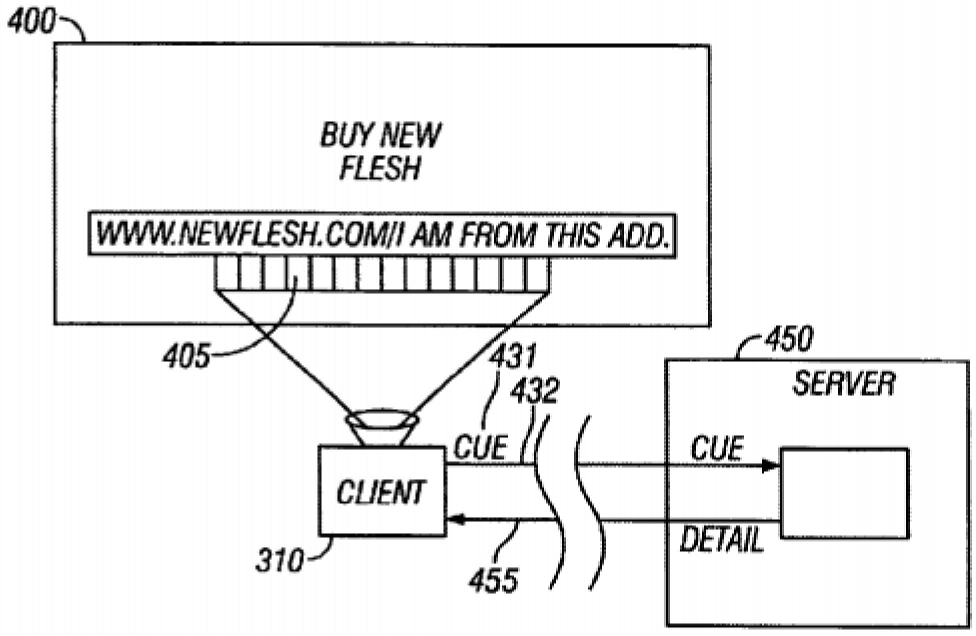 FIG. 4 Harris at 4:15-30 • "The advertisement 400 includes a bar code 405 therein. The bar code 405 is associated with the advertisement, and includes some additional information about the advertisement. For example, the bar code may include the web site address of the company preparing the advertisement, or appointment information about the advertisement, or a "vcf file" or the

Element	'252 Claim Recitation	Exemplary Citations in Reference
		like which is an importable file with address information about the company or author sponsoring the advertisement. In operation, the user brings one of the clients, either the cellular telephone 310 or PDA 315, into range of the bar code 405. The client reads the bar code and decodes it as noted above. The decoded information can represent ASCII information, a compressed file such as a zip file, G code information, or any other compressed or non-compressed information." Harris at 4:41-48 • "In this system, the information can represent a cue, e.g. a number, that is associated with a special function. The cue is used during a hot sync to access a publicly available information network, e.g. by direct connection or by the Internet. The cue 431 is sent over the Internet 432 to the server 450, and addresses more detailed information in a memory of the server. The server returns that information as 455, and the client receives the more detailed information." Harris at 4:55-58 • "During hot sync, the database returns the full text of the detailed information, e.g., "visit the website at http://www.pdascan.com/˜more to get a free gift." Any desired length or size of information can be returned." To the extent it is found that Harris does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Harris does not anticipate this claim, Harris renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Harris with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the

Element	'252 Claim Recitation	Exemplary Citations in Reference
		reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

2. Claim 26 (Non-asserted)

Claim	'252 Claim Recitation	Exemplary Citations in Reference
26	The method of claim 18, wherein the data relating to the image comprises the image.	Harris discloses the method of claim 18, wherein the data relating to the image comprises the image. Harris at FIG. 4

		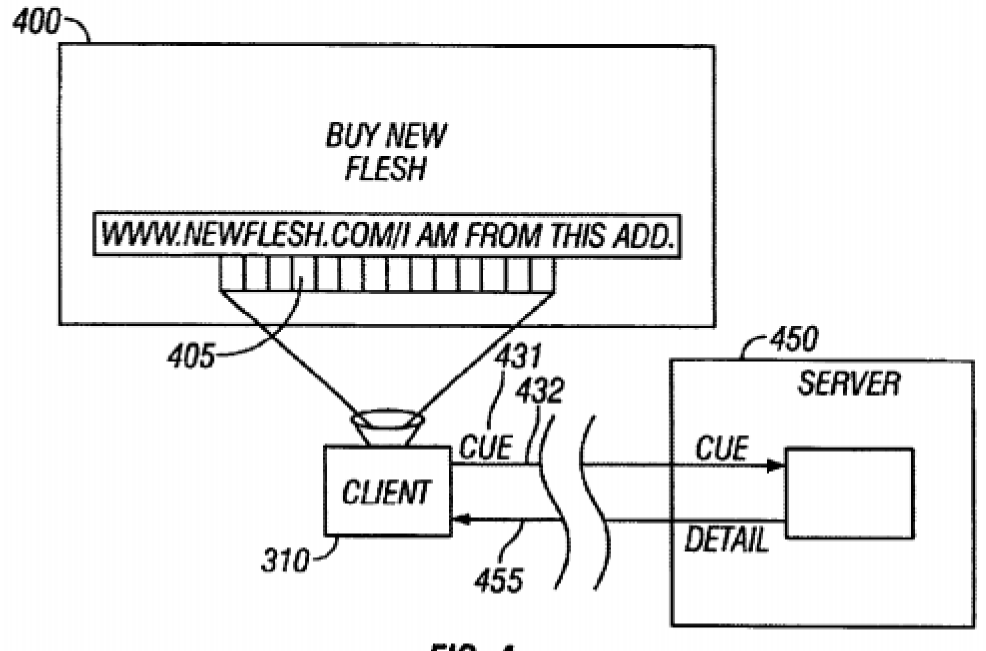 FIG. 4 Harris at 4:15-30 - "The advertisement 400 includes a bar code 405 therein. The bar code 405 is associated with the advertisement, and includes some additional information about the advertisement. For example, the bar code may include the web site address of the company preparing the advertisement, or appointment information about the advertisement, or a "vcf file" or the like which is an importable file with address information about the company or author sponsoring the advertisement. In operation, the user brings one of the clients, either the cellular telephone 310 or PDA 315, into range of the bar code 405. The client reads the bar code and decodes it as noted above. The decoded information can represent ASCII information, a compressed file such as a

		zip file, G code information, or any other compressed or non-compressed information." To the extent it is found that Harris does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Harris does not anticipate this claim, Harris renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Harris with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

3. Claim 27

Claim	'252 Claim Recitation	Exemplary Citations in Reference
27	The method of claim 26, wherein image processing further comprises compressing the image to form the data relating to the image.	Harris discloses the method of claim 26, wherein image processing further comprises compressing the image to form the data relating to the image. Harris at 4:15-30 • "The advertisement 400 includes a bar code 405 therein. The bar code 405 is associated with the advertisement, and includes some additional information about the advertisement. For example, the bar code may include the web site address of the company preparing the advertisement, or appointment information about the advertisement, or a "vcf file" or the like which is an importable file with address information about the company or author sponsoring the advertisement. In operation, the user brings one of the clients, either the cellular telephone 310 or PDA 315, into range of the bar code 405. The client

		reads the bar code and decodes it as noted above. The decoded information can represent ASCII information, a compressed file such as a zip file, G code information, or any other compressed or non-compressed information." To the extent it is found that Harris does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Harris does not anticipate this claim, Harris renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Harris with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

4. Claim 29

Claim	'252 Claim Recitation	Exemplary Citations in Reference
29	The method of claim 18, wherein image processing further includes converting at least a portion of the data relating to the image into a grayscale image.	To the extent it is found that Harris does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Harris does not anticipate this claim, Harris renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Harris with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to

		address similar known problems, according to known methods which would yield predictable results.

5. Claim 31

Claim	'252 Claim Recitation	Exemplary Citations in Reference
31	The method of claim 18, further comprising the distal server using database to identify the pertinent information based on image characteristics of the image derived by the image processing platform based on the data relating to the image.	Harris discloses the method of claim 18, further comprising the distal server using database to identify the pertinent information based on image characteristics of the image derived by the image processing platform based on the data relating to the image. Harris at FIG. 4 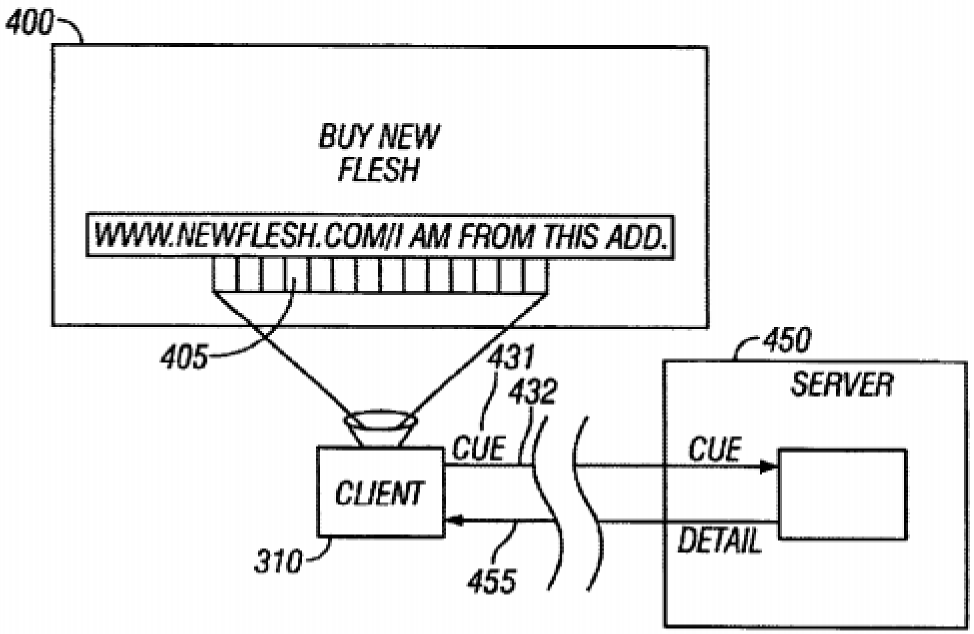 FIG. 4

		Harris at 4:41-48 • "In this system, the information can represent a cue, e.g. a number, that is associated with a special function. The cue is used during a hot sync to access a publicly available information network, e.g. by direct connection or by the Internet. The cue 431 is sent over the Internet 432 to the server 450, and addresses more detailed information in a memory of the server. The server returns that information as 455, and the client receives the more detailed information." Harris at 4:55-58 • "During hot sync, the database returns the full text of the detailed information, e.g., "visit the website at http://www.pdascan.com/˜more to get a free gift." Any desired length or size of information can be returned." To the extent it is found that Harris does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Harris does not anticipate this claim, Harris renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Harris with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

6.

Exhibit E-11

Bank of America Corporation ("BofA" or "Defendant") provides this chart subject to all reservations, objections, statements, and disclaimers set forth herein and in BofA's Final Invalidity Contentions Cover Pleading, as well as any amendment, supplement, or modification thereof, which are incorporated herein by reference in their entirety.

On information and belief, the application that matured into Ehrhart was filed on July 13, 2001, claiming the priority date of July 13, 2001. Ehrhart published at least as of April 20, 2004, and is available as prior art at least under 35 U.S.C. § 102(e).

As shown in the chart below, Ehrhart anticipates asserted claims 18, 27, 29, 31, of the '252 Patent. To the extent Ehrhart is found not to expressly disclose certain limitations in the asserted claims, such limitations are inherent. To the extent Ehrhart is found not to anticipate any asserted claims or claim elements of the '252 Patent, the reference nevertheless renders those claims or claim elements obvious under 35 U.S.C. § 103, either alone or in combination with other art identified in the cover pleading of BofA's Invalidity Contentions, for the reasons identified in the cover pleading. These Invalidity Contentions are not an admission by BofA that the accused products, including any current or past versions of the accused products, are covered by, or infringe the asserted claims, but are based instead on the recognition that if the claims are interpreted to be broad enough to encompass the accused products, the claims must also be construed to have that same scope when considering whether they are invalid.

The following chart is based on the Court's claim construction rulings (Dkt. Nos. 145, 153, 236), and to the extent not contradictory, Plaintiffs' apparent interpretations of the claims in Plaintiffs' "Preliminary" Final Infringement Contentions. BofA notes that in many instances, Plaintiffs' "Preliminary" Final Infringement Contentions ignore and/or contradict the Court's claim construction rulings. BofA does not accept the assumptions concerning the scope and meaning implicit in Plaintiffs' Preliminary Final Infringement Contentions, to the extent those assumptions are discernible, and reserves the right to challenge Plaintiffs' proposed (or implied) constructions. Furthermore, Plaintiffs' "Preliminary" Final Contentions fail to comply with the level of specificity required by SPR 2.1.2 and 2.1.3, and include improper reservation of rights to supplement with "additional ways in which BoA Accused Products infringe." Any alleged deficient discovery is of Plaintiffs own making. Most if not all of the deficiencies in Plaintiffs' Preliminary Final Contentions have been documented by Defendants through multiple letters since at least May 18, 2021 and up to Bank of America's most recent letter of August 19, 2023. Plaintiffs' shifting infringement contentions, deficiencies in its infringement contentions, and changing claim scope have prejudiced BofA from being able to reasonably prepare its defenses, including its invalidity contentions. BofA also reserves the right to revise and supplement these charts if and when Plaintiffs are permitted to provide revised or supplemental Infringement Contentions. Citations given in the chart below are merely representative of the respective elements to provide notice to Plaintiffs of BofA's invalidity theories and are not meant to be exhaustive.

1. Claim 18

Element	'252 Claim Recitation	Exemplary Citations in Reference
[18Pre]	A method for retrieving information from image processing, the method comprising:	Ehrhart discloses a method for retrieving information from image processing, the method comprising: "In another aspect, the present invention includes a method for acquiring an image of an object with an optical reader." Ehrhart at 2:62-64.
[18A]	providing a mobile device having a camera the mobile device configured to capture an image and configured to transmit data relating to the image an image processing platform;	Ehrhart discloses a mobile device having a camera, the mobile device configured to capture an image, and configured to transmit data relating to the image an image processing platform. *Id.* at [18Pre]. "The optical reader of the present invention automatically, or through manual Selection, determines whether a captured image is a color photographic image or, a color image that includes a graphical Symbol. Subsequently, the optical reader of the present invention processes the acquired imaging data in accordance with that determination. The optical reader of the present invention is operative to acquire and associate a plurality of acquired images. One aspect of the present invention is an optical reader. The optical reader includes a color imaging assembly for acquiring an image of an object, the color imaging assembly generating imaging data corresponding to the image. An image analysis circuit is coupled to the color imaging assembly. The image analysis circuit being configured to determine if the color imaging data includes at least one graphical Symbol. The image is classified as a

Element	'252 Claim Recitation	Exemplary Citations in Reference
		graphical Symbol, or the image is classified as a color photograph if the color imaging data does not include at least one graphical Symbol. A processing circuit is coupled to the image analysis circuit. The processing circuit is operative to process the imaging data based on the determination. In another aspect, the present invention includes an optical reader for capturing an image of an object. The optical reader includes a color imaging assembly for converting the image of the object into color digital data corresponding to the image." Ehrhart at 1:57-2:16. "In another aspect, the present invention includes an optical reader for capturing an image of an object. The optical reader includes a color imaging assembly for capturing the image as color imaging data. A classification circuit is coupled to the color imaging assembly, the classification circuit being configured to process at least a portion of the color imaging data to thereby Select one of a plurality of classifications, whereby the image is classified as a color photographic image, or as an image that includes at least one graphical Symbol. An automatic mode Selector is coupled to the classification circuit, the automatic mode Selector being configured to Select an optical reader mode in accordance with the Selected classification. A processor is coupled to the classification circuit, the processor being programmed to process the color imaging data in accordance with the optical reader mode Selected by the automatic mode Selector." Ehrhart at 2:26-42. "As embodied herein, and depicted in FIGS. 1A-1D, perspective views of the optical reader in accordance with various embodiments of the present invention are disclosed. FIG. 1A shows the underside of hand held wireless optical reader 10. FIG. 1B shows the top of the optical reader depicted in FIG. 1A. Optical reader 10 includes housing 100, antenna 102, window 104 and trigger 12. Window 104 accommodates illumination assembly 20 and imaging assembly 30. As shown in FIG. 1B, the top side of reader 10 includes function keys 14,

Element	'252 Claim Recitation	Exemplary Citations in Reference
		alphanumeric key pad 16, and display 60. In one embodiment, function keys 14 include an enter key and up and down cursor keys. FIG. 1C is also a hand held wireless optical reader 10. Reader 10 includes function keys 14, alphanumeric key pad 16, writing stylus 18, display 60, and signature block 62. Stylus 18 is employed by a user to write his signature in signature block 62. FIG. 1D shows yet another embodiment of optical reader 10 of the present invention. In this embodiment, reader 10 includes a gun-shaped housing 100. Display 60 and keypad 16 are disposed on a top portion of gun-shaped housing 100, whereas trigger 12 is disposed on the underside of the top portion of housing 100. Housing 100 also includes window 104 that accommodates illumination assembly 20 and imaging assembly 30. Wire 106 is disposed at the butt-end of housing 100. Wire 106 provides optical reader 10 with a hard wired communication link for external devices such as a host processor or other data collection devices." Ehrhart at 5:51-6:10. To the extent it is found that Ehrhart does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Ehrhart does not anticipate this claim, Ehrhart renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Ehrhart with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[18B]	configuring the image processing platform to receive the data relating to the image and to conduct image	Ehrhart discloses configuring the image processing platform to receive the data relating to the image and to conduct image processing.

Element	'252 Claim Recitation	Exemplary Citations in Reference
	processing, including:	*Id.* at [18Pre-A]. "In accordance with the invention, the present invention for an optical reader includes a color imaging assembly for acquiring color imaging data. An image analysis circuit determines if the acquired image includes at least one graphical symbol. A processing circuit processes the imaging data based on the determination of whether the image includes at least one graphical symbol. The present invention allows a user to read graphical symbols, such as bar codes, text, OCR characters or signatures using a color imager. The color optical reader of the present invention is configured to automatically determine whether a color image includes a graphical symbol, or is merely a color photographic image. The optical reader of the present invention also is operative to associate one acquired image with at least one subsequently acquired image." Ehrhart at 5:36-50. "In the embodiment depicted in FIG. 2, ASIC 44 is implemented using a programmable logic array (PLA) device. In a similar embodiment, ASIC 44 is implemented using a field programmable gate array (FPGA) device. ASIC 44 is tasked with controlling the image acquisition process, and the storage of image data. As part of the image acquisition process, ASIC 44 performs various timing and control functions including control of light source 24, control of color imager 34, and control of external interface 56. It will be apparent to those of ordinary skill in the pertinent art that modifications and variations can be made to processor 40 of the present invention depending on the cost, availability, and performance of off-the-shelf microprocessors, and the type of color imager used. In one embodiment, microprocessor 42 and ASIC 44 are replaced by a single microprocessor 40. In one embodiment, microprocessor 40 is implemented using a single RISC processor. In yet another embodiment, microprocessor 40 is implemented using a RISC and DSP hybrid processor." Ehrhart at 7:32-50.

Element	'252 Claim Recitation	Exemplary Citations in Reference
		To the extent it is found that Ehrhart does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Ehrhart does not anticipate this claim, Ehrhart renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Ehrhart with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results
[18B.i]	operating on the data relating to the image to determine if the image contains one or more recognizable symbols; and	Ehrhart discloses operating on the data relating to the image to determine if the image contains one or more recognizable symbols. *Id.* at [18Pre-B]. "In accordance with the invention, the present invention for an optical reader includes a color imaging assembly for acquiring color imaging data. An image analysis circuit determines if the acquired image includes at least one graphical symbol. A processing circuit processes the imaging data based on the determination of whether the image includes at least one graphical symbol. The present invention allows a user to read graphical symbols, such as bar codes, text, OCR characters or signatures using a color imager. The color optical reader of the present invention is configured to automatically determine whether a color image includes a graphical symbol, or is merely a color photographic image. The optical reader of the present invention also is operative to associate one acquired image with at least one subsequently acquired image." Ehrhart at 5:36-50. "In step 1004, processor 40 identifies the page features. Processor 40 analyzes the page and divides it into blank and non-blank portions. The non-blank portions are

Element	'252 Claim Recitation	Exemplary Citations in Reference
		analyzed to distinguish text regions from non-text regions. After determining the layout of the page, processor 40 uses black-to-white transitions to determine degrees of skew. In step 1008, horizontal white spaces are identified to separate lines of text. In step 1010, vertical white spaces are identified within each line of text to thereby separate individual words and characters from each other. In Step 1014, a character recognition algorithm is used in an attempt to recognize each individual character." Ehrhart at 12:55-66.

		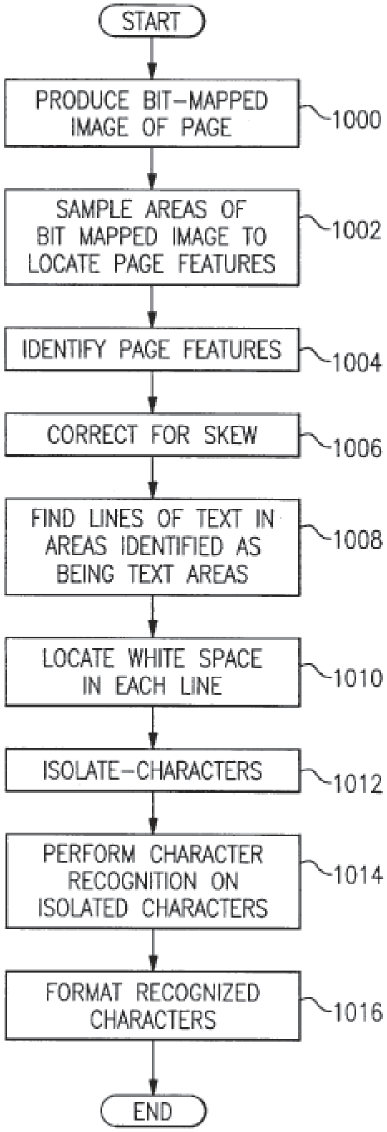START PRODUCE BIT-MAPPED IMAGE OF PAGE 1000 SAMPLE AREAS OF BIT MAPPED IMAGE TO LOCATE PAGE FEATURES 1002 IDENTIFY PAGE FEATURES 1004 CORRECT FOR SKEW 1006 FIND LINES OF TEXT IN AREAS IDENTIFIED AS BEING TEXT AREAS 1008 LOCATE WHITE SPACE IN EACH LINE 1010 ISOLATE-CHARACTERS 1012 PERFORM CHARACTER RECOGNITION ON ISOLATED CHARACTERS 1014 FORMAT RECOGNIZED CHARACTERS 1016 END FIG.10

Element	'252 Claim Recitation	Exemplary Citations in Reference
		Ehrhart at Fig. 10.

		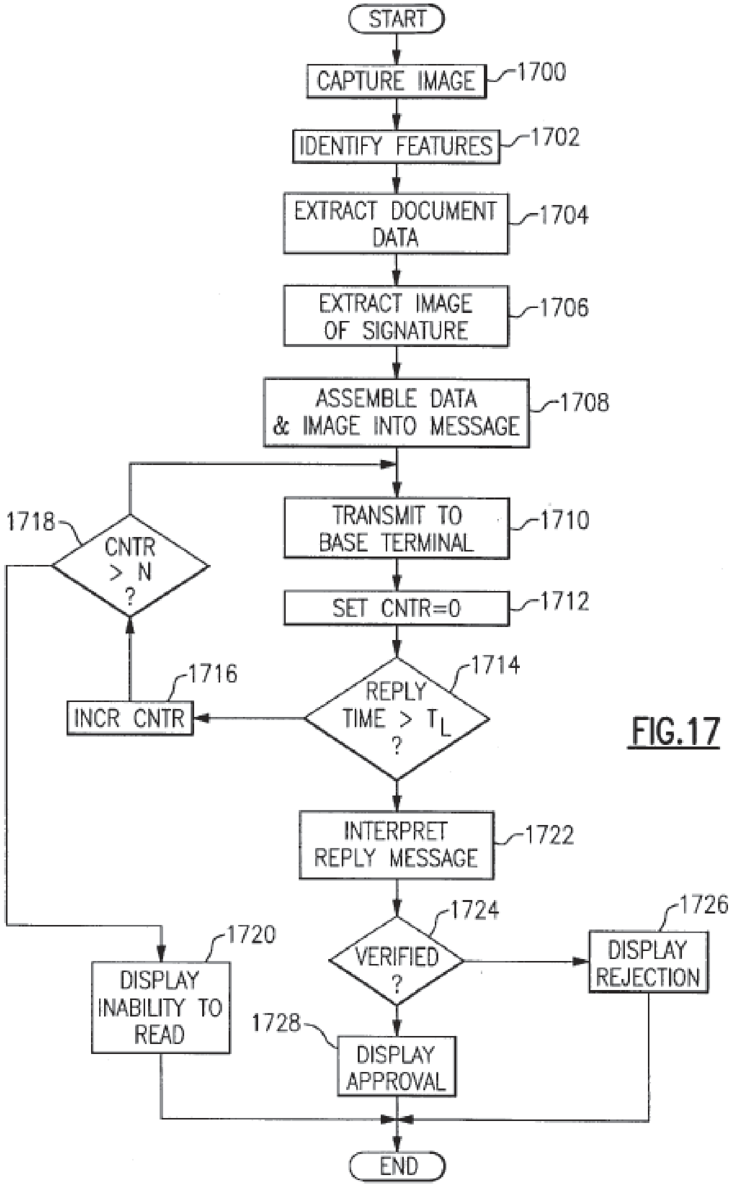START CAPTURE IMAGE 1700 IDENTIFY FEATURES 1702 EXTRACT DOCUMENT DATA 1704 EXTRACT IMAGE OF SIGNATURE 1706 ASSEMBLE DATA & IMAGE INTO MESSAGE 1708 TRANSMIT TO BASE TERMINAL 1710 SET CNTR=0 1712 REPLY TIME > T_L ? 1714 INCR CNTR 1716 CNTR > N ? 1718 INTERPRET REPLY MESSAGE 1722 VERIFIED ? 1724 DISPLAY REJECTION 1726 DISPLAY INABILITY TO READ 1720 DISPLAY APPROVAL 1728 END FIG.17

Element	'252 Claim Recitation	Exemplary Citations in Reference
		Ehrhart at Fig. 17. To the extent it is found that Ehrhart does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Ehrhart does not anticipate this claim, Ehrhart renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Ehrhart with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[18B.ii]	decoding the recognizable symbols to extract symbol information by analyzing the recognizable symbols according to type;	Ehrhart discloses decoding the recognizable symbols to extract symbol information by analyzing the recognizable symbols according to type. *Id.* at [18Pre-B]. "As embodied herein and depicted in FIG. 13, an example of image association in accordance with the present invention is disclosed. One or ordinary skill in the art will recognize that associated images 1300 can be disposed on paper, displayed electronically on display 60, or displayed electronically sing other electronic means, such as a computer monitor. In this example, the first image captured is color photograph 1302 which shows a damaged parcel. The second image captured is bar code 1304 affixed to the side of the damaged parcel. Processor 40 decodes bar code 1304 and associates decoded bar code data 1306 with color photograph 1302. In this example, the user elected to associate a third image, signature 1308. Thus, personnel viewing record 1300 may reasonably conclude that a damaged parcel was delivered

Element	'252 Claim Recitation	Exemplary Citations in Reference
		to Company XYZ, and that the person signing for the parcel delivery was someone named John W. Smith." Ehrhart at 13:27-43. "In the Automatic imaging mode, processor 40 is programmed to analyze the color imaging data to determine if an acquired image includes a graphical symbol or is merely a color photographic image. If it makes the determination that the color image includes a graphical symbol, it further analyzes the acquired image and classifies it as a bar code, OCR symbol, text, or a signature. Based on the classification, optical reader 10 jumps to the appropriate routine in EROM 48. The semi-automatic mode is similar. Thus, in the automatic or semi-automatic modes, the bar code scanning mode, the OCR/text mode, the signature capture mode, the color photography mode, and the association mode are controlled by the application program, not by the user." Ehrhart at 8:36-49. "In step 1004, processor 40 identifies the page features. Processor 40 analyzes the page and divides it into blank and non-blank portions. The non-blank portions are analyzed to distinguish text regions from non-text regions. After determining the layout of the page, processor 40 uses black-to-white transitions to determine degrees of skew. In step 1008, horizontal white spaces are identified to separate lines of text. In step 1010, vertical white spaces are identified within each line of text to thereby separate individual words and characters from each other. In Step 1014, a character recognition algorithm is used in an attempt to recognize each individual character." Ehrhart at 12:55-66. "As embodied herein and depicted in FIG. 11, a flow chart showing a method for performing OCR in accordance with yet another embodiment of the present invention is disclosed. In step 1100, reader 10 produces a bit-mapped image of the page. Subsequently, processor 40 finds lines of text in the image, locates the white spaces in each line, and isolates the characters. In step 1108, processor 40 performs character recognition, either OCR-A or OCR-B, as desired. The decoded characters are stored in memory." Ehrhart at 13:1-9.

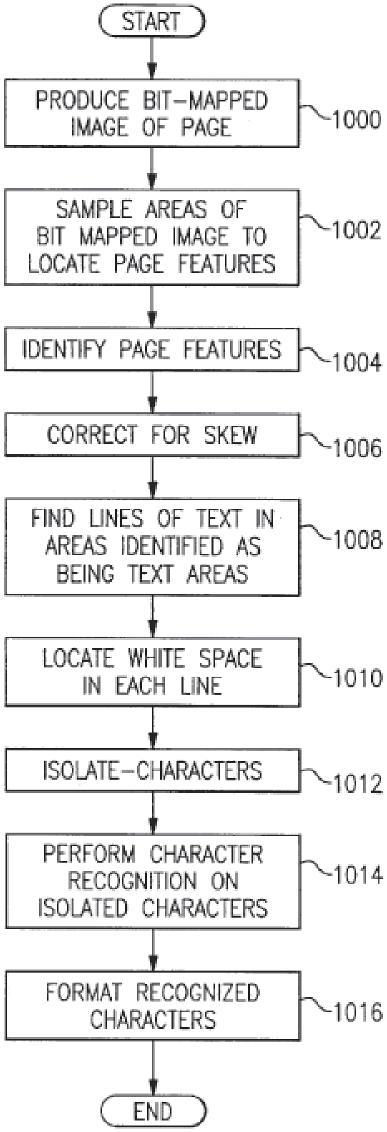
START
PRODUCE BIT-MAPPED IMAGE OF PAGE
1000
SAMPLE AREAS OF BIT MAPPED IMAGE TO LOCATE PAGE FEATURES
1002
IDENTIFY PAGE FEATURES
1004
CORRECT FOR SKEW
1006
FIND LINES OF TEXT IN AREAS IDENTIFIED AS BEING TEXT AREAS
1008
LOCATE WHITE SPACE IN EACH LINE
1010
ISOLATE-CHARACTERS
1012
PERFORM CHARACTER RECOGNITION ON ISOLATED CHARACTERS
1014
FORMAT RECOGNIZED CHARACTERS
1016
END

FIG.10

Element	'252 Claim Recitation	Exemplary Citations in Reference
		Ehrhart at Fig. 10.

		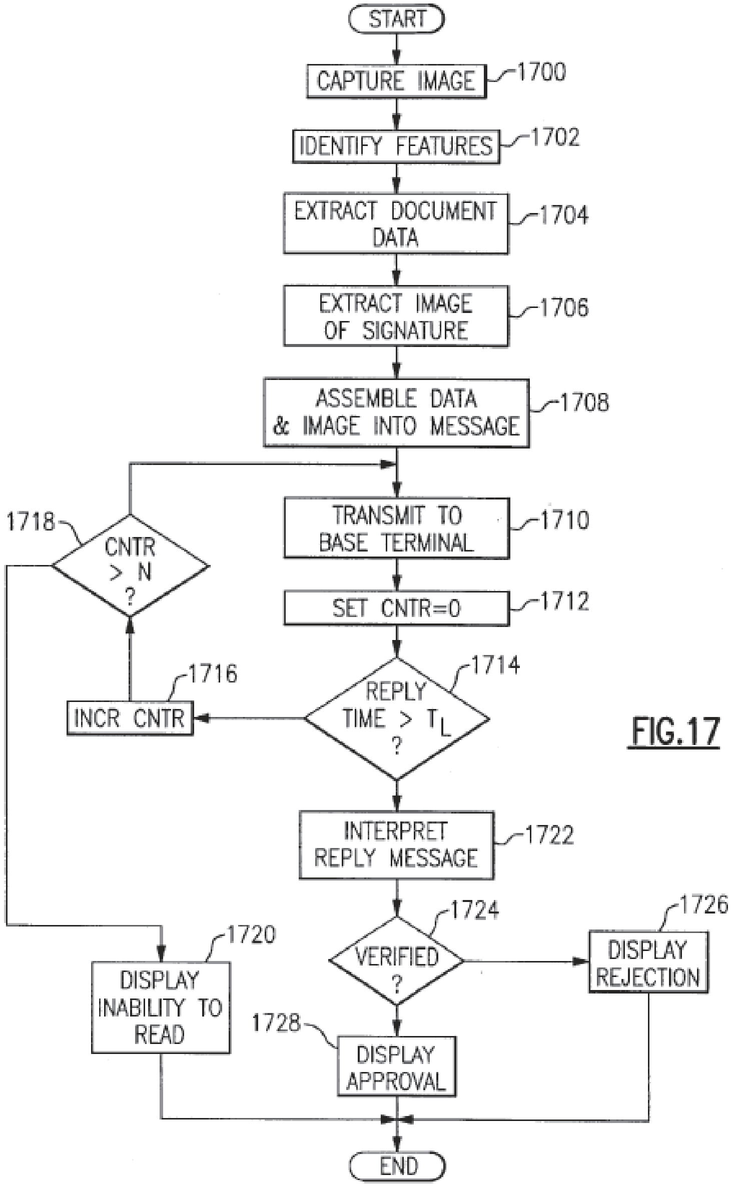 START CAPTURE IMAGE 1700 IDENTIFY FEATURES 1702 EXTRACT DOCUMENT DATA 1704 EXTRACT IMAGE OF SIGNATURE 1706 ASSEMBLE DATA & IMAGE INTO MESSAGE 1708 1718 CNTR > N ? TRANSMIT TO BASE TERMINAL 1710 SET CNTR=0 1712 1716 1714 INCR CNTR REPLY TIME > T_L ? INTERPRET REPLY MESSAGE 1722 1724 1726 1720 VERIFIED ? DISPLAY REJECTION DISPLAY INABILITY TO READ 1728 DISPLAY APPROVAL END FIG.17

Element	'252 Claim Recitation	Exemplary Citations in Reference
		Ehrhart at Fig. 17. To the extent it is found that Ehrhart does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Ehrhart does not anticipate this claim, Ehrhart renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Ehrhart with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[18C]	providing access to a distal server configured to use a database to identify pertinent information associated with the recognizable symbols based on the symbol information; and	Ehrhart discloses providing access to a distal server configured to use a database to identify pertinent information associated with the recognizable symbols based on the symbol information. *Id.* at [18Pre-B]. "LAN 1820 includes server 1822, computer 1824, database 1826, and a plurality of optical readers 10. Database 1826 is used to store associated images along with other data fields. For example, it would be rather useful to store additional information with the associated images shown in FIG. 13. One may want to associate the delivery means, route, driver, and other related information for subsequent analysis. Network 1810 allows reader 10, PAN 1850, and wireless system 1400 a way to store such data in database 1826. System analysts can access this information via personal computer 1802 connected to network 1810. In one embodiment, LAN 1820 includes an Internet website. In this embodiment, users are authenticated before gaining access to database 1826." Ehrhart at 16:1-12.

Element	'252 Claim Recitation	Exemplary Citations in Reference
		To the extent it is found that Ehrhart does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Ehrhart does not anticipate this claim, Ehrhart renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Ehrhart with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[18D]	allowing the mobile device to receive the pertinent information over a network.	Ehrhart discloses allowing the mobile device to receive the pertinent information over a network. *Id.* at [18Pre-C]. "As embodied herein, and depicted in FIGS. 1A-1D, perspective views of the optical reader in accordance with various embodiments of the present invention are disclosed. FIG. 1A shows the underside of hand held wireless optical reader 10. FIG. 1B shows the top of the optical reader depicted in FIG. 1A. Optical reader 10 includes housing 100, antenna 102, window 104 and trigger 12. Window 104 accommodates illumination assembly 20 and imaging assembly 30. As shown in FIG. 1B, the top side of reader 10 includes function keys 14, alphanumeric key pad 16, and display 60. In one embodiment, function keys 14 include an enter key and up and down cursor keys. FIG. 1C is also a hand held wireless optical reader 10. Reader 10 includes function keys 14, alphanumeric key pad 16, writing stylus 18, display 60, and signature block 62. Stylus 18 is employed by a user to write his signature in signature block 62. FIG.

Element	'252 Claim Recitation	Exemplary Citations in Reference
		1D shows yet another embodiment of optical reader 10 of the present invention. In this embodiment, reader 10 includes a gun-shaped housing 100. Display 60 and keypad 16 are disposed on a top portion of gun-shaped housing 100, whereas trigger 12 is disposed on the underside of the top portion of housing 100. Housing 100 also includes window 104 that accommodates illumination assembly 20 and imaging assembly 30. Wire 106 is disposed at the butt-end of housing 100. Wire 106 provides optical reader 10 with a hard wired communication link for external devices such as a host processor or other data collection devices." Ehrhart at 5:51-6:10. To the extent it is found that Ehrhart does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Ehrhart does not anticipate this claim, Ehrhart renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Ehrhart with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

2. Claim 26

Claim	'252 Claim Recitation	Exemplary Citations in Reference
26	The method of claim 18, wherein the data relating to the image comprises the image.	Ehrhart discloses the method of claim 18, wherein the data relating to the image comprises the image.

		"In the Automatic imaging mode, processor 40 is programmed to analyze the color imaging data to determine if an acquired image includes a graphical symbol or is merely a color photographic image. If it makes the determination that the color image includes a graphical symbol, it further analyzes the acquired image and classifies it as a bar code, OCR symbol, text, or a signature. Based on the classification, optical reader 10 jumps to the appropriate routine in EROM 48. The semi-automatic mode is similar. Thus, in the automatic or semi-automatic modes, the bar code scanning mode, the OCR/text mode, the signature capture mode, the color photography mode, and the association mode are controlled by the application program, not by the user." Ehrhart at 8:36-49. "If processor 40 does not locate a bar code symbol it moves onto step 420 and looks for OCR-A or OCR-B characters. If it finds these characters it performs optical character recognition in step 422. If it does not, processor evaluates the image for the presence of text. If text is located, the image is cropped, and the text is compressed and stored in steps 428 and 430. If the image does not include text, processor 40 evaluates the image for the presence of a signature. If one is present, the image is cropped, and the data is compressed and stored in steps 428 and 430. In another embodiment, optical reader 10 is networked, and processor 40 communicates with remote network resources to provide signature verification services. If processor 40 cannot detect a bar code symbol, OCR symbols, text, or a signature, the user is asked in step 432 if he desires to store the image. If he does, the color photographic image is stored in memory in step 434." Ehrhart at 10:5-20. "As embodied herein and depicted in FIG. 11, a flow chart showing a method for performing OCR in accordance with yet another embodiment of the present invention is disclosed. In step 1100, reader 10 produces a bit-mapped image of the page. Subsequently, processor 40 finds lines of text in the image, locates the white spaces in each line, and isolates the characters. In step 1108, processor 40 performs character recognition, either OCR-A or OCR-B, as desired. The decoded characters are stored in memory." Ehrhart at 13:1-9.

		To the extent it is found that Ehrhart does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Ehrhart does not anticipate this claim, Ehrhart renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified in the cover pleading of BofA's Preliminary Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Ehrhart with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

3. Claim 27

Claim	'252 Claim Recitation	Exemplary Citations in Reference
27	The method of claim 26, wherein image processing further comprises compressing the image to form the data relating to the image.	Ehrhart discloses the method of claim 26, wherein image processing further comprises compressing the image to form the data relating to the image. "If processor 40 does not locate a bar code symbol it moves onto step 420 and looks for OCR-A or OCR-B characters. If it finds these characters it performs optical character recognition in step 422. If it does not, processor evaluates the image for the presence of text. If text is located, the image is cropped, and the text is compressed and stored in steps 428 and 430. If the image does not include text, processor 40 evaluates the image for the presence of a signature. If one is present, the image is cropped, and the data is compressed and stored in steps 428 and 430. In another embodiment, optical reader 10 is networked, and processor 40 communicates with remote network resources to provide signature verification services. If processor 40 cannot detect a bar code symbol, OCR symbols, text, or a

		signature, the user is asked in step 432 if he desires to store the image. If he does, the color photographic image is stored in memory in step 434." Ehrhart at 10:5-20. To the extent it is found that Ehrhart does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Ehrhart does not anticipate this claim, Ehrhart renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Ehrhart with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

4. Claim 29

Claim	'252 Claim Recitation	Exemplary Citations in Reference
29	The method of claim 18, wherein image processing further includes converting at least a portion of the data relating to the image into a grayscale image.	Ehrhart discloses converting at least a portion of the data relating to the image into a grayscale image. *Id.* at 18. "In another embodiment, processor 40 creates a gray-Scale image to determine whether the image includes a graphical Symbol." Ehrhart at 9:54-56. "As embodied herein and depicted in FIG. 10, a flow chart showing a method for reading text in accordance with yet another embodiment of the present invention is disclosed. This routine can be accessed in a number of ways as described above. In step 1000, a bit-map image of the page is produced. In step 1002, the bit mapped

		image is sampled. In one embodiment, this is performed by analyzing every Nth scan line of the bit mapped image. The value of integer N is dependent on the resolution of the scanned image. In one embodiment the image is sampled every {fraction (1/40)}th of an inch. This provides sufficient resolution to locate and classify the various regions on the page. By sampling every {fraction (1/40)}th of an inch instead of every scan line, the processing and memory requirements of reader 10 are substantially reduced. In step 1004, processor 40 identifies the page features. Processor 40 analyzes the page and divides it into blank and non-blank portions. The non-blank portions are analyzed to distinguish text regions from non-text regions. After determining the layout of the page, processor 40 uses black-to-white transitions to determine degrees of skew. In step 1008, horizontal white spaces are identified to separate lines of text. In step 1010, vertical white spaces are identified within each line of text to thereby separate individual words and characters from each other. In step 1014, a character recognition algorithm is used in an attempt to recognize each individual character. Finally, in step 1016, processor 40 formats the recovered text before storing the text in memory." Ehrhart at 12:41-67. To the extent it is found that Ehrhart does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Ehrhart does not anticipate this claim, Ehrhart renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Ehrhart with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

5. Claim 31

Claim	'252 Claim Recitation	Exemplary Citations in Reference
31	The method of claim 18, further comprising the distal server using database to identify the pertinent information based on image characteristics of the image derived by the image processing platform based on the data relating to the image.	Ehrhart discloses the method of claim 18, further comprising the distal server using database to identify the pertinent information based on image characteristics of the image derived by the image processing platform based on the data relating to the image. "LAN 1820 includes server 1822, computer 1824, database 1826, and a plurality of optical readers 10. Database 1826 is used to store associated images along with other data fields. For example, it would be rather useful to store additional information with the associated images shown in FIG. 13. One may want to associate the delivery means, route, driver, and other related information for subsequent analysis. Network 1810 allows reader 10, PAN 1850, and wireless system 1400 a way to store such data in database 1826. System analysts can access this information via personal computer 1802 connected to network 1810. In one embodiment, LAN 1820 includes an Internet website. In this embodiment, users are authenticated before gaining access to database 1826." Ehrhart at 16:1-12. To the extent it is found that Ehrhart does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Ehrhart does not anticipate this claim, Ehrhart renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Ehrhart with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

Exhibit E-23

Bank of America Corporation (“BofA” or “Defendant”) provides this chart subject to all reservations, objections, statements, and disclaimers set forth herein and in BofA’s Final Invalidity Contentions Cover Pleading, as well as any amendment, supplement, or modification thereof, which are incorporated herein by reference in their entirety. A statement of reasons and good cause for each supplement in this chart is identified in the cover pleading.

QBIC was known or used by others in the United States and publicly used or on sale in the United States by November 5, 2000. *See generally* W. Niblack et al., The QBIC Project: Querying Images By Content Using Color, Texture, and Shape, SPIE Vol. 1908 (1993), pp. 173–187 (IBM0002390–404); M. Flickner et al., Query by Image and Video Content: The QBIC System (“QBIC”), IEEE Sept. 1995, pp. 23–32 (IBM0000893–902); W. Niblack et al, Updates to the QBIC System, SPIE Vol. 3312 (1997), pp. 150–161 (IBM 0002419–430); and Query By Image Content QBIC Demonstration Program, IBM Research, Sept. 1998 (IBM000747). QBIC is available as prior art at least under 35 U.S.C. § 102(a), (b), and (g), having been known or used by others in the United States and publicly used or on sale in the United States by at least 1995 and certainly by November 5, 2000, including, for example, at http://libra.uc.davis.edu through the Art History Department at U.C. Davis; at www..thinker.org/imagebase/indox-2 .htm3 through the Fine Arts Museum of San Francisco; through demos at http://wwwqbic.almaden.ibm.com; and through IBM’s commercial products, including in IBM’s Ultimedia Manager, the DB2 series of products such as the Image Extender in DB2 Universal Database, and the Enterprise Information Portal. *See generally* Deposition of IBM Corporate Representative Myron Flickner, *Nantworks, LLC and Nant Holdings IP, LLC v. Bank of America Corporation and Bank of America, N.A.*, 2:20-cv-07872-GW-PVC (C.D.C.A. 2020) (“Flickner 8/29/23 Depo”) at e.g., 32–33, 39–40, 44, 62–63, 66–67,72–74, 103–107, 110–112, 127, 163–174, 183–189, 191–195, 215, 219–220, 241, 307–309, 311–317, Exs. 3–13; *see also* Deposition of IBM Corporate Representative Myron Flickner, *Network-1 Technologies, Inc., v. Google LLC and Youtube LLC*, 1:14-cv-09558-PGG (S.D.N.Y. 2019) (“Flickner Depo”), including, e.g., at 13, 14, 15–18, 29–30, 34, 36, 37, 43–46, 48, 58, 60, 61–62, 67–69, 72, 74–76, 84–86, 88–89, 98–100, 102, 106, 113, 125, 135, 152–154, 160–161, 170–172, 180–181, 194, 198–201, 203–205, 215–216, 220–221, and Exs. 1–16.

As shown in the chart below, QBIC anticipates asserted claims 18, 27, 29, and 31 of the ’252 Patent. The disclosures/exemplary citations for each element incorporate the disclosures/exemplary citations of the proceeding limitations. To the extent QBIC is found not to expressly disclose certain limitations in the asserted claims, such limitations are inherent. To the extent QBIC is found not to anticipate any asserted claims or claim elements of the ’252 Patent, the reference nevertheless renders those claims or claim elements obvious under 35 U.S.C. § 103, either alone or in combination with other art identified in the cited references and/or or in the cover pleading of BofA’s Final Invalidity Contentions, for the reasons identified in the cover pleading. These Invalidity Contentions are not an admission by BofA that the accused products, including any current or past versions of the accused products, are covered by, or infringe the asserted claims, but are based instead on the recognition that if the claims are interpreted to be broad enough to encompass the accused products, the claims must also be construed to have that same scope when considering whether they are invalid. The following chart incorporates by reference the analysis from Exhibit E-21 for each limitation.

The following chart is based on the Court's claim construction rulings (Dkt. Nos. 145, 153, 236), and to the extent not contradictory, Plaintiffs' apparent interpretations of the claims in Plaintiffs' "Preliminary" Final Infringement Contentions. BofA notes that in many instances, Plaintiffs' "Preliminary" Final Infringement Contentions ignore and/or contradict the Court's claim construction rulings. BofA does not accept the assumptions concerning the scope and meaning implicit in Plaintiffs' Preliminary Final Infringement Contentions, to the extent those assumptions are discernible, and reserves the right to challenge Plaintiffs' proposed (or implied) constructions. Furthermore, Plaintiffs' "Preliminary" Final Contentions fail to comply with the level of specificity required by SPR 2.1.2 and 2.1.3, and include improper reservation of rights to supplement with "additional ways in which BoA Accused Products infringe." Any alleged deficient discovery is of Plaintiffs own making. Most if not all of the deficiencies in Plaintiffs' Preliminary Final Contentions have been documented by Defendants through multiple letters since at least May 18, 2021 and up to Bank of America's most recent letter of August 19, 2023. Plaintiffs' shifting infringement contentions, deficiencies in its infringement contentions, and changing claim scope have prejudiced BofA from being able to reasonably prepare its defenses, including its invalidity contentions. BofA also reserves the right to revise and supplement these charts if and when Plaintiffs are permitted to provide revised or supplemental Infringement Contentions. Citations given in the chart below are merely representative of the respective elements to provide notice to Plaintiffs of BofA's invalidity theories and are not meant to be exhaustive.

1. Claim 18

Element	'252 Claim Recitation	Exemplary Citations in Reference
[18Pre]	A method for retrieving information from image processing, the method comprising:	QBIC discloses a method for retrieving information from image processing, the method comprising: The QBIC Project: Querying Images by Content Using Color, Texture, and Shape (IBM 0002390–404) at 2390 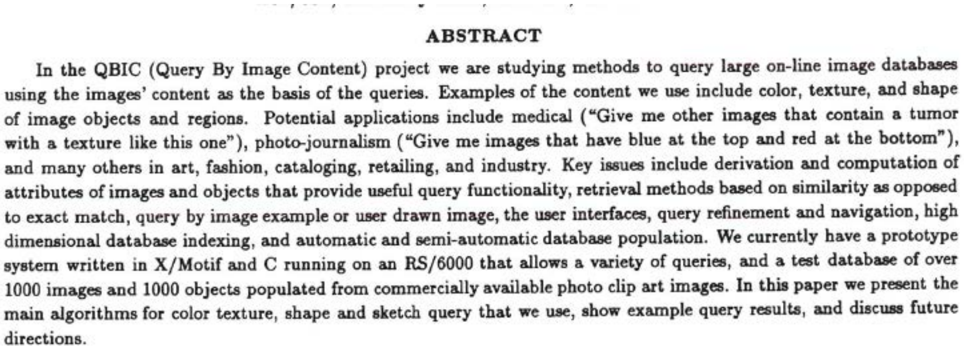 ABSTRACT In the QBIC (Query By Image Content) project we are studying methods to query large on-line image databases using the images' content as the basis of the queries. Examples of the content we use include color, texture, and shape of image objects and regions. Potential applications include medical ("Give me other images that contain a tumor with a texture like this one"), photo-journalism ("Give me images that have blue at the top and red at the bottom"), and many others in art, fashion, cataloging, retailing, and industry. Key issues include derivation and computation of attributes of images and objects that provide useful query functionality, retrieval methods based on similarity as opposed to exact match, query by image example or user drawn image, the user interfaces, query refinement and navigation, high dimensional database indexing, and automatic and semi-automatic database population. We currently have a prototype system written in X/Motif and C running on an RS/6000 that allows a variety of queries, and a test database of over 1000 images and 1000 objects populated from commercially available photo clip art images. In this paper we present the main algorithms for color texture, shape and sketch query that we use, show example query results, and discuss future directions. QBIC Demo (IBM000747)

Element	'252 Claim Recitation	Exemplary Citations in Reference
		 QBIC$^{(TM)}$ -- IBM's Query By Image Content QBIC allows you to make queries of image databases large and small based on visual characteristics -- properties such as color percentages, color layout and textures occurring in the images. The queries match image colors, textures and their positions without having to describe them with in words. QBIC's content-based queries can also be combined with keywords to provide a powerful multimedia search and retrieval tool. As image collections grow larger and become more prolific, QBIC is the best tool to more efficiently search and navigate through them. In this demo, we use QBIC to index about 470 color images based on their color and texture information. You can evaluate the QBIC Search capabilty by following the link located at left called IBM Image Search Demo. For license and business related questins, please contact Ted Loewenberg by phone 408-927-1202, or by email QBIC Demo (IBM000747) 

Element	'252 Claim Recitation	Exemplary Citations in Reference
		Flickner Depo at 13 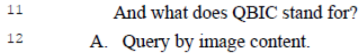 And what does QBIC stand for? A. Query by image content. Flickner Depo at 40 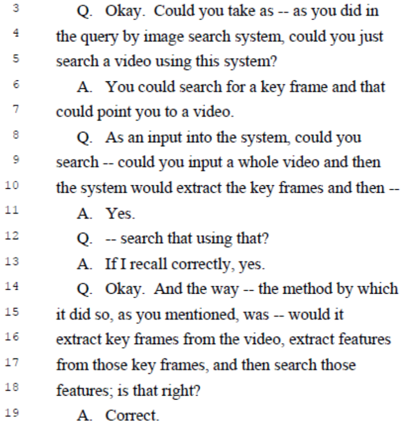 Q. Okay. Could you take as -- as you did in the query by image search system, could you just search a video using this system? A. You could search for a key frame and that could point you to a video. Q. As an input into the system, could you search -- could you input a whole video and then the system would extract the key frames and then -- A. Yes. Q. -- search that using that? A. If I recall correctly, yes. Q. Okay. And the way -- the method by which it did so, as you mentioned, was -- would it extract key frames from the video, extract features from those key frames, and then search those features; is that right? A. Correct. Flickner Depo at 76 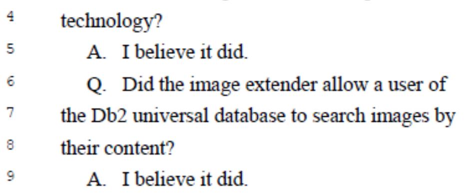 Q. Did the image extender incorporate QBIC technology? A. I believe it did. Q. Did the image extender allow a user of the Db2 universal database to search images by their content? A. I believe it did. Flickner Depo at 89-90

Element	'252 Claim Recitation	Exemplary Citations in Reference
		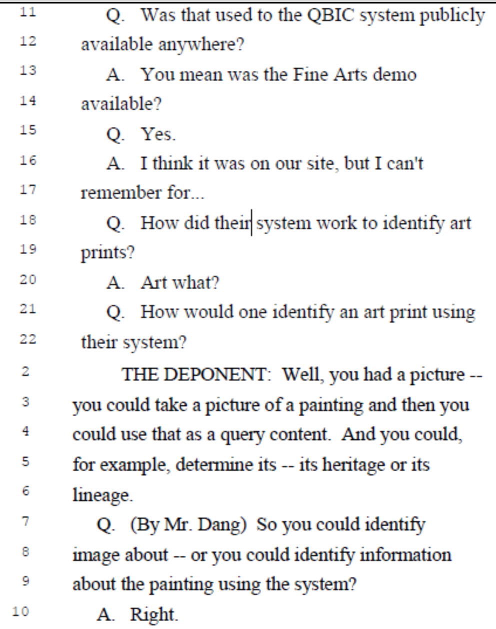 Q. Was that used to the QBIC system publicly available anywhere? A. You mean was the Fine Arts demo available? Q. Yes. A. I think it was on our site, but I can't remember for... Q. How did their system work to identify art prints? A. Art what? Q. How would one identify an art print using their system? THE DEPONENT: Well, you had a picture -- you could take a picture of a painting and then you could use that as a query content. And you could, for example, determine its -- its heritage or its lineage. Q. (By Mr. Dang) So you could identify image about -- or you could identify information about the painting using the system? A. Right. Flickner Depo at, e.g., 24, 41-42, 62-63 (user could take a picture or video and submit it to the QBIC system).

Element	'252 Claim Recitation	Exemplary Citations in Reference
		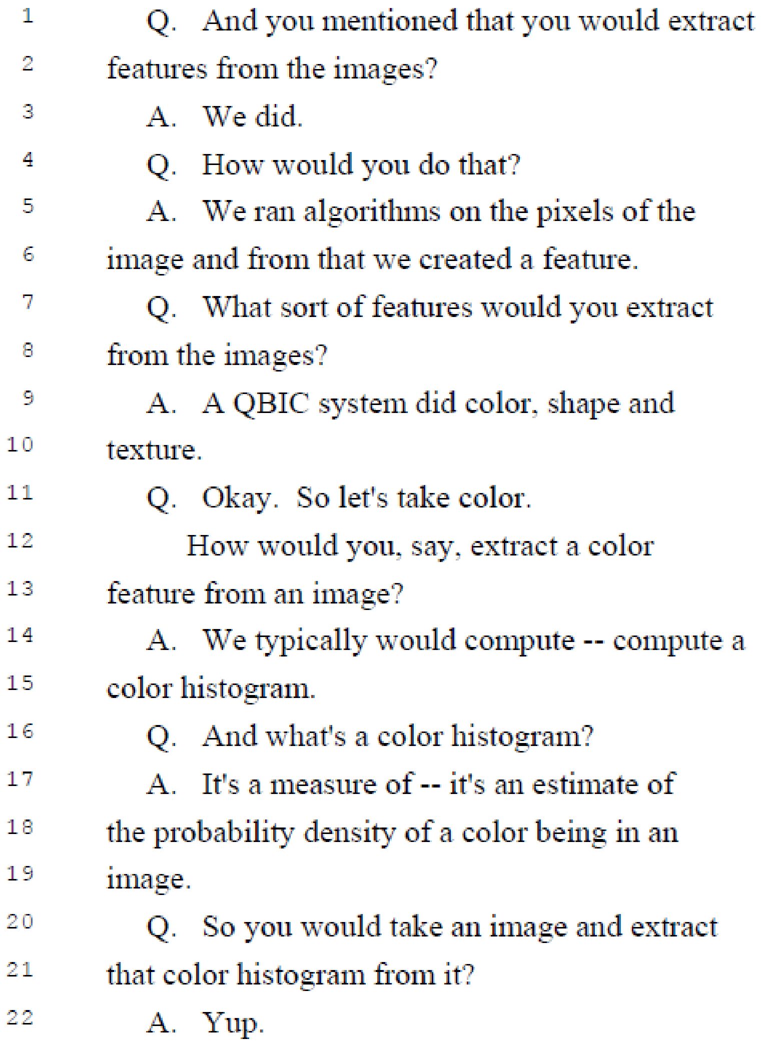 Q. And you mentioned that you would extract features from the images? A. We did. Q. How would you do that? A. We ran algorithms on the pixels of the image and from that we created a feature. Q. What sort of features would you extract from the images? A. A QBIC system did color, shape and texture. Q. Okay. So let's take color. How would you, say, extract a color feature from an image? A. We typically would compute -- compute a color histogram. Q. And what's a color histogram? A. It's a measure of -- it's an estimate of the probability density of a color being in an image. Q. So you would take an image and extract that color histogram from it? A. Yup.

		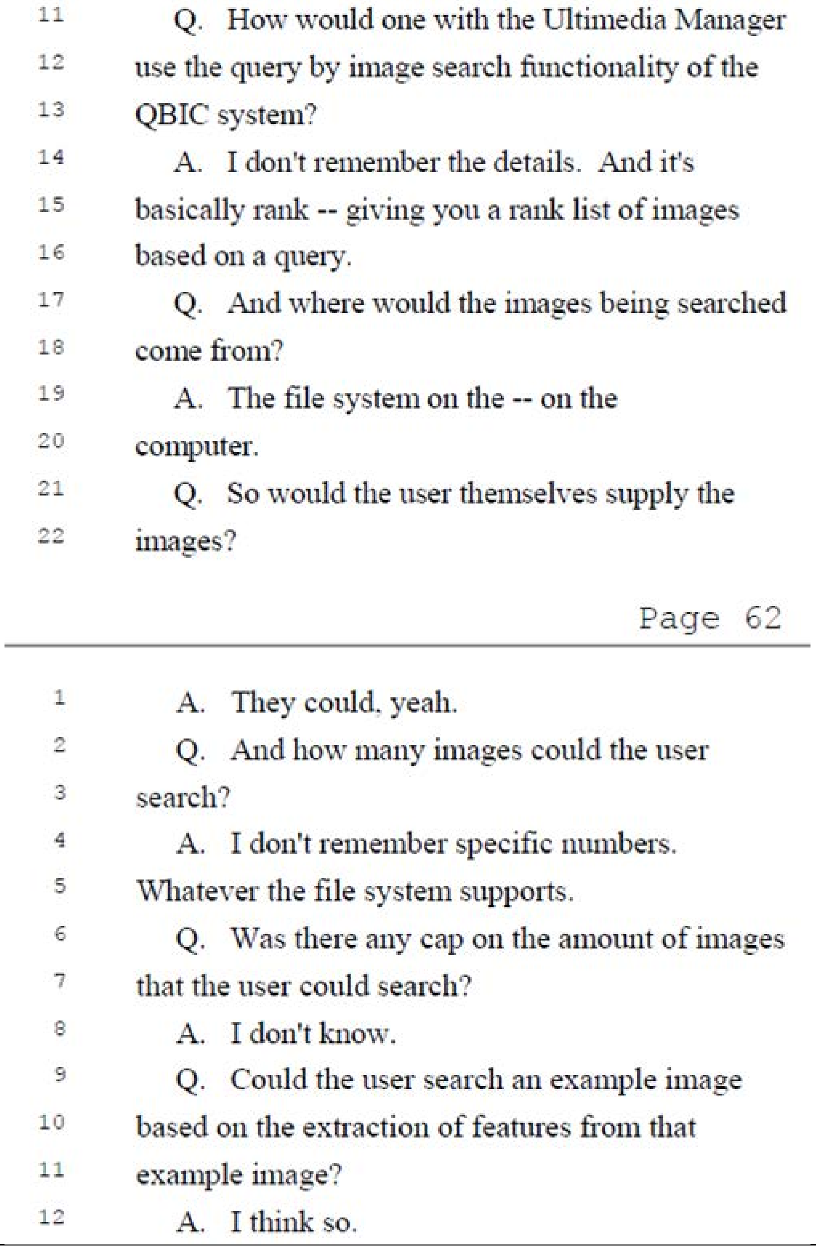 Q. How would one with the Ultimedia Manager use the query by image search functionality of the QBIC system? A. I don't remember the details. And it's basically rank -- giving you a rank list of images based on a query. Q. And where would the images being searched come from? A. The file system on the -- on the computer. Q. So would the user themselves supply the images? Page 62 A. They could, yeah. Q. And how many images could the user search? A. I don't remember specific numbers. Whatever the file system supports. Q. Was there any cap on the amount of images that the user could search? A. I don't know. Q. Could the user search an example image based on the extraction of features from that example image? A. I think so.

Element	'252 Claim Recitation	Exemplary Citations in Reference
		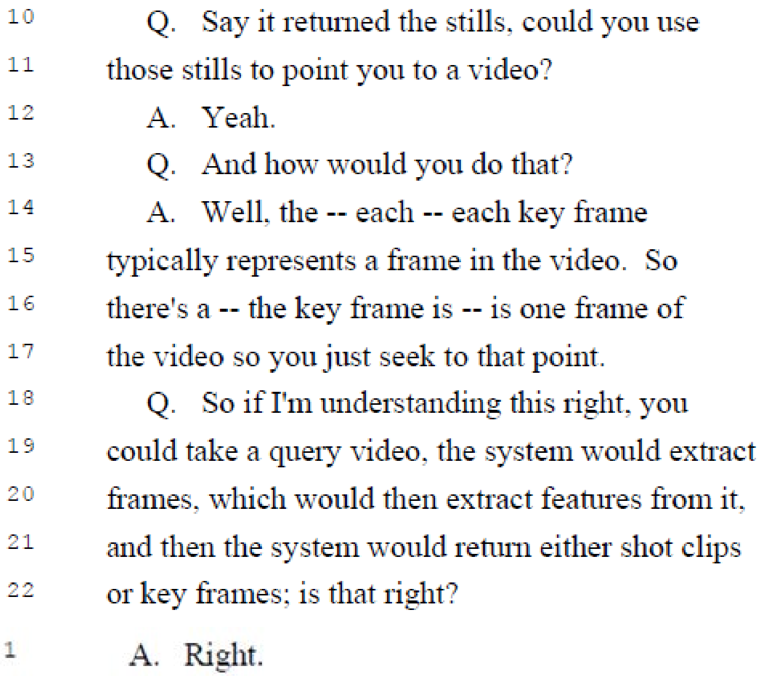 Q. Say it returned the stills, could you use those stills to point you to a video? A. Yeah. Q. And how would you do that? A. Well, the -- each -- each key frame typically represents a frame in the video. So there's a -- the key frame is -- is one frame of the video so you just seek to that point. Q. So if I'm understanding this right, you could take a query video, the system would extract frames, which would then extract features from it, and then the system would return either shot clips or key frames; is that right? A. Right. Flickner 8/29/23 Depo at 42

Element	'252 Claim Recitation	Exemplary Citations in Reference
		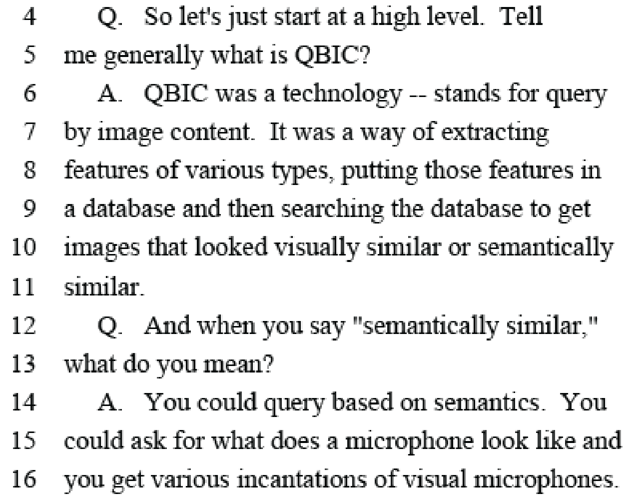 Q. So let's just start at a high level. Tell me generally what is QBIC? A. QBIC was a technology -- stands for query by image content. It was a way of extracting features of various types, putting those features in a database and then searching the database to get images that looked visually similar or semantically similar. Q. And when you say "semantically similar," what do you mean? A. You could query based on semantics. You could ask for what does a microphone look like and you get various incantations of visual microphones. *Id.* at 205–206

Element	'252 Claim Recitation	Exemplary Citations in Reference
		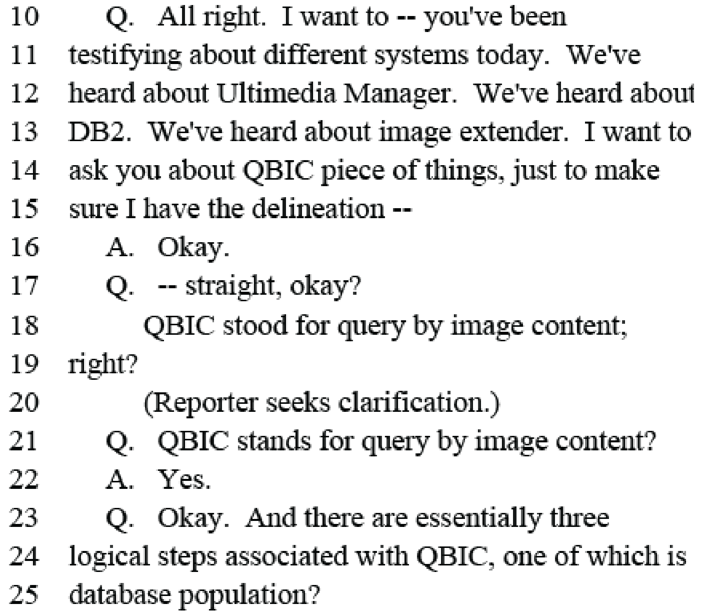 Q. All right. I want to -- you've been testifying about different systems today. We've heard about Ultimedia Manager. We've heard about DB2. We've heard about image extender. I want to ask you about QBIC piece of things, just to make sure I have the delineation -- A. Okay. Q. -- straight, okay? QBIC stood for query by image content; right? (Reporter seeks clarification.) Q. QBIC stands for query by image content? A. Yes. Q. Okay. And there are essentially three logical steps associated with QBIC, one of which is database population? 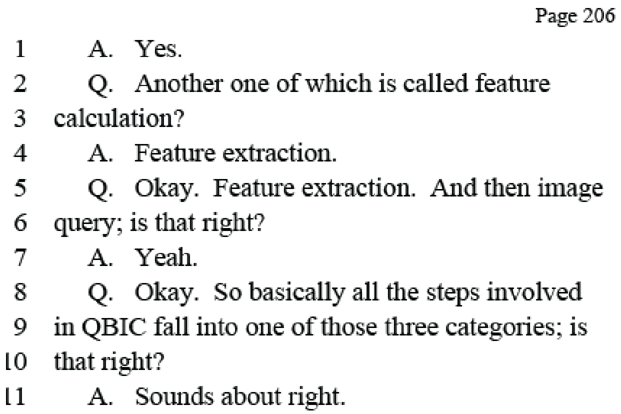 Page 206 A. Yes. Q. Another one of which is called feature calculation? A. Feature extraction. Q. Okay. Feature extraction. And then image query; is that right? A. Yeah. Q. Okay. So basically all the steps involved in QBIC fall into one of those three categories; is that right? A. Sounds about right. The QBIC system was implemented in IBM commercial products, including in

Element	'252 Claim Recitation	Exemplary Citations in Reference
		IBM's Ultimedia Manager, the DB2 series of products such as the Image Extender in DB2 Universal Database, and the Enterprise Information Portal. All of the features referenced in this chart were incorporated into these commercial implementations before November 6, 2000, as reflected in a combination of scientific papers, IBM's internal documentation, and the QBIC demo software applications. *See, e.g.*, Flickner 8/29/23 Depo Tr. at 33:9-24, 71:21-72:10, 164:21-166:2, 166:11-175:24, 178:5-186:17, 186:21-192:4, 192:14-194:1, 195:16-196:4, 194:1-308:1; IBM 000001-617; IBM 000747; IBM 000777-880; IBM 000893-902; IBM 0002315-320; IBM 0002390-404; IBM 0002418-430.
[18A]	providing a mobile device having a camera the mobile device configured to capture an image and configured to transmit data relating to the image an image processing platform;	QBIC discloses a mobile device having a camera, where the mobile device is configured to capture an image and configured to transmit data relating to the image an image processing platform. 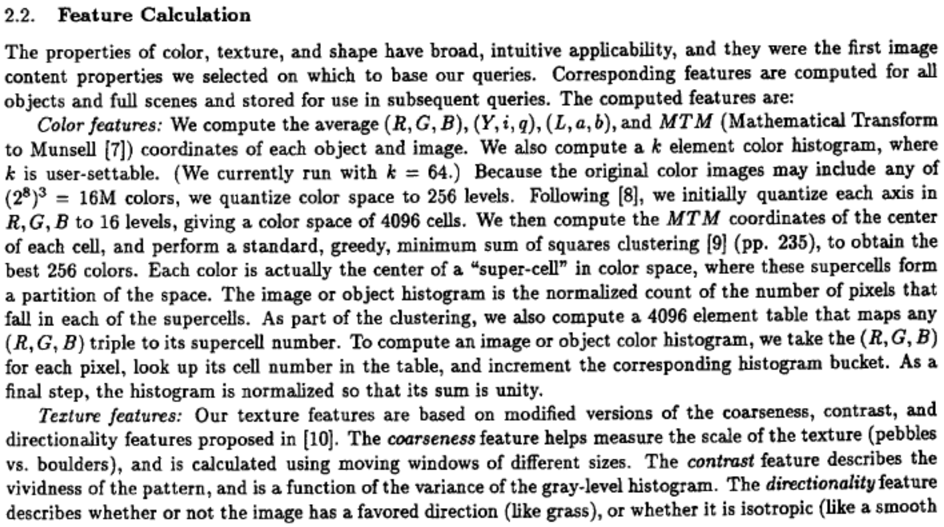 2.2. **Feature Calculation** The properties of color, texture, and shape have broad, intuitive applicability, and they were the first image content properties we selected on which to base our queries. Corresponding features are computed for all objects and full scenes and stored for use in subsequent queries. The computed features are: *Color features:* We compute the average (R, G, B), (Y, i, q), (L, a, b), and MTM (Mathematical Transform to Munsell [7]) coordinates of each object and image. We also compute a k element color histogram, where k is user-settable. (We currently run with $k = 64$.) Because the original color images may include any of $(2^8)^3 = 16$M colors, we quantize color space to 256 levels. Following [8], we initially quantize each axis in R, G, B to 16 levels, giving a color space of 4096 cells. We then compute the MTM coordinates of the center of each cell, and perform a standard, greedy, minimum sum of squares clustering [9] (pp. 235), to obtain the best 256 colors. Each color is actually the center of a "super-cell" in color space, where these supercells form a partition of the space. The image or object histogram is the normalized count of the number of pixels that fall in each of the supercells. As part of the clustering, we also compute a 4096 element table that maps any (R, G, B) triple to its supercell number. To compute an image or object color histogram, we take the (R, G, B) for each pixel, look up its cell number in the table, and increment the corresponding histogram bucket. As a final step, the histogram is normalized so that its sum is unity. *Texture features:* Our texture features are based on modified versions of the coarseness, contrast, and directionality features proposed in [10]. The *coarseness* feature helps measure the scale of the texture (pebbles vs. boulders), and is calculated using moving windows of different sizes. The *contrast* feature describes the vividness of the pattern, and is a function of the variance of the gray-level histogram. The *directionality* feature describes whether or not the image has a favored direction (like grass), or whether it is isotropic (like a smooth

Element	'252 Claim Recitation	Exemplary Citations in Reference
		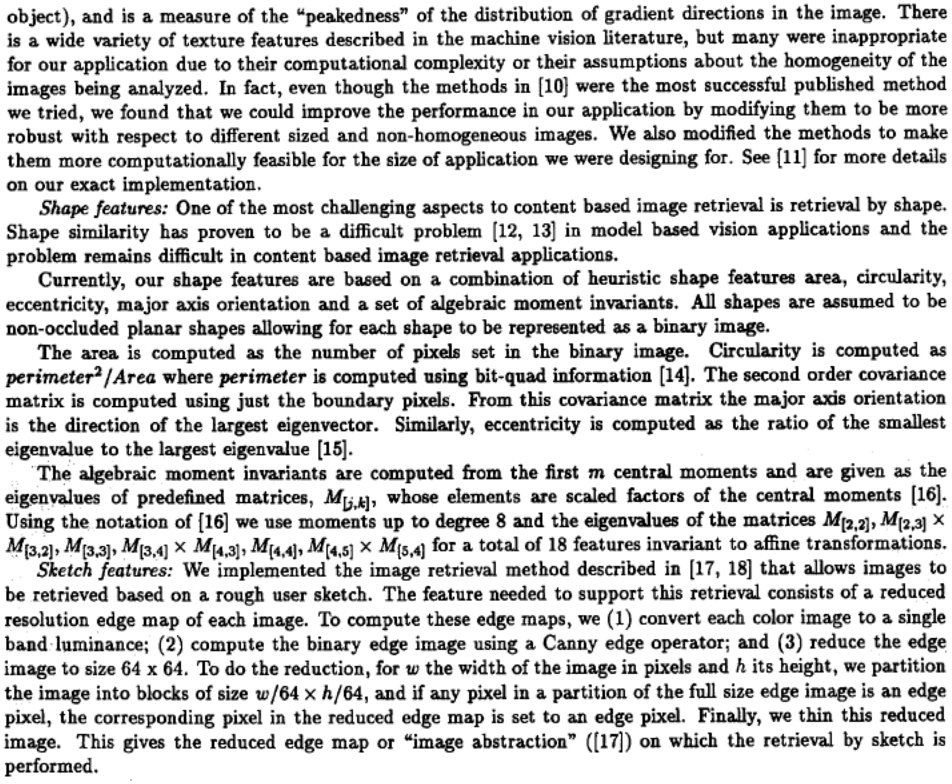 object), and is a measure of the "peakedness" of the distribution of gradient directions in the image. There is a wide variety of texture features described in the machine vision literature, but many were inappropriate for our application due to their computational complexity or their assumptions about the homogeneity of the images being analyzed. In fact, even though the methods in [10] were the most successful published method we tried, we found that we could improve the performance in our application by modifying them to be more robust with respect to different sized and non-homogeneous images. We also modified the methods to make them more computationally feasible for the size of application we were designing for. See [11] for more details on our exact implementation. *Shape features:* One of the most challenging aspects to content based image retrieval is retrieval by shape. Shape similarity has proven to be a difficult problem [12, 13] in model based vision applications and the problem remains difficult in content based image retrieval applications. Currently, our shape features are based on a combination of heuristic shape features area, circularity, eccentricity, major axis orientation and a set of algebraic moment invariants. All shapes are assumed to be non-occluded planar shapes allowing for each shape to be represented as a binary image. The area is computed as the number of pixels set in the binary image. Circularity is computed as $perimeter^2/Area$ where $perimeter$ is computed using bit-quad information [14]. The second order covariance matrix is computed using just the boundary pixels. From this covariance matrix the major axis orientation is the direction of the largest eigenvector. Similarly, eccentricity is computed as the ratio of the smallest eigenvalue to the largest eigenvalue [15]. The algebraic moment invariants are computed from the first m central moments and are given as the eigenvalues of predefined matrices, $M_{[j,k]}$, whose elements are scaled factors of the central moments [16]. Using the notation of [16] we use moments up to degree 8 and the eigenvalues of the matrices $M_{[2,2]}, M_{[2,3]} \times M_{[3,2]}, M_{[3,3]}, M_{[3,4]} \times M_{[4,3]}, M_{[4,4]}, M_{[4,5]} \times M_{[5,4]}$ for a total of 18 features invariant to affine transformations. *Sketch features:* We implemented the image retrieval method described in [17, 18] that allows images to be retrieved based on a rough user sketch. The feature needed to support this retrieval consists of a reduced resolution edge map of each image. To compute these edge maps, we (1) convert each color image to a single band luminance; (2) compute the binary edge image using a Canny edge operator; and (3) reduce the edge image to size 64 x 64. To do the reduction, for w the width of the image in pixels and h its height, we partition the image into blocks of size $w/64 \times h/64$, and if any pixel in a partition of the full size edge image is an edge pixel, the corresponding pixel in the reduced edge map is set to an edge pixel. Finally, we thin this reduced image. This gives the reduced edge map or "image abstraction" ([17]) on which the retrieval by sketch is performed. IBM0002390 at 2392-93

Element	’252 Claim Recitation	Exemplary Citations in Reference
		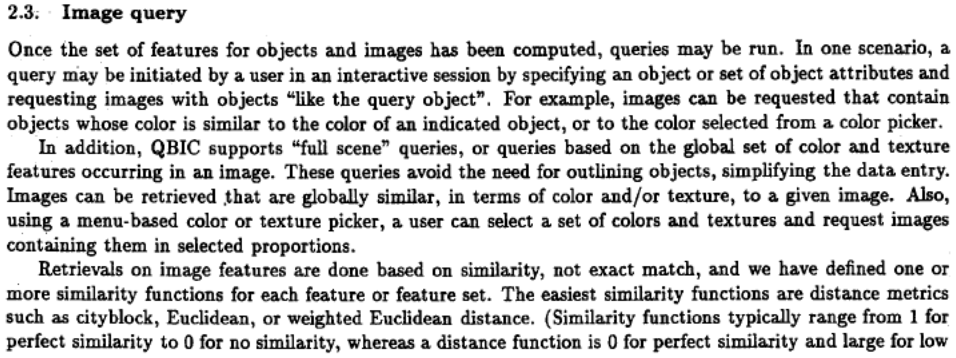 **2.3. Image query** Once the set of features for objects and images has been computed, queries may be run. In one scenario, a query may be initiated by a user in an interactive session by specifying an object or set of object attributes and requesting images with objects “like the query object”. For example, images can be requested that contain objects whose color is similar to the color of an indicated object, or to the color selected from a color picker. In addition, QBIC supports “full scene” queries, or queries based on the global set of color and texture features occurring in an image. These queries avoid the need for outlining objects, simplifying the data entry. Images can be retrieved that are globally similar, in terms of color and/or texture, to a given image. Also, using a menu-based color or texture picker, a user can select a set of colors and textures and request images containing them in selected proportions. Retrievals on image features are done based on similarity, not exact match, and we have defined one or more similarity functions for each feature or feature set. The easiest similarity functions are distance metrics such as cityblock, Euclidean, or weighted Euclidean distance. (Similarity functions typically range from 1 for perfect similarity to 0 for no similarity, whereas a distance function is 0 for perfect similarity and large for low

Element	'252 Claim Recitation	Exemplary Citations in Reference
		similarity, so we must convert as necessary to obtain consistent measures.) Different similarity measures may be used, and they may vary from application to application. The matching methods and similarity functions for the different feature types are described below: *Color:* For average color, the distance between a query object and database object is weighted Euclidean distance. The weights are user adjustable so that we may, for example, downweight the L component of an (L, a, b) query. Best results, as judged by our subjective evaluation in on-going experiments, are obtained by representing the average color in the MTM color space, and by inversely weighting each component by its standard deviation over the samples in the database. Improved results are obtained when color matching is done not on average color, but on the distribution of colors that occur in an object or image. In this way, a red and blue object matches other red and blue objects better than to a purple object. To match color histograms, we use a method similar to the one described in [8]. Let X be the query histogram and Y the histogram of an item in the database, both normalized. We compute the element by element difference histogram Z. Then the similarity between X and Y is given by $\|Z\| = Z^T A Z$, where A is a symmetric color similarity matrix with $$a(i,j) \equiv 1 - d(c_i, c_j)/d_{max}$$ where c_i and c_j are the ith and jth colors in the color histograms, and $d(c_i, c_j)$ is the MTM color distance, and d_{max} is the maximum distance between any two colors ([8] used Luv distance.) This metric gives the (weighted) length of the difference vector between X and Y, weighted by A, which accounts for the perceptual distance between different pairs of colors. The result is a measure that accounts for both the difference between the amounts in any given color (e.g. a particular shade of red), as well as differences between similar colors (e.g. red and orange). Based on our experience, we also found it useful for the user to be able to request images with x% of color 1, y% of color 2, etc. where both the percentages and the colors are explicitly specified by the user. This allows searches such as for a beach scene (25% white (for the sand) and 50% blue (for the sky and water)), and for many scenes for which the user has seen before, and for which he/she can remember approximate color percentages. The color picker that allows this type of query is shown in Figure 2. 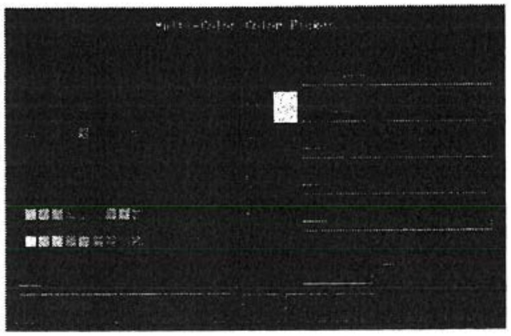 Figure 2: Multi-Color Color Picker with two selected colors. The palette on the left displays the set of selectable colors. Once a color is selected, it appears in a box on the right, and the amount of the color can be specified by the sliders. In this case, 25% light blue and 15% white have been selected. *Texture:* Texture distance is computed as weighted Euclidean distance in the three dimensional texture space. The most common weighting factors are the inverse variances for each component, computed over the

Element	'252 Claim Recitation	Exemplary Citations in Reference
		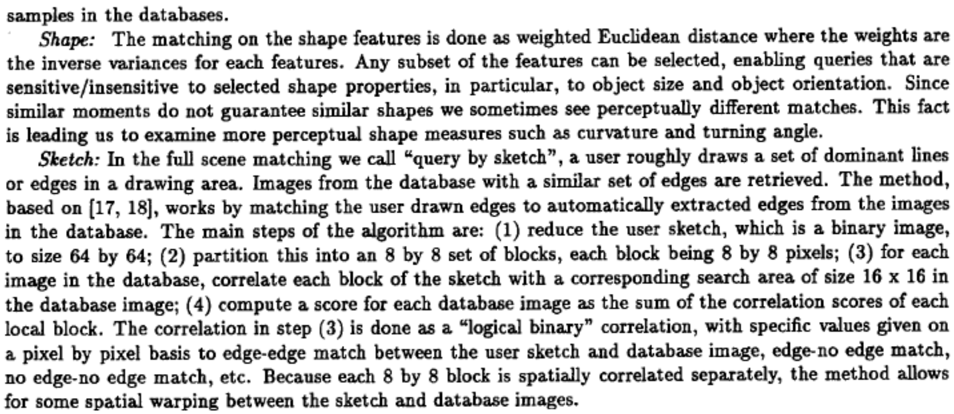 samples in the databases. *Shape:* The matching on the shape features is done as weighted Euclidean distance where the weights are the inverse variances for each features. Any subset of the features can be selected, enabling queries that are sensitive/insensitive to selected shape properties, in particular, to object size and object orientation. Since similar moments do not guarantee similar shapes we sometimes see perceptually different matches. This fact is leading us to examine more perceptual shape measures such as curvature and turning angle. *Sketch:* In the full scene matching we call "query by sketch", a user roughly draws a set of dominant lines or edges in a drawing area. Images from the database with a similar set of edges are retrieved. The method, based on [17, 18], works by matching the user drawn edges to automatically extracted edges from the images in the database. The main steps of the algorithm are: (1) reduce the user sketch, which is a binary image, to size 64 by 64; (2) partition this into an 8 by 8 set of blocks, each block being 8 by 8 pixels; (3) for each image in the database, correlate each block of the sketch with a corresponding search area of size 16 x 16 in the database image; (4) compute a score for each database image as the sum of the correlation scores of each local block. The correlation in step (3) is done as a "logical binary" correlation, with specific values given on a pixel by pixel basis to edge-edge match between the user sketch and database image, edge-no edge match, no edge-no edge match, etc. Because each 8 by 8 block is spatially correlated separately, the method allows for some spatial warping between the sketch and database images. IBM0002390 at 2393-95 Query by Image and Video Content: The QBIC System (IBM 0000893–902): "QBIC has two main components: database population (the process of creating an image database) and database query. During the population, images and videos are processed to extract features describing their content-colors, textures, shapes, and camera and object motion-and the features are stored in a database. During the query, the user composes a query graphically. Features are generated from the graphical query and then input to a matching engine that finds images or videos from the database with similar features. Figure 2 shows the system architecture." *Id*. at p. 3.

Element	'252 Claim Recitation	Exemplary Citations in Reference
		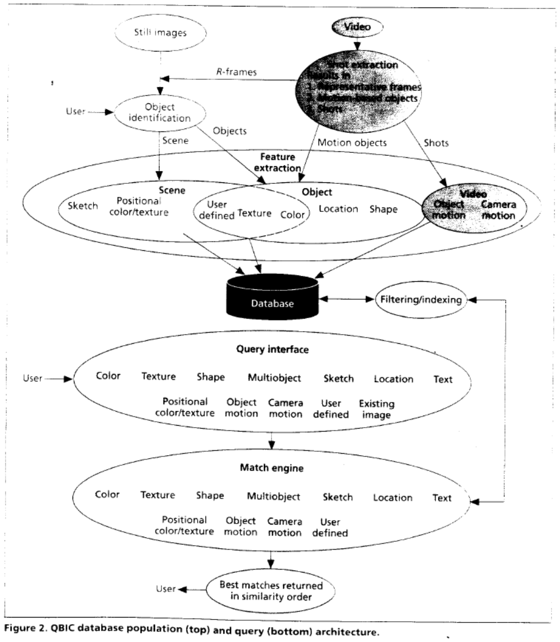 Figure 2. QBIC database population (top) and query (bottom) architecture. *Id*. at Fig. 2. Image, Audio, and Video Extenders Administration and Programming (IBM 000001–617): "DB2 (DB2) Universal Database (UDB) is a powerful, object-relational database manager. It stores and protects traditional numeric and character data, as well as large, complex objects (LOBs). DB2 extenders help you exploit DB2 's object-relational features. The extenders define distinct data types and special functions for image, audio, video, and text objects. … Because of that, the extenders give your applications a single point of access to any or all of these types of data, along with traditional numeric and character data. In addition, the extenders give your applications new ways to search for information." *Id*. at p. 3.

Element	'252 Claim Recitation	Exemplary Citations in Reference
		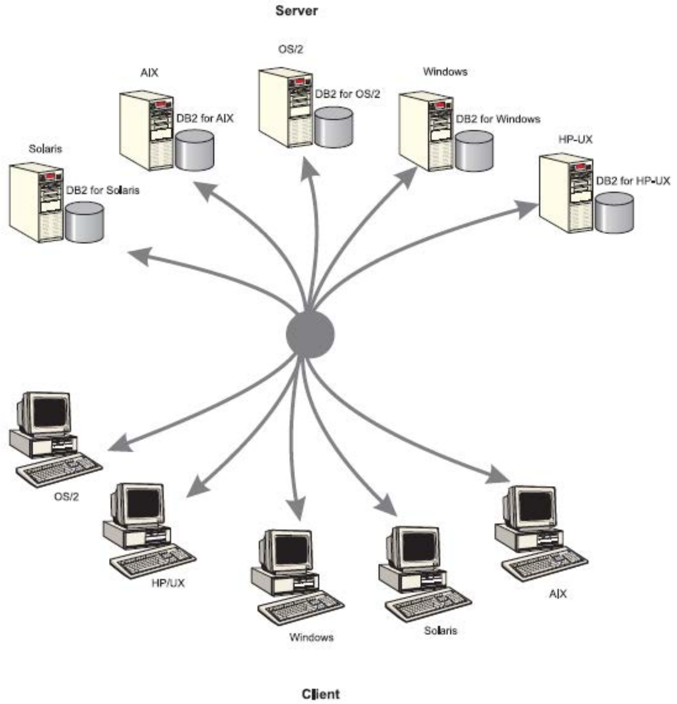 *Figure 6. DB2 extender platforms* *Id.* at p. 12 (Fig. 6). "The DB2 Extenders run in the DB2 client/server environment. This environment consists of a database server and one or more remote database clients. The DB2 extender services run on the server. Before you can access them, you have to start them." *Id.* at p. 47. OS/2 2.1 Ultimedia Tools (IBM 0001550–1805): "Perfect Image/2 can capture or import images and enhance the images for later use in Builder/2. To capture images, the Video Capture Adapter/A must be installed in your micro channel PC. The Video Capture Adapter/A works in conjunction with the video feature in Perfect Image/2 to capture images from live

Element	'252 Claim Recitation	Exemplary Citations in Reference
		motion video, paused motion video, or a still video player. Input sources can include television, broadcast signals, video camera, VCR, videodisc player, or any other video source that produces the proper video signal." *Id*. at p. 63. "Preparing for video capture: A video input device (for example, a VCR, video camera, or laserdisc) must be connected to the VCA before you can capture video images. Refer to your Video Capture Adapter/A reference manual for the correct procedure to connect a video source to the adapter card." *Id*. at p. 64. IBM Ultimedia Video I/O Adapter (IBM 0002217–221): "Video Capture: The adapter provides video capture on a frame-by-frame (field-by-field) basis of "live" video from a single source. Selection of S-video or CVBS input formats is enabled. The captured video is converted from NTSC (PAL) to YUV. The YUV digital video can be converted to 24-bit RGB, 8-bit dithered RGB, or remain unchanged. . . . Video Monitoring: Monitoring of incoming video data is supported with a graphics adapter and/or a video monitor supported by IBM. The image sizes shown in table #1 are supported for viewing with a supported graphics adapter." *Id*. at p. 2.

Element	'252 Claim Recitation	Exemplary Citations in Reference
		Flickner Depo at 38-41 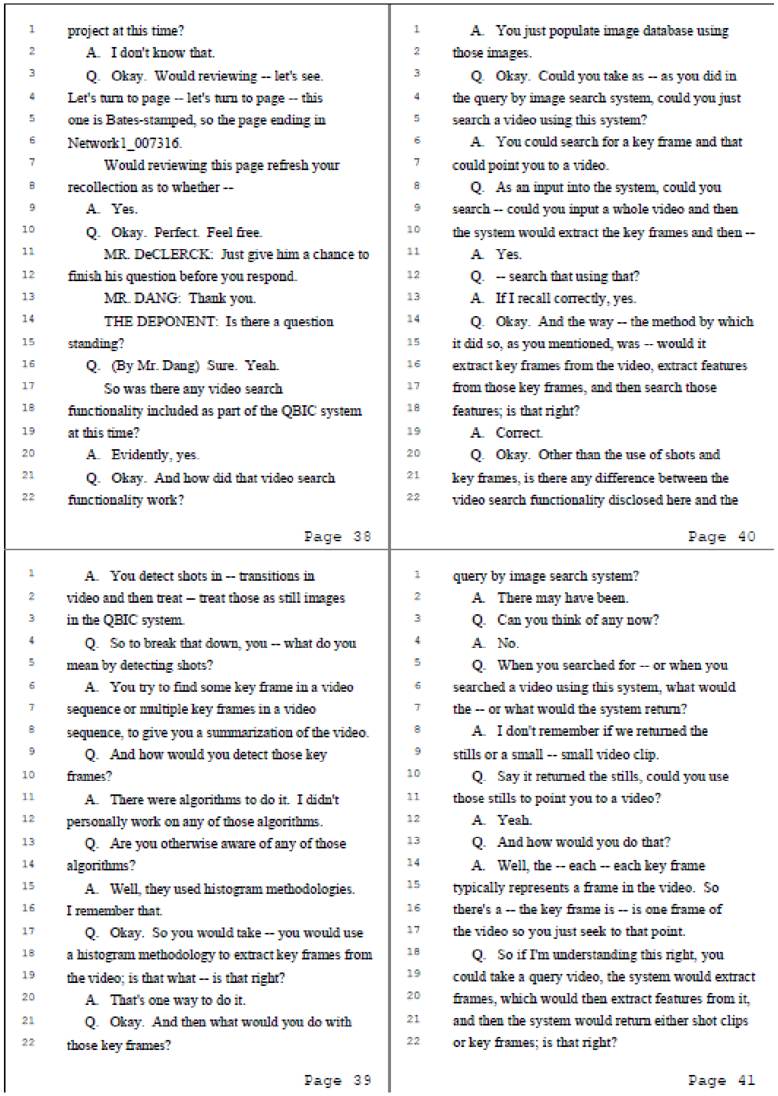 project at this time? A. I don't know that. Q. Okay. Would reviewing -- let's see. Let's turn to page -- let's turn to page -- this one is Bates-stamped, so the page ending in Network1_007316. Would reviewing this page refresh your recollection as to whether -- A. Yes. Q. Okay. Perfect. Feel free. MR. DeCLERCK: Just give him a chance to finish his question before you respond. MR. DANG: Thank you. THE DEPONENT: Is there a question standing? Q. (By Mr. Dang) Sure. Yeah. So was there any video search functionality included as part of the QBIC system at this time? A. Evidently, yes. Q. Okay. And how did that video search functionality work? Page 38 A. You detect shots in -- transitions in video and then treat -- treat those as still images in the QBIC system. Q. So to break that down, you -- what do you mean by detecting shots? A. You try to find some key frame in a video sequence or multiple key frames in a video sequence, to give you a summarization of the video. Q. And how would you detect those key frames? A. There were algorithms to do it. I didn't personally work on any of those algorithms. Q. Are you otherwise aware of any of those algorithms? A. Well, they used histogram methodologies. I remember that. Q. Okay. So you would take -- you would use a histogram methodology to extract key frames from the video; is that what -- is that right? A. That's one way to do it. Q. Okay. And then what would you do with those key frames? Page 39 A. You just populate image database using those images. Q. Okay. Could you take as -- as you did in the query by image search system, could you just search a video using this system? A. You could search for a key frame and that could point you to a video. Q. As an input into the system, could you search -- could you input a whole video and then the system would extract the key frames and then -- A. Yes. Q. -- search that using that? A. If I recall correctly, yes. Q. Okay. And the way -- the method by which it did so, as you mentioned, was -- would it extract key frames from the video, extract features from those key frames, and then search those features; is that right? A. Correct. Q. Okay. Other than the use of shots and key frames, is there any difference between the video search functionality disclosed here and the Page 40 query by image search system? A. There may have been. Q. Can you think of any now? A. No. Q. When you searched for -- or when you searched a video using this system, what would the -- or what would the system return? A. I don't remember if we returned the stills or a small -- small video clip. Q. Say it returned the stills, could you use those stills to point you to a video? A. Yeah. Q. And how would you do that? A. Well, the -- each -- each key frame typically represents a frame in the video. So there's a -- the key frame is -- is one frame of the video so you just seek to that point. Q. So if I'm understanding this right, you could take a query video, the system would extract frames, which would then extract features from it, and then the system would return either shot clips or key frames; is that right? Page 41

Element	'252 Claim Recitation	Exemplary Citations in Reference
		Flickner Depo at 139-140 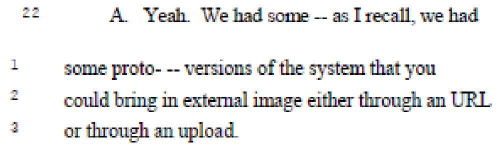 A. Yeah. We had some -- as I recall, we had some proto- -- versions of the system that you could bring in external image either through an URL or through an upload. Flickner Depo at 125 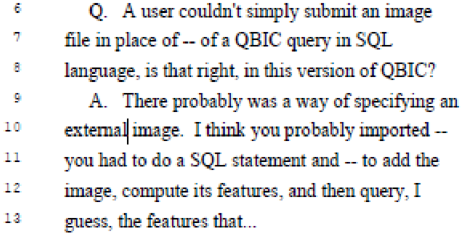 Q. A user couldn't simply submit an image file in place of -- of a QBIC query in SQL language, is that right, in this version of QBIC? A. There probably was a way of specifying an external image. I think you probably imported -- you had to do a SQL statement and -- to add the image, compute its features, and then query, I guess, the features that... Flickner 8/29/23 Depo at 53 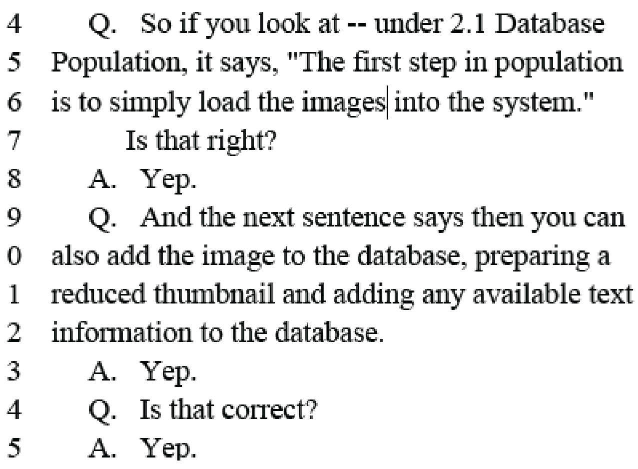 Q. So if you look at -- under 2.1 Database Population, it says, "The first step in population is to simply load the images into the system." Is that right? A. Yep. Q. And the next sentence says then you can also add the image to the database, preparing a reduced thumbnail and adding any available text information to the database. A. Yep. Q. Is that correct? A. Yep. *Id.* at 60–61

Element	'252 Claim Recitation	Exemplary Citations in Reference
		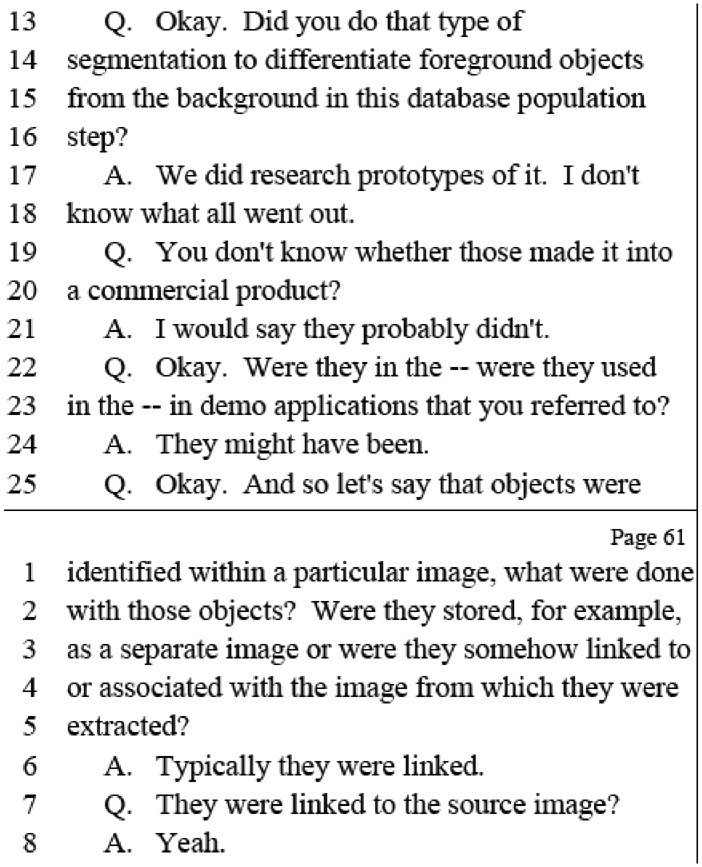 Q. Okay. Did you do that type of segmentation to differentiate foreground objects from the background in this database population step? A. We did research prototypes of it. I don't know what all went out. Q. You don't know whether those made it into a commercial product? A. I would say they probably didn't. Q. Okay. Were they in the -- were they used in the -- in demo applications that you referred to? A. They might have been. Q. Okay. And so let's say that objects were Page 61 identified within a particular image, what were done with those objects? Were they stored, for example, as a separate image or were they somehow linked to or associated with the image from which they were extracted? A. Typically they were linked. Q. They were linked to the source image? A. Yeah. *Id.* at 32–33

Element	'252 Claim Recitation	Exemplary Citations in Reference
		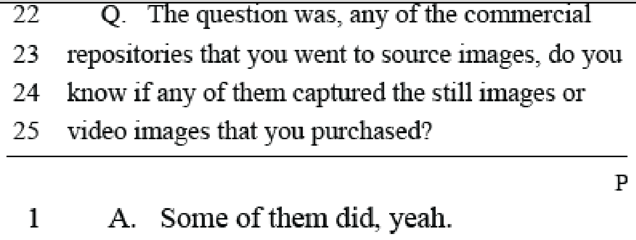 Q. The question was, any of the commercial repositories that you went to source images, do you know if any of them captured the still images or video images that you purchased? P A. Some of them did, yeah. *Id.* at 74–75

			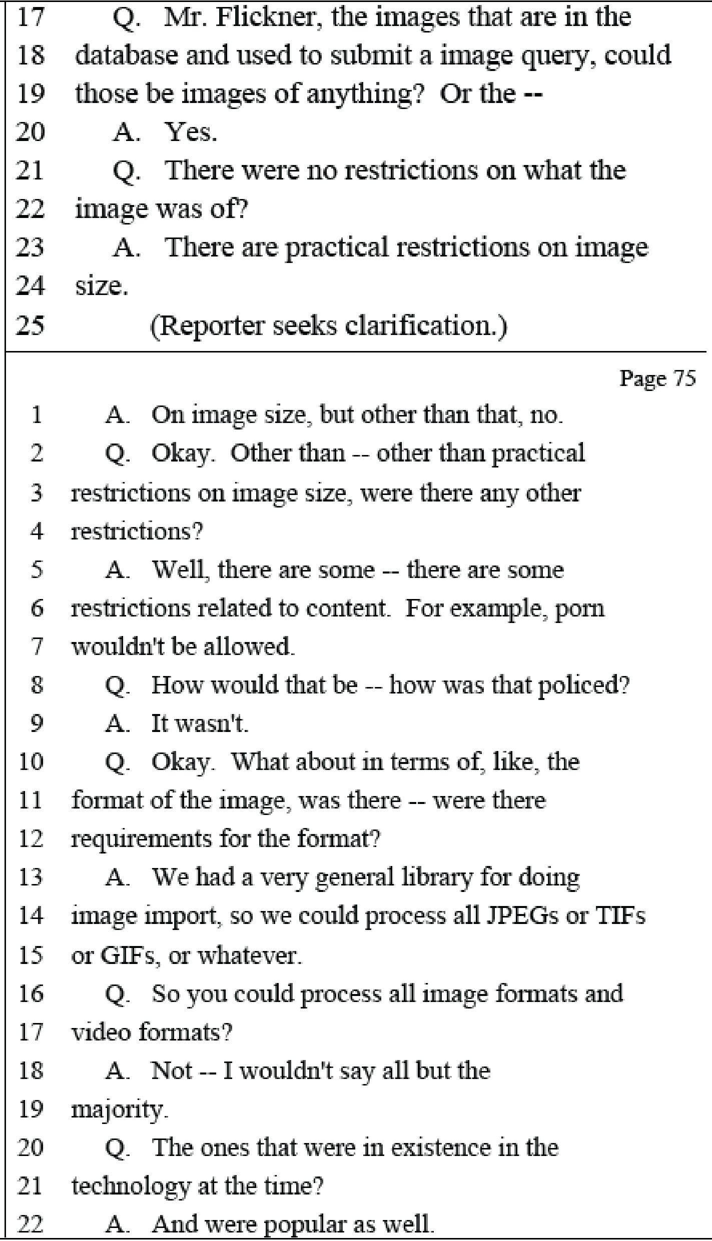 Q. Mr. Flickner, the images that are in the database and used to submit a image query, could those be images of anything? Or the -- A. Yes. Q. There were no restrictions on what the image was of? A. There are practical restrictions on image size. (Reporter seeks clarification.) Page 75 A. On image size, but other than that, no. Q. Okay. Other than -- other than practical restrictions on image size, were there any other restrictions? A. Well, there are some -- there are some restrictions related to content. For example, porn wouldn't be allowed. Q. How would that be -- how was that policed? A. It wasn't. Q. Okay. What about in terms of, like, the format of the image, was there -- were there requirements for the format? A. We had a very general library for doing image import, so we could process all JPEGs or TIFs or GIFs, or whatever. Q. So you could process all image formats and video formats? A. Not -- I wouldn't say all but the majority. Q. The ones that were in existence in the technology at the time? A. And were popular as well.

Element	'252 Claim Recitation	Exemplary Citations in Reference
		Id. at 114 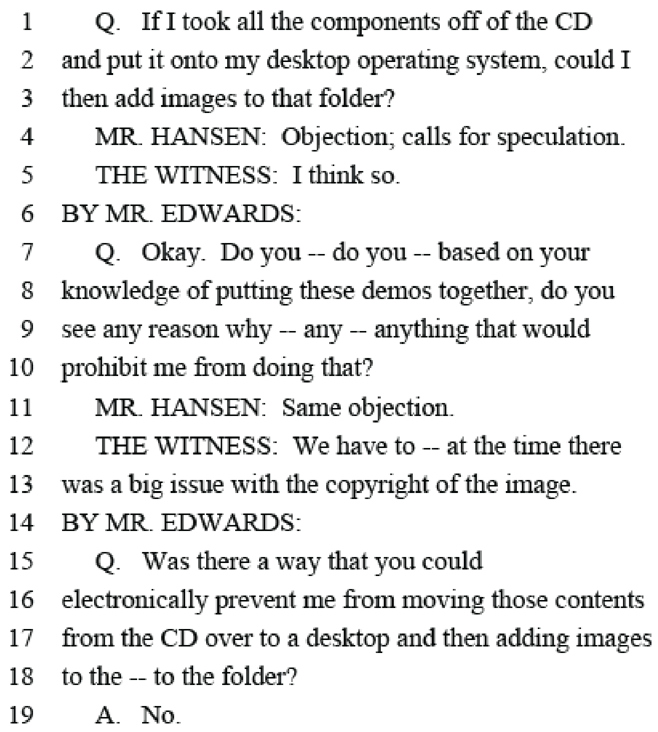 Q. If I took all the components off of the CD and put it onto my desktop operating system, could I then add images to that folder? MR. HANSEN: Objection; calls for speculation. THE WITNESS: I think so. BY MR. EDWARDS: Q. Okay. Do you -- do you -- based on your knowledge of putting these demos together, do you see any reason why -- any -- anything that would prohibit me from doing that? MR. HANSEN: Same objection. THE WITNESS: We have to -- at the time there was a big issue with the copyright of the image. BY MR. EDWARDS: Q. Was there a way that you could electronically prevent me from moving those contents from the CD over to a desktop and then adding images to the -- to the folder? A. No. *Id.* at 202

Element	'252 Claim Recitation	Exemplary Citations in Reference
		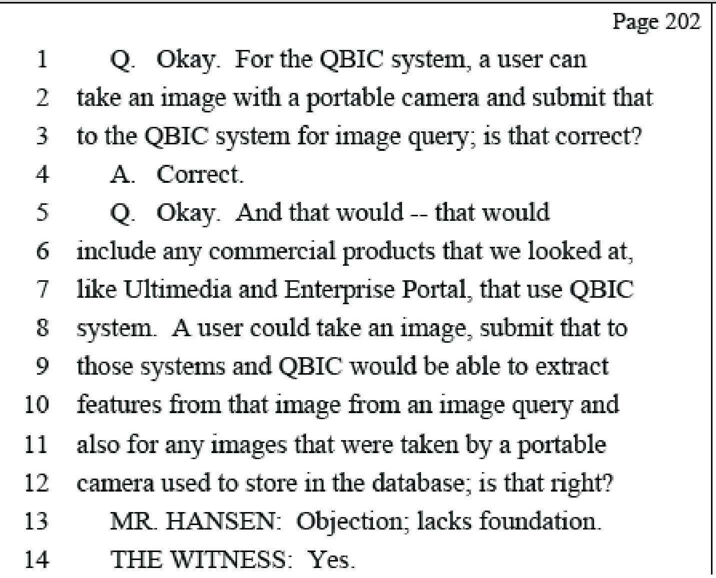Page 202 Q. Okay. For the QBIC system, a user can take an image with a portable camera and submit that to the QBIC system for image query; is that correct? A. Correct. Q. Okay. And that would -- that would include any commercial products that we looked at, like Ultimedia and Enterprise Portal, that use QBIC system. A user could take an image, submit that to those systems and QBIC would be able to extract features from that image from an image query and also for any images that were taken by a portable camera used to store in the database; is that right? MR. HANSEN: Objection; lacks foundation. THE WITNESS: Yes. *Id.* at 300 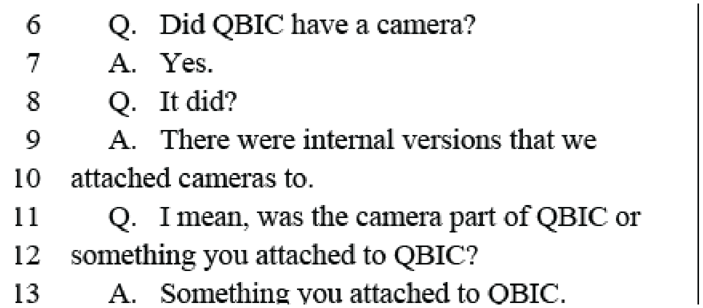Q. Did QBIC have a camera? A. Yes. Q. It did? A. There were internal versions that we attached cameras to. Q. I mean, was the camera part of QBIC or something you attached to QBIC? A. Something you attached to QBIC. *Id.* at 312

Element	'252 Claim Recitation	Exemplary Citations in Reference
		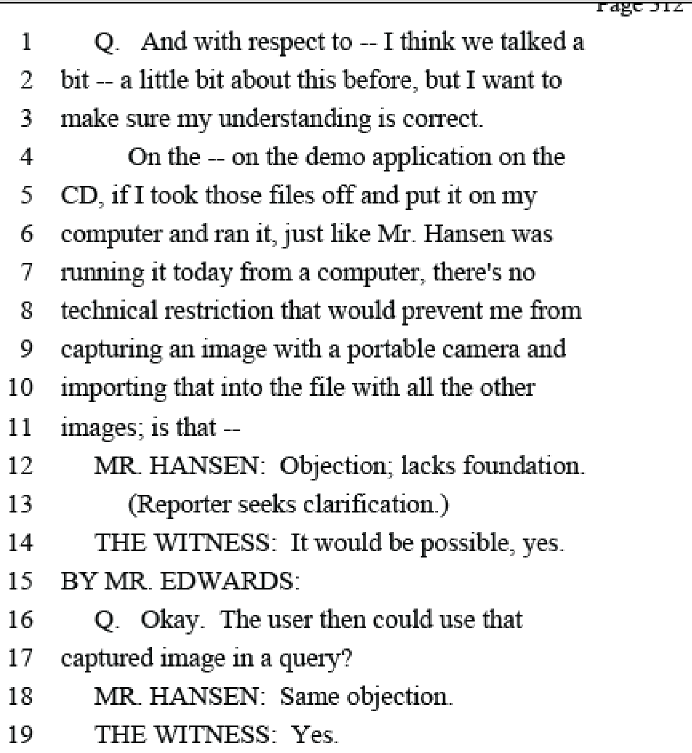 Page 312 Q. And with respect to -- I think we talked a bit -- a little bit about this before, but I want to make sure my understanding is correct. On the -- on the demo application on the CD, if I took those files off and put it on my computer and ran it, just like Mr. Hansen was running it today from a computer, there's no technical restriction that would prevent me from capturing an image with a portable camera and importing that into the file with all the other images; is that -- MR. HANSEN: Objection; lacks foundation. (Reporter seeks clarification.) THE WITNESS: It would be possible, yes. BY MR. EDWARDS: Q. Okay. The user then could use that captured image in a query? MR. HANSEN: Same objection. THE WITNESS: Yes. *See also* Flickner Depo at, e.g., 24, 41-42, 62-63 (user could take a picture or video and submit it to the QBIC system). To the extent it is found that QBIC does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that QBIC does not anticipate this claim, QBIC renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to

Element	'252 Claim Recitation	Exemplary Citations in Reference
		combine QBIC with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results. For example, the '252 patent acknowledges that acquiring information about objects through symbols, such as barcodes or some other means to encode detectable information in the image or on the object, were traditional methods known in the art and not part the present invention. '252 Patent, 3:7-16. Indeed, the '252 patent acknowledges that algorithms, software, and hardware components were commercially available at the time of the invention that could detect and decode symbols, such as barcodes or text, in an input image. *Id.* at 11:56-61. This limitation was thus known in the art by patentee's own admission, and by asserting this claim, Plaintiff is attempting to reclaim disavowed claim scope and claim subject matter which patentee has acknowledged is not part of the invention.
[18B]	configuring the image processing platform to receive the data relating to the image and to conduct image processing, including:	QBIC discloses configuring the image processing platform to receive the data relating to the image and to conduct image processing.

Element	'252 Claim Recitation	Exemplary Citations in Reference
		The QBIC Project: Querying Images by Content Using Color, Texture, and Shape (IBM 0002390–404) at 2390. 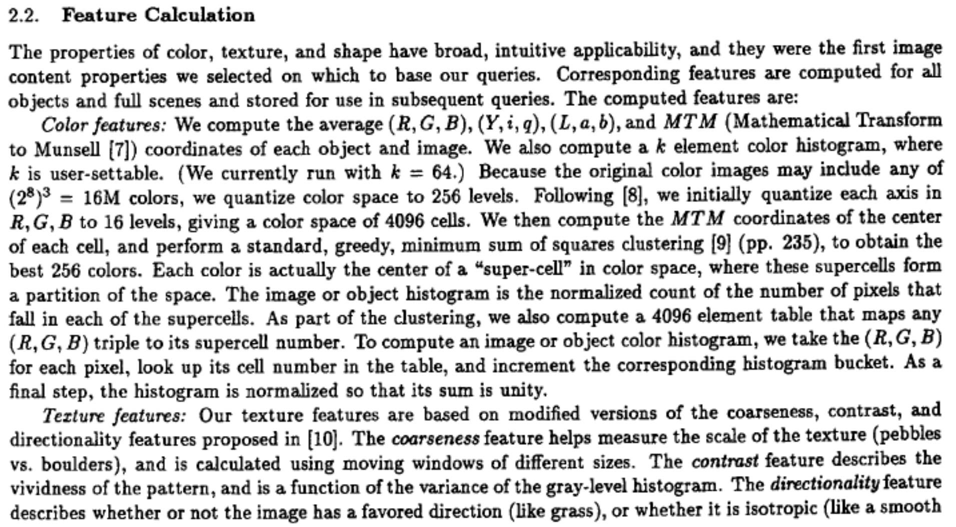 2.2. **Feature Calculation** The properties of color, texture, and shape have broad, intuitive applicability, and they were the first image content properties we selected on which to base our queries. Corresponding features are computed for all objects and full scenes and stored for use in subsequent queries. The computed features are: *Color features:* We compute the average (R, G, B), (Y, i, q), (L, a, b), and MTM (Mathematical Transform to Munsell [7]) coordinates of each object and image. We also compute a k element color histogram, where k is user-settable. (We currently run with $k = 64$.) Because the original color images may include any of $(2^8)^3 = 16$M colors, we quantize color space to 256 levels. Following [8], we initially quantize each axis in R, G, B to 16 levels, giving a color space of 4096 cells. We then compute the MTM coordinates of the center of each cell, and perform a standard, greedy, minimum sum of squares clustering [9] (pp. 235), to obtain the best 256 colors. Each color is actually the center of a "super-cell" in color space, where these supercells form a partition of the space. The image or object histogram is the normalized count of the number of pixels that fall in each of the supercells. As part of the clustering, we also compute a 4096 element table that maps any (R, G, B) triple to its supercell number. To compute an image or object color histogram, we take the (R, G, B) for each pixel, look up its cell number in the table, and increment the corresponding histogram bucket. As a final step, the histogram is normalized so that its sum is unity. *Texture features:* Our texture features are based on modified versions of the coarseness, contrast, and directionality features proposed in [10]. The *coarseness* feature helps measure the scale of the texture (pebbles vs. boulders), and is calculated using moving windows of different sizes. The *contrast* feature describes the vividness of the pattern, and is a function of the variance of the gray-level histogram. The *directionality* feature describes whether or not the image has a favored direction (like grass), or whether it is isotropic (like a smooth

Element	'252 Claim Recitation	Exemplary Citations in Reference
		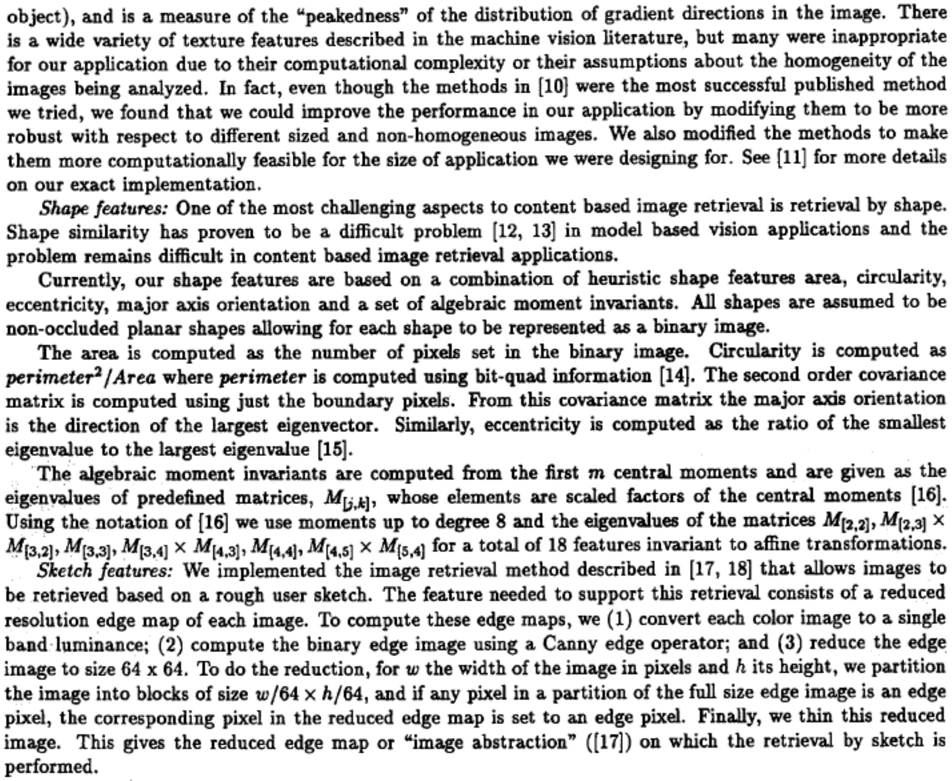 object), and is a measure of the "peakedness" of the distribution of gradient directions in the image. There is a wide variety of texture features described in the machine vision literature, but many were inappropriate for our application due to their computational complexity or their assumptions about the homogeneity of the images being analyzed. In fact, even though the methods in [10] were the most successful published method we tried, we found that we could improve the performance in our application by modifying them to be more robust with respect to different sized and non-homogeneous images. We also modified the methods to make them more computationally feasible for the size of application we were designing for. See [11] for more details on our exact implementation. *Shape features:* One of the most challenging aspects to content based image retrieval is retrieval by shape. Shape similarity has proven to be a difficult problem [12, 13] in model based vision applications and the problem remains difficult in content based image retrieval applications. Currently, our shape features are based on a combination of heuristic shape features area, circularity, eccentricity, major axis orientation and a set of algebraic moment invariants. All shapes are assumed to be non-occluded planar shapes allowing for each shape to be represented as a binary image. The area is computed as the number of pixels set in the binary image. Circularity is computed as $perimeter^2/Area$ where $perimeter$ is computed using bit-quad information [14]. The second order covariance matrix is computed using just the boundary pixels. From this covariance matrix the major axis orientation is the direction of the largest eigenvector. Similarly, eccentricity is computed as the ratio of the smallest eigenvalue to the largest eigenvalue [15]. The algebraic moment invariants are computed from the first m central moments and are given as the eigenvalues of predefined matrices, $M_{[j,k]}$, whose elements are scaled factors of the central moments [16]. Using the notation of [16] we use moments up to degree 8 and the eigenvalues of the matrices $M_{[2,2]}, M_{[2,3]} \times M_{[3,2]}, M_{[3,3]}, M_{[3,4]} \times M_{[4,3]}, M_{[4,4]}, M_{[4,5]} \times M_{[5,4]}$ for a total of 18 features invariant to affine transformations. *Sketch features:* We implemented the image retrieval method described in [17, 18] that allows images to be retrieved based on a rough user sketch. The feature needed to support this retrieval consists of a reduced resolution edge map of each image. To compute these edge maps, we (1) convert each color image to a single band luminance; (2) compute the binary edge image using a Canny edge operator; and (3) reduce the edge image to size 64 x 64. To do the reduction, for w the width of the image in pixels and h its height, we partition the image into blocks of size $w/64 \times h/64$, and if any pixel in a partition of the full size edge image is an edge pixel, the corresponding pixel in the reduced edge map is set to an edge pixel. Finally, we thin this reduced image. This gives the reduced edge map or "image abstraction" ([17]) on which the retrieval by sketch is performed. IBM0002390 at 2392-93

Element	'252 Claim Recitation	Exemplary Citations in Reference
		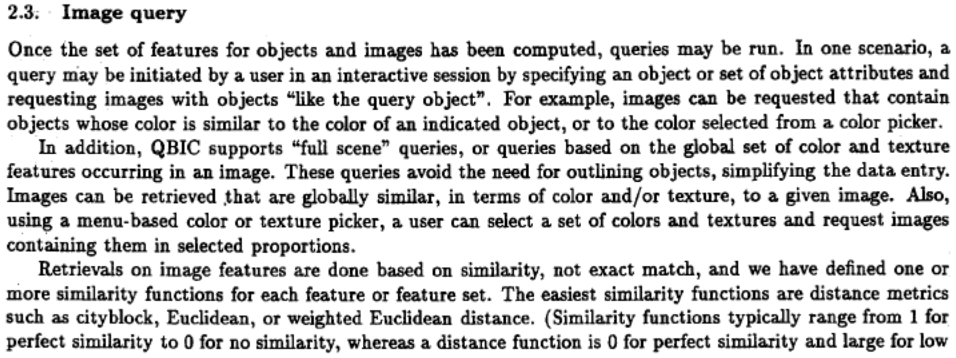 2.3. **Image query** Once the set of features for objects and images has been computed, queries may be run. In one scenario, a query may be initiated by a user in an interactive session by specifying an object or set of object attributes and requesting images with objects "like the query object". For example, images can be requested that contain objects whose color is similar to the color of an indicated object, or to the color selected from a color picker. In addition, QBIC supports "full scene" queries, or queries based on the global set of color and texture features occurring in an image. These queries avoid the need for outlining objects, simplifying the data entry. Images can be retrieved that are globally similar, in terms of color and/or texture, to a given image. Also, using a menu-based color or texture picker, a user can select a set of colors and textures and request images containing them in selected proportions. Retrievals on image features are done based on similarity, not exact match, and we have defined one or more similarity functions for each feature or feature set. The easiest similarity functions are distance metrics such as cityblock, Euclidean, or weighted Euclidean distance. (Similarity functions typically range from 1 for perfect similarity to 0 for no similarity, whereas a distance function is 0 for perfect similarity and large for low

Element	'252 Claim Recitation	Exemplary Citations in Reference
		similarity, so we must convert as necessary to obtain consistent measures.) Different similarity measures may be used, and they may vary from application to application. The matching methods and similarity functions for the different feature types are described below: *Color:* For average color, the distance between a query object and database object is weighted Euclidean distance. The weights are user adjustable so that we may, for example, downweight the L component of an (L, a, b) query. Best results, as judged by our subjective evaluation in on-going experiments, are obtained by representing the average color in the MTM color space, and by inversely weighting each component by its standard deviation over the samples in the database. Improved results are obtained when color matching is done not on average color, but on the distribution of colors that occur in an object or image. In this way, a red and blue object matches other red and blue objects better than to a purple object. To match color histograms, we use a method similar to the one described in [8]. Let X be the query histogram and Y the histogram of an item in the database, both normalized. We compute the element by element difference histogram Z. Then the similarity between X and Y is given by $\|Z\| = Z^T A Z$, where A is a symmetric color similarity matrix with $$a(i,j) \equiv 1 - d(c_i, c_j)/d_{max}$$ where c_i and c_j are the ith and jth colors in the color histograms, and $d(c_i, c_j)$ is the MTM color distance, and d_{max} is the maximum distance between any two colors ([8] used Luv distance.) This metric gives the (weighted) length of the difference vector between X and Y, weighted by A, which accounts for the perceptual distance between different pairs of colors. The result is a measure that accounts for both the difference between the amounts in any given color (e.g. a particular shade of red), as well as differences between similar colors (e.g. red and orange). Based on our experience, we also found it useful for the user to be able to request images with x% of color 1, y% of color 2, etc. where both the percentages and the colors are explicitly specified by the user. This allows searches such as for a beach scene (25% white (for the sand) and 50% blue (for the sky and water)), and for many scenes for which the user has seen before, and for which he/she can remember approximate color percentages. The color picker that allows this type of query is shown in Figure 2. 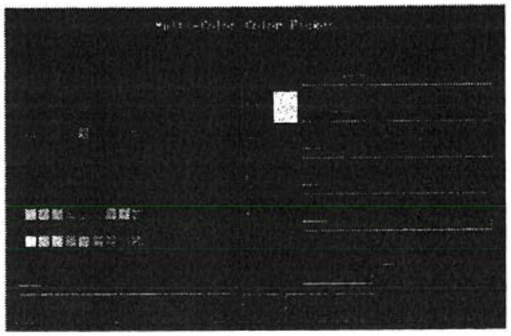 Figure 2: Multi-Color Color Picker with two selected colors. The palette on the left displays the set of selectable colors. Once a color is selected, it appears in a box on the right, and the amount of the color can be specified by the sliders. In this case, 25% light blue and 15% white have been selected. *Texture:* Texture distance is computed as weighted Euclidean distance in the three dimensional texture space. The most common weighting factors are the inverse variances for each component, computed over the

Element	'252 Claim Recitation	Exemplary Citations in Reference
		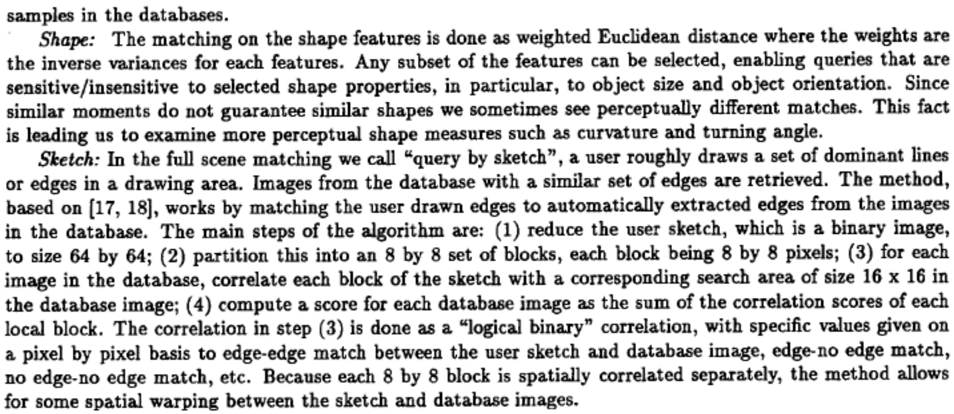samples in the databases. *Shape:* The matching on the shape features is done as weighted Euclidean distance where the weights are the inverse variances for each features. Any subset of the features can be selected, enabling queries that are sensitive/insensitive to selected shape properties, in particular, to object size and object orientation. Since similar moments do not guarantee similar shapes we sometimes see perceptually different matches. This fact is leading us to examine more perceptual shape measures such as curvature and turning angle. *Sketch:* In the full scene matching we call "query by sketch", a user roughly draws a set of dominant lines or edges in a drawing area. Images from the database with a similar set of edges are retrieved. The method, based on [17, 18], works by matching the user drawn edges to automatically extracted edges from the images in the database. The main steps of the algorithm are: (1) reduce the user sketch, which is a binary image, to size 64 by 64; (2) partition this into an 8 by 8 set of blocks, each block being 8 by 8 pixels; (3) for each image in the database, correlate each block of the sketch with a corresponding search area of size 16 x 16 in the database image; (4) compute a score for each database image as the sum of the correlation scores of each local block. The correlation in step (3) is done as a "logical binary" correlation, with specific values given on a pixel by pixel basis to edge-edge match between the user sketch and database image, edge-no edge match, no edge-no edge match, etc. Because each 8 by 8 block is spatially correlated separately, the method allows for some spatial warping between the sketch and database images. IBM0002390 at 2393-95 Query by Image and Video Content: The QBIC System (IBM 0000893–902): "To achieve this functionality, QBIC has two main components: database population (the process of creating an image database) and database query. During the population, images and videos are processed to extract features describing their content-colors, textures, shapes, and camera and object motion-and the features are stored in a database. During the query, the user composes a query graphically. Features are generated from the graphical query and then input to a matching engine that finds images or videos from the database with similar features. Figure 2 shows the system architecture." *Id.* at p. 3. "For both population and query, the QBIC data model has • still images or scenes (full images) that contain objects (subsets of an image), and • video shots that consist of sets of contiguous frames and contain motion objects. For still images, the QBIC data model distinguishes between 'scenes' (or images) and 'objects.' A scene is an image or single representative frame of video. An object is a part of a scene-for example, the fox in Figure 3-or a moving entity in a video. For still image database population, features are extracted from images and objects and stored in a database as shown in the top left part of Figure 2. Videos

Element	'252 Claim Recitation	Exemplary Citations in Reference
		are broken into clips called shots. Representative frames, or r-frames, are generated for each extracted shot. R-frames are treated as still images, and features are extracted and stored in the database. Further processing of shots generates motion objects-for example, a car moving across the screen. Queries are allowed on objects ('Find images with a red, round object'), scenes ('Find images that have approximately 30-percent red and IS-percent blue colors'), shots ('Find all shots panning from left to right'), or any combination ('Find images that have 30 percent red and contain a blue textured object')." *Id.* at p. 3. "**DATABASE POPULATION** In still image database population, the images are reduced to a standard-sized icon called a thumbnail and annotated with any available text information. Object identification is an optional but key part of this step. It lets users manually, semiautomatically, or fully automatically identify interesting regions—which we call objects—in the images. Internally, each object is represented as a binary mask. There may be an arbitrary number of objects per image. Objects can overlap and can consist of multiple disconnected components like the set of dots on a polka-dot dress. Text, like 'baby on beach,' can be associated with an outlined object or with the scene as a whole. **Object-outlining tools** Ideally, object identification would be automatic, but this is generally difficult. The alternative—manual identification—is tedious and can inhibit query-by-content applications. As a result, we have devoted considerable effort to developing tools to aid in this step. In recent work, we have successfully used fully automatic unsupervised segmentation methods along with a foreground/background model to identify objects in a restricted class of images. The images, typical of museums and retail catalogs, have a small number of foreground objects on a generally separable background. Figure 4 shows example results. Even in this domain, robust algorithms are required because of the textured and variegated backgrounds. We also provide semiautomatic tools for identifying objects. One is an enhanced flood-fill technique. Flood-fill methods, found in most photo-editing programs,

Element	'252 Claim Recitation	Exemplary Citations in Reference
		start from a single object pixel and repeatedly add adjacent pixels whose values are within some given threshold of the original pixel. Selecting the threshold, which must change from image to image and object to object, is tedious. We automatically calculate a dynamic threshold by having the user click on background as well as object points. For reasonably uniform objects that are distinct from the background, this operation allows fast object identification without manually adjusting a threshold. The example in Figure 3 shows an object, a fox, identified by using only a few clicks. We designed another outlining tool to help users track object edges. This tool takes a user-drawn curve and automatically aligns it with nearby image edges. Based on the 'snakes' concept developed in recent computer vision research, the tool finds the curve that maximizes the image gradient magnitude along the curve. The spline snake formulation we use allows for smooth solutions to the resulting nonlinear minimization problem. The computation is done at interactive speeds so that, as the user draws a curve, it is 'rubber-banded' to lie along object boundaries. . . . **Video data** For video data, database population has three major components: • shot detection, • representative frame creation for each shot, and • derivation of a layered representation of coherently moving structures/ objects. . . . Once the shot boundaries have been detected, each shot is represented using an r-frame. R-frames are used for several purposes. First, during database population, r-frames are treated as still images in which objects can be identified by using the previously described methods." *Id.* at pp. 3–6. Image, Audio, and Video Extenders Administration and Programming (IBM 000001–617): "Your applications can even search for images by content. Imagine an application that uses visual examples to search for images. With such an application, users could select an example image and have the application find other images that have colors or textures similar to those in the example. With DB2 extenders' Query by Image Content (QBIC) capability, you can create applications that

Element	'252 Claim Recitation	Exemplary Citations in Reference
		search for images in this visual way." *Id.* at p. 4. " A QBIC catalog is a set of files that hold data about the visual features of images. The Image Extender uses this data to search for images by content. You create a QBIC catalog for each column of images in a user table that you want to make available for searching by content. When you create a QBIC catalog you identify the features for which you want the Image Extender to analyze, store, and later query data. You can also add or drop features from a QBIC catalog after the catalog is created. A QBIC catalog can hold data for the following image features: **Average color** The sum of the color values for all pixels in an image divided by the number of pixels in the image. (A pixel is the smallest element of an image that can be assigned color and intensity.) For example, if 50% of an image consists of blue pixels and the other 50% red pixels, the image has an average color value of purple. Average color is used to search for images that have a predominant color. If an image has a predominant color, the average color will be similar to the predominant color. **Histogram color** Measures the distribution of colors in an image against a spectrum of 64 colors. For each of the 64 colors, histogram color identifies the percentage of pixels in an image that have that color. For example, the histogram color of an image might be 40% white pixels, 50% blue, and 10% red; none of the pixels in the image have any of the remaining colors in the histogram spectrum. Histogram color is used to search for images that have a variety of colors. **Positional color** The average color value for the pixels in a specified area in an image. For example, the upper right-hand corner of an image might show a bright yellow sun; the positional color of this area of the image is bright yellow. Positional color is used to search for images that have a predominant color in a particular area. **Texture** Measures the coarseness, contrast, and directionality of an image. Coarseness

Element	'252 Claim Recitation	Exemplary Citations in Reference
		indicates the size of repeating items in an image (for example, pebbles versus boulders). Contrast identifies the brightness variations in an image (light versus dark). Directionality indicates whether a direction predominates in an image (as in the vertical direction of a picket fence) or does not predominate (as in an image of sand). Texture is used to search for images that have a particular pattern. To make an image available for searching by content, you catalog the image. When you catalog an image, the Image Extender analyzes the image, by computing the feature values for the image, and stores the values in a QBIC catalog. When you search for an image by content, your query identifies one or more features for the search (such as average color), a source for each feature (such as an example image), and a target set of cataloged images. The Image Extender computes the feature value of the source and compares it to the cataloged feature values for the target images. It then computes a score that indicates how similar the feature values of the target images are to the source. You can have the Image Extender return the images whose features are most similar to the source. The Image Extender will return the handle of each image and the image score. You can also have the Image Extender return only the score of a single image." *Id.* at pp. 20–21. Flickner Depo at 21

Element	'252 Claim Recitation	Exemplary Citations in Reference
		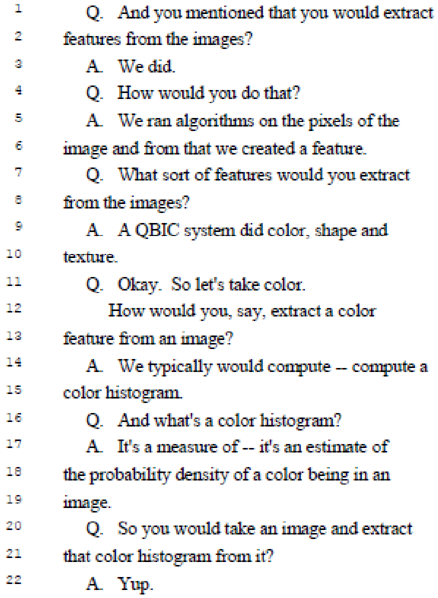 Q. And you mentioned that you would extract features from the images? A. We did. Q. How would you do that? A. We ran algorithms on the pixels of the image and from that we created a feature. Q. What sort of features would you extract from the images? A. A QBIC system did color, shape and texture. Q. Okay. So let's take color. How would you, say, extract a color feature from an image? A. We typically would compute -- compute a color histogram. Q. And what's a color histogram? A. It's a measure of -- it's an estimate of the probability density of a color being in an image. Q. So you would take an image and extract that color histogram from it? A. Yup. Flickner Depo at 212-213

Element	'252 Claim Recitation	Exemplary Citations in Reference
		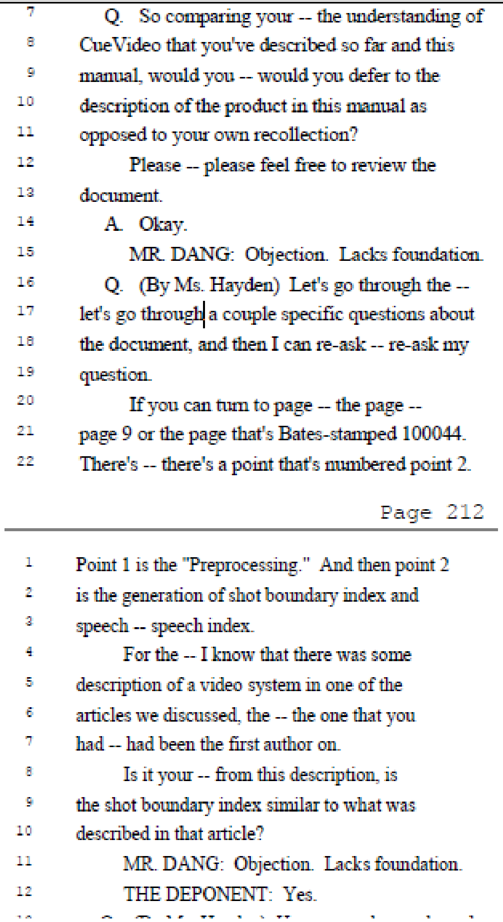 Q. So comparing your -- the understanding of CueVideo that you've described so far and this manual, would you -- would you defer to the description of the product in this manual as opposed to your own recollection? Please -- please feel free to review the document. A. Okay. MR. DANG: Objection. Lacks foundation. Q. (By Ms. Hayden) Let's go through the -- let's go through a couple specific questions about the document, and then I can re-ask -- re-ask my question. If you can turn to page -- the page -- page 9 or the page that's Bates-stamped 100044. There's -- there's a point that's numbered point 2. Page 212 Point 1 is the "Preprocessing." And then point 2 is the generation of shot boundary index and speech -- speech index. For the -- I know that there was some description of a video system in one of the articles we discussed, the -- the one that you had -- had been the first author on. Is it your -- from this description, is the shot boundary index similar to what was described in that article? MR. DANG: Objection. Lacks foundation. THE DEPONENT: Yes. Flickner Depo at 72-73

Element	'252 Claim Recitation	Exemplary Citations in Reference
		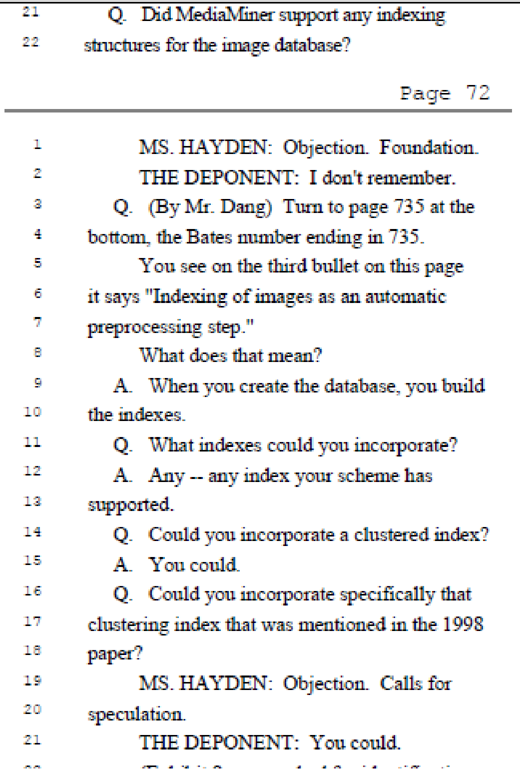 Q. Did MediaMiner support any indexing structures for the image database? Page 72 MS. HAYDEN: Objection. Foundation. THE DEPONENT: I don't remember. Q. (By Mr. Dang) Turn to page 735 at the bottom, the Bates number ending in 735. You see on the third bullet on this page it says "Indexing of images as an automatic preprocessing step." What does that mean? A. When you create the database, you build the indexes. Q. What indexes could you incorporate? A. Any -- any index your scheme has supported. Q. Could you incorporate a clustered index? A. You could. Q. Could you incorporate specifically that clustering index that was mentioned in the 1998 paper? MS. HAYDEN: Objection. Calls for speculation. THE DEPONENT: You could. Flickner Depo at 54-55

Element	'252 Claim Recitation	Exemplary Citations in Reference
		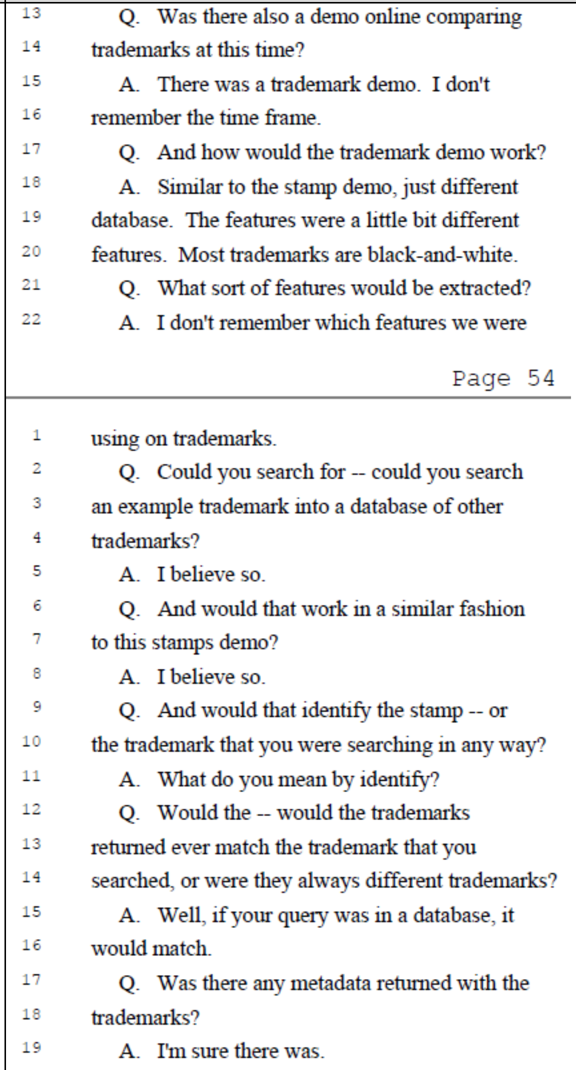 Q. Was there also a demo online comparing trademarks at this time? A. There was a trademark demo. I don't remember the time frame. Q. And how would the trademark demo work? A. Similar to the stamp demo, just different database. The features were a little bit different features. Most trademarks are black-and-white. Q. What sort of features would be extracted? A. I don't remember which features we were Page 54 using on trademarks. Q. Could you search for -- could you search an example trademark into a database of other trademarks? A. I believe so. Q. And would that work in a similar fashion to this stamps demo? A. I believe so. Q. And would that identify the stamp -- or the trademark that you were searching in any way? A. What do you mean by identify? Q. Would the -- would the trademarks returned ever match the trademark that you searched, or were they always different trademarks? A. Well, if your query was in a database, it would match. Q. Was there any metadata returned with the trademarks? A. I'm sure there was.

Element	'252 Claim Recitation	Exemplary Citations in Reference
		Flickner Depo at 89-90 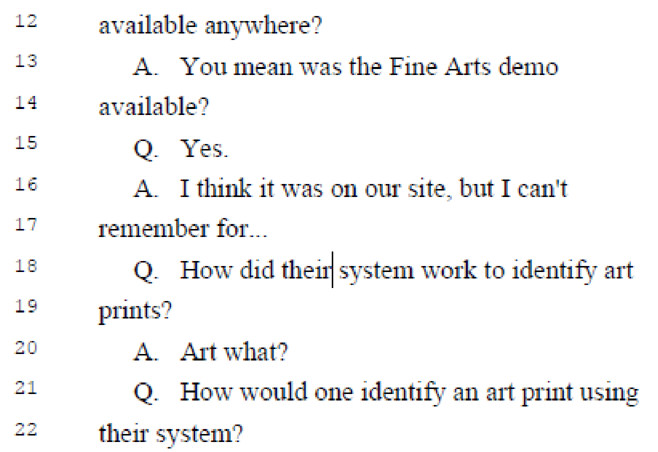Q. Was that used to the QBIC system publicly available anywhere? A. You mean was the Fine Arts demo available? Q. Yes. A. I think it was on our site, but I can't remember for... Q. How did their system work to identify art prints? A. Art what? Q. How would one identify an art print using their system? 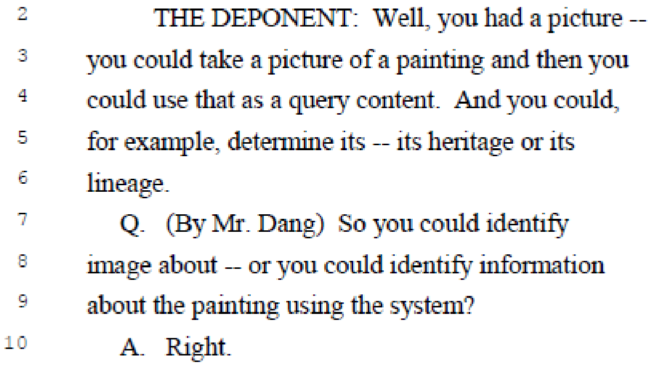THE DEPONENT: Well, you had a picture -- you could take a picture of a painting and then you could use that as a query content. And you could, for example, determine its -- its heritage or its lineage. Q. (By Mr. Dang) So you could identify image about -- or you could identify information about the painting using the system? A. Right.

Element	'252 Claim Recitation	Exemplary Citations in Reference
		Flickner Depo at 88 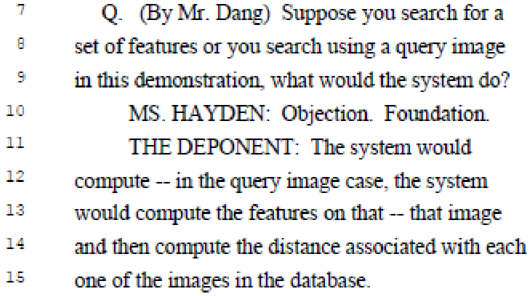 Q. (By Mr. Dang) Suppose you search for a set of features or you search using a query image in this demonstration, what would the system do? MS. HAYDEN: Objection. Foundation. THE DEPONENT: The system would compute -- in the query image case, the system would compute the features on that -- that image and then compute the distance associated with each one of the images in the database. Flickner Depo at 50 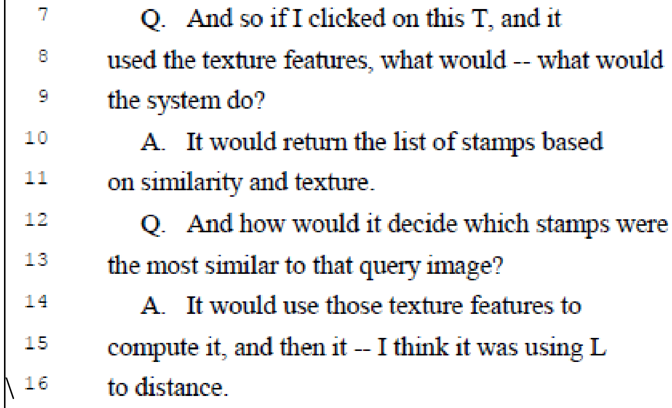 Q. And so if I clicked on this T, and it used the texture features, what would -- what would the system do? A. It would return the list of stamps based on similarity and texture. Q. And how would it decide which stamps were the most similar to that query image? A. It would use those texture features to compute it, and then it -- I think it was using L to distance. Flickner Depo at 78-79

Element	'252 Claim Recitation	Exemplary Citations in Reference
		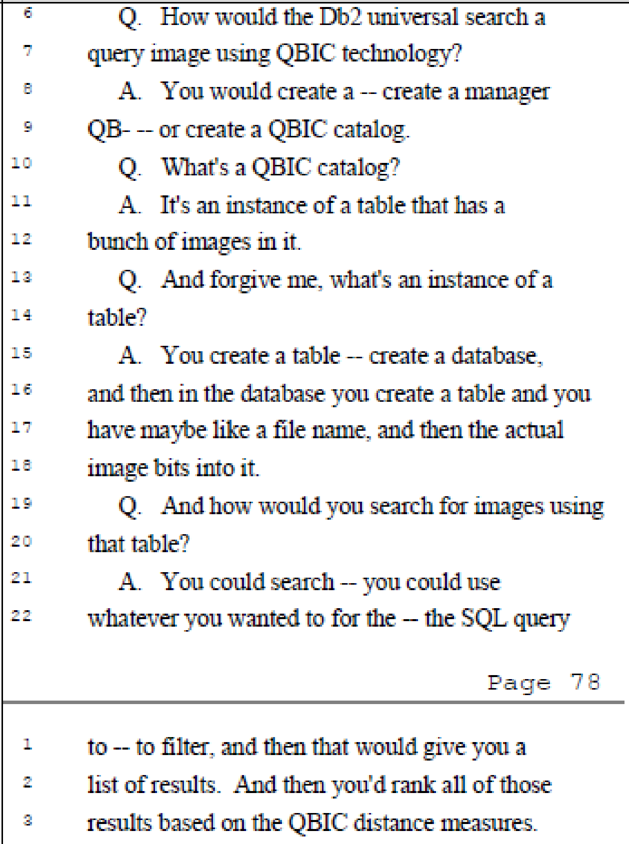 Q. How would the Db2 universal search a query image using QBIC technology? A. You would create a -- create a manager QB- -- or create a QBIC catalog. Q. What's a QBIC catalog? A. It's an instance of a table that has a bunch of images in it. Q. And forgive me, what's an instance of a table? A. You create a table -- create a database, and then in the database you create a table and you have maybe like a file name, and then the actual image bits into it. Q. And how would you search for images using that table? A. You could search -- you could use whatever you wanted to for the -- the SQL query Page 78 to -- to filter, and then that would give you a list of results. And then you'd rank all of those results based on the QBIC distance measures. Flickner Depo at 201-202

Element	'252 Claim Recitation	Exemplary Citations in Reference
		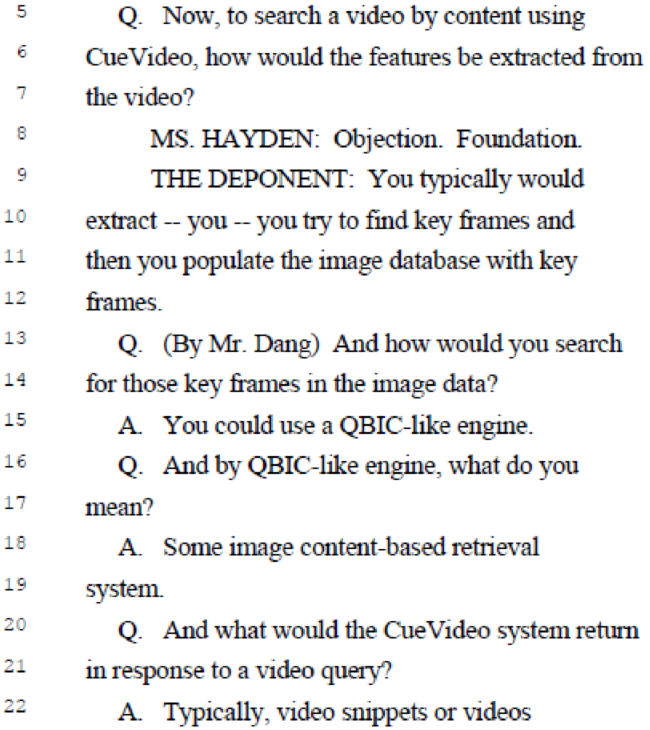 Q. Now, to search a video by content using CueVideo, how would the features be extracted from the video? MS. HAYDEN: Objection. Foundation. THE DEPONENT: You typically would extract -- you -- you try to find key frames and then you populate the image database with key frames. Q. (By Mr. Dang) And how would you search for those key frames in the image data? A. You could use a QBIC-like engine. Q. And by QBIC-like engine, what do you mean? A. Some image content-based retrieval system. Q. And what would the CueVideo system return in response to a video query? A. Typically, video snippets or videos

Element	’252 Claim Recitation	Exemplary Citations in Reference
		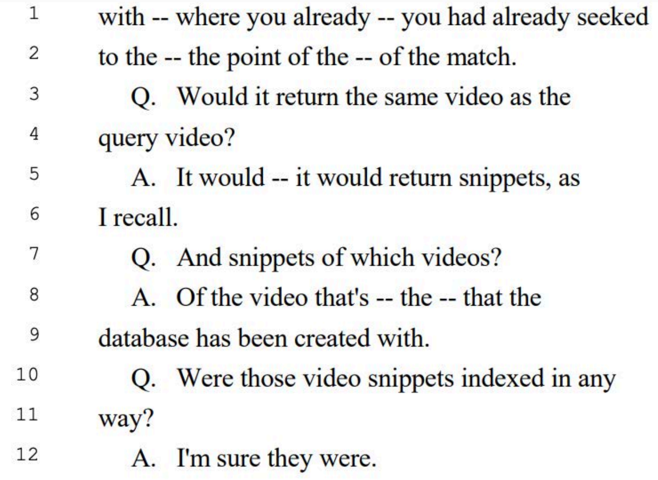 with -- where you already -- you had already seeked to the -- the point of the -- of the match. Q. Would it return the same video as the query video? A. It would -- it would return snippets, as I recall. Q. And snippets of which videos? A. Of the video that's -- the -- that the database has been created with. Q. Were those video snippets indexed in any way? A. I'm sure they were. Flickner Depo at 22-25

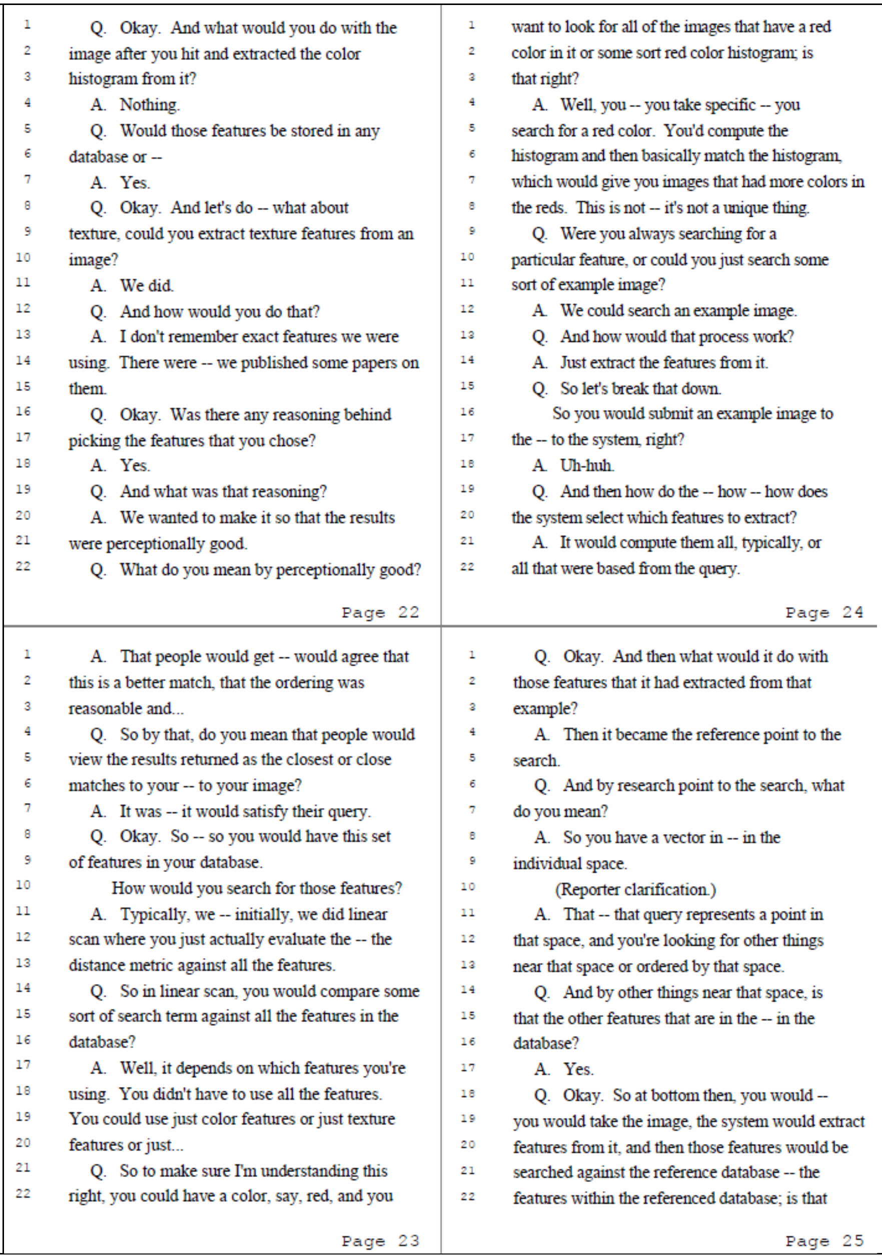

Q. Okay. And what would you do with the image after you hit and extracted the color histogram from it?
A. Nothing.
Q. Would those features be stored in any database or --
A. Yes.
Q. Okay. And let's do -- what about texture, could you extract texture features from an image?
A. We did.
Q. And how would you do that?
A. I don't remember exact features we were using. There were -- we published some papers on them.
Q. Okay. Was there any reasoning behind picking the features that you chose?
A. Yes.
Q. And what was that reasoning?
A. We wanted to make it so that the results were perceptionally good.
Q. What do you mean by perceptionally good?

Page 22

A. That people would get -- would agree that this is a better match, that the ordering was reasonable and...
Q. So by that, do you mean that people would view the results returned as the closest or close matches to your -- to your image?
A. It was -- it would satisfy their query.
Q. Okay. So -- so you would have this set of features in your database.
How would you search for those features?
A. Typically, we -- initially, we did linear scan where you just actually evaluate the -- the distance metric against all the features.
Q. So in linear scan, you would compare some sort of search term against all the features in the database?
A. Well, it depends on which features you're using. You didn't have to use all the features. You could use just color features or just texture features or just...
Q. So to make sure I'm understanding this right, you could have a color, say, red, and you

Page 23

want to look for all of the images that have a red color in it or some sort red color histogram; is that right?
A. Well, you -- you take specific -- you search for a red color. You'd compute the histogram and then basically match the histogram, which would give you images that had more colors in the reds. This is not -- it's not a unique thing.
Q. Were you always searching for a particular feature, or could you just search some sort of example image?
A. We could search an example image.
Q. And how would that process work?
A. Just extract the features from it.
Q. So let's break that down.
So you would submit an example image to the -- to the system, right?
A. Uh-huh.
Q. And then how do the -- how -- how does the system select which features to extract?
A. It would compute them all, typically, or all that were based from the query.

Page 24

Q. Okay. And then what would it do with those features that it had extracted from that example?
A. Then it became the reference point to the search.
Q. And by research point to the search, what do you mean?
A. So you have a vector in -- in the individual space.
(Reporter clarification.)
A. That -- that query represents a point in that space, and you're looking for other things near that space or ordered by that space.
Q. And by other things near that space, is that the other features that are in the -- in the database?
A. Yes.
Q. Okay. So at bottom then, you would -- you would take the image, the system would extract features from it, and then those features would be searched against the reference database -- the features within the referenced database; is that

Page 25

Element	'252 Claim Recitation	Exemplary Citations in Reference
		Flickner Depo at 159 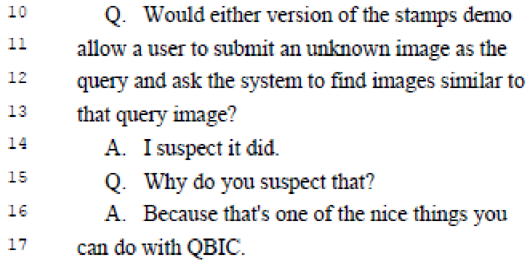 Q. Would either version of the stamps demo allow a user to submit an unknown image as the query and ask the system to find images similar to that query image? A. I suspect it did. Q. Why do you suspect that? A. Because that's one of the nice things you can do with QBIC. Flickner Depo at 209-210 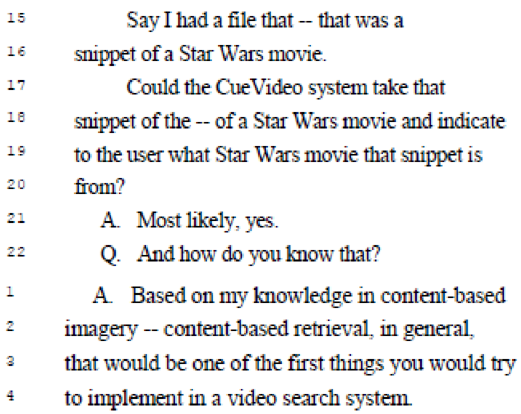 Q. Okay. Let me ask -- let me ask that question using -- using an example. Say I had a file that -- that was a snippet of a Star Wars movie. Could the CueVideo system take that snippet of the -- of a Star Wars movie and indicate to the user what Star Wars movie that snippet is from? A. Most likely, yes. Q. And how do you know that? A. Based on my knowledge in content-based imagery -- content-based retrieval, in general, that would be one of the first things you would try to implement in a video search system. Flickner Depo at 36

Element	'252 Claim Recitation	Exemplary Citations in Reference
		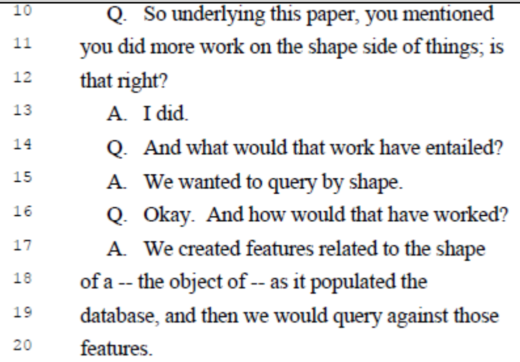 Q. So underlying this paper, you mentioned you did more work on the shape side of things; is that right? A. I did. Q. And what would that work have entailed? A. We wanted to query by shape. Q. Okay. And how would that have worked? A. We created features related to the shape of a -- the object of -- as it populated the database, and then we would query against those features. Flickner Depo at 38-41

Element	'252 Claim Recitation	Exemplary Citations in Reference
		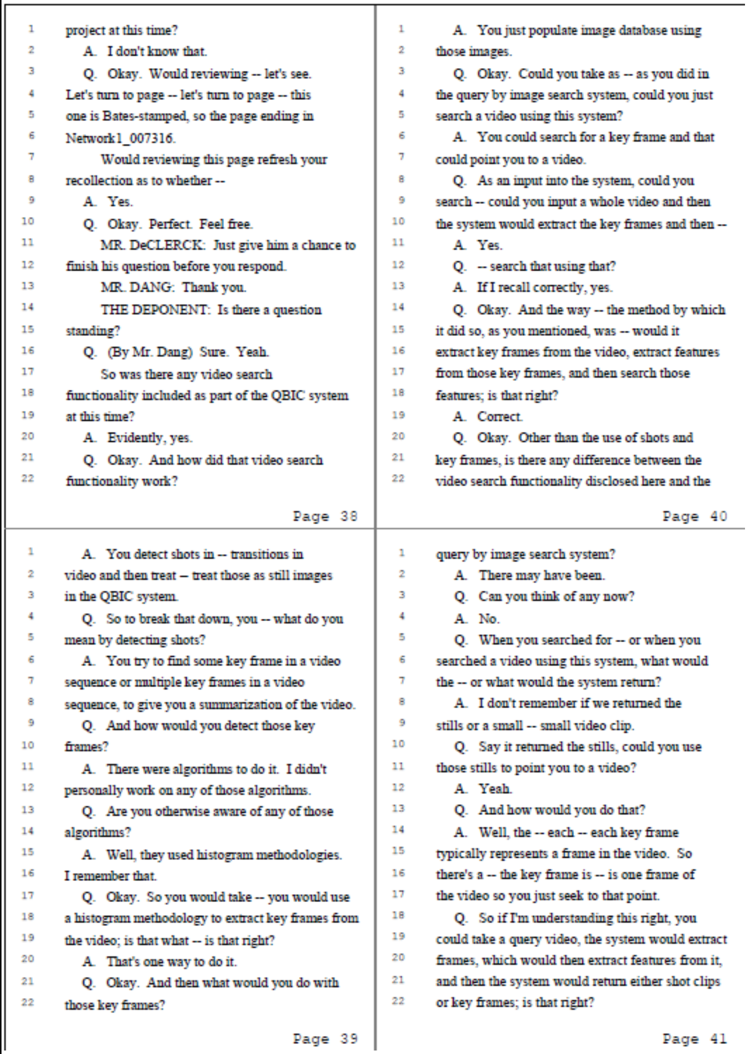

project at this time?
A. I don't know that.
Q. Okay. Would reviewing -- let's see. Let's turn to page -- let's turn to page -- this one is Bates-stamped, so the page ending in Network1_007316.
Would reviewing this page refresh your recollection as to whether --
A. Yes.
Q. Okay. Perfect. Feel free.
MR. DeCLERCK: Just give him a chance to finish his question before you respond.
MR. DANG: Thank you.
THE DEPONENT: Is there a question standing?
Q. (By Mr. Dang) Sure. Yeah.
So was there any video search functionality included as part of the QBIC system at this time?
A. Evidently, yes.
Q. Okay. And how did that video search functionality work?

Page 38

A. You detect shots in -- transitions in video and then treat -- treat those as still images in the QBIC system.
Q. So to break that down, you -- what do you mean by detecting shots?
A. You try to find some key frame in a video sequence or multiple key frames in a video sequence, to give you a summarization of the video.
Q. And how would you detect those key frames?
A. There were algorithms to do it. I didn't personally work on any of those algorithms.
Q. Are you otherwise aware of any of those algorithms?
A. Well, they used histogram methodologies. I remember that.
Q. Okay. So you would take -- you would use a histogram methodology to extract key frames from the video; is that what -- is that right?
A. That's one way to do it.
Q. Okay. And then what would you do with those key frames?

Page 39

A. You just populate image database using those images.
Q. Okay. Could you take as -- as you did in the query by image search system, could you just search a video using this system?
A. You could search for a key frame and that could point you to a video.
Q. As an input into the system, could you search -- could you input a whole video and then the system would extract the key frames and then --
A. Yes.
Q. -- search that using that?
A. If I recall correctly, yes.
Q. Okay. And the way -- the method by which it did so, as you mentioned, was -- would it extract key frames from the video, extract features from those key frames, and then search those features; is that right?
A. Correct.
Q. Okay. Other than the use of shots and key frames, is there any difference between the video search functionality disclosed here and the

Page 40

query by image search system?
A. There may have been.
Q. Can you think of any now?
A. No.
Q. When you searched for -- or when you searched a video using this system, what would the -- or what would the system return?
A. I don't remember if we returned the stills or a small -- small video clip.
Q. Say it returned the stills, could you use those stills to point you to a video?
A. Yeah.
Q. And how would you do that?
A. Well, the -- each -- each key frame typically represents a frame in the video. So there's a -- the key frame is -- is one frame of the video so you just seek to that point.
Q. So if I'm understanding this right, you could take a query video, the system would extract frames, which would then extract features from it, and then the system would return either shot clips or key frames; is that right?

Page 41

Element	'252 Claim Recitation	Exemplary Citations in Reference
		See also Flickner Depo at 40-42, 87. *See also* Flickner Depo at, e.g., 24, 41-42, 62-63 (user could take a picture or video and submit it to the QBIC system). Flickner 8/29/23 Depo at 79–80

Element	'252 Claim Recitation	Exemplary Citations in Reference
		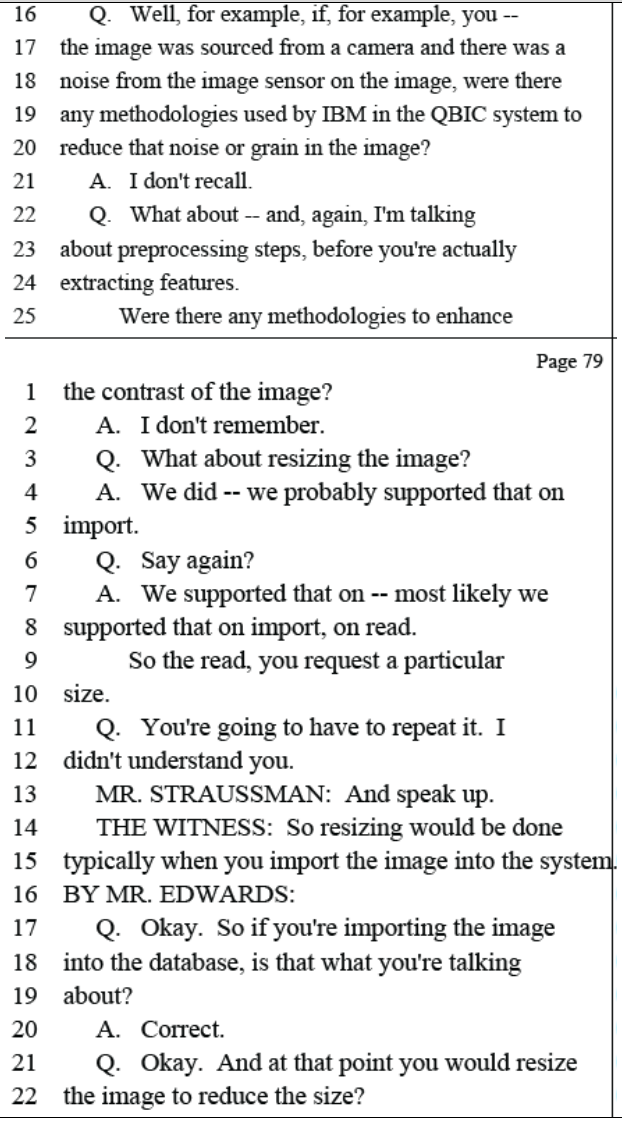 Q. Well, for example, if, for example, you -- the image was sourced from a camera and there was a noise from the image sensor on the image, were there any methodologies used by IBM in the QBIC system to reduce that noise or grain in the image? A. I don't recall. Q. What about -- and, again, I'm talking about preprocessing steps, before you're actually extracting features. Page 79 Were there any methodologies to enhance the contrast of the image? A. I don't remember. Q. What about resizing the image? A. We did -- we probably supported that on import. Q. Say again? A. We supported that on -- most likely we supported that on import, on read. So the read, you request a particular size. Q. You're going to have to repeat it. I didn't understand you. MR. STRAUSSMAN: And speak up. THE WITNESS: So resizing would be done typically when you import the image into the system. BY MR. EDWARDS: Q. Okay. So if you're importing the image into the database, is that what you're talking about? A. Correct. Q. Okay. And at that point you would resize the image to reduce the size?

Element	'252 Claim Recitation	Exemplary Citations in Reference
		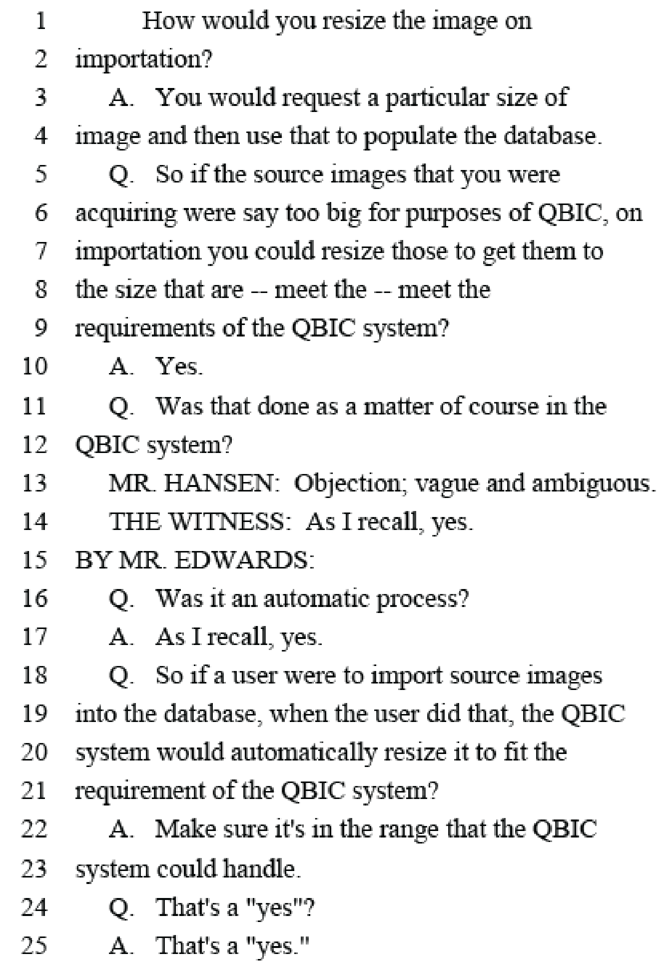 How would you resize the image on importation? A. You would request a particular size of image and then use that to populate the database. Q. So if the source images that you were acquiring were say too big for purposes of QBIC, on importation you could resize those to get them to the size that are -- meet the -- meet the requirements of the QBIC system? A. Yes. Q. Was that done as a matter of course in the QBIC system? MR. HANSEN: Objection; vague and ambiguous. THE WITNESS: As I recall, yes. BY MR. EDWARDS: Q. Was it an automatic process? A. As I recall, yes. Q. So if a user were to import source images into the database, when the user did that, the QBIC system would automatically resize it to fit the requirement of the QBIC system? A. Make sure it's in the range that the QBIC system could handle. Q. That's a "yes"? A. That's a "yes." *Id.* at 80–81

Element	'252 Claim Recitation	Exemplary Citations in Reference
		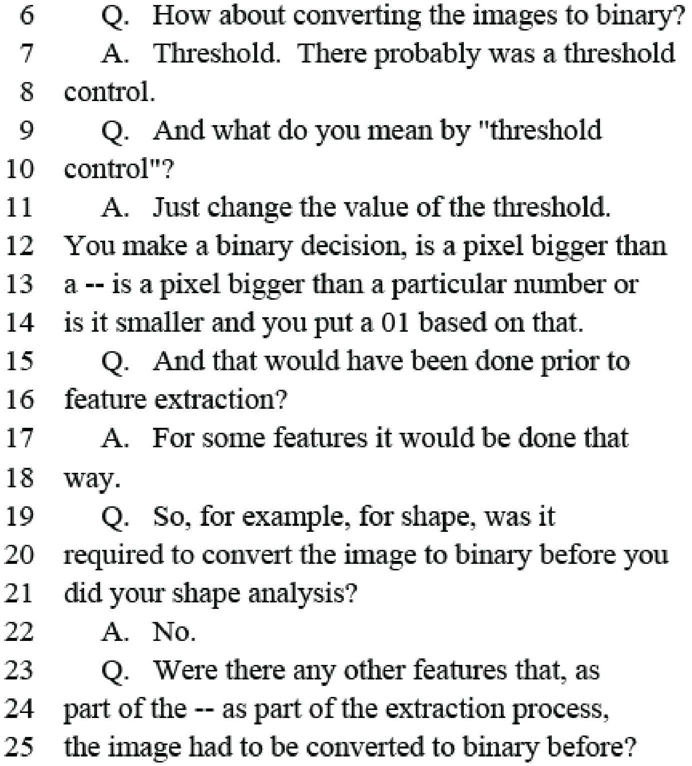 Q. How about converting the images to binary? A. Threshold. There probably was a threshold control. Q. And what do you mean by "threshold control"? A. Just change the value of the threshold. You make a binary decision, is a pixel bigger than a -- is a pixel bigger than a particular number or is it smaller and you put a 01 based on that. Q. And that would have been done prior to feature extraction? A. For some features it would be done that way. Q. So, for example, for shape, was it required to convert the image to binary before you did your shape analysis? A. No. Q. Were there any other features that, as part of the -- as part of the extraction process, the image had to be converted to binary before?

Element	'252 Claim Recitation	Exemplary Citations in Reference
		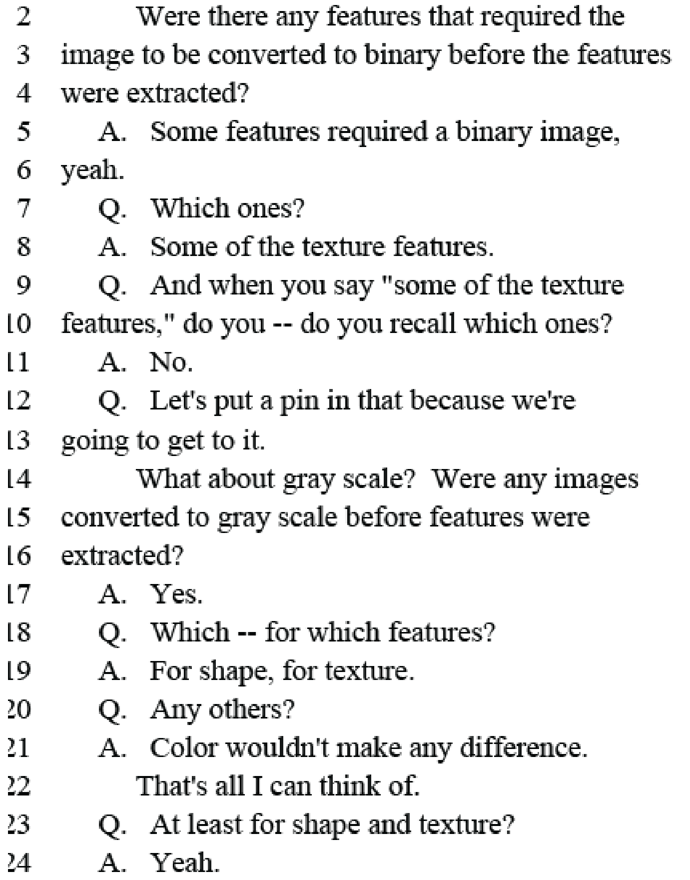 Were there any features that required the image to be converted to binary before the features were extracted? A. Some features required a binary image, yeah. Q. Which ones? A. Some of the texture features. Q. And when you say "some of the texture features," do you -- do you recall which ones? A. No. Q. Let's put a pin in that because we're going to get to it. What about gray scale? Were any images converted to gray scale before features were extracted? A. Yes. Q. Which -- for which features? A. For shape, for texture. Q. Any others? A. Color wouldn't make any difference. That's all I can think of. Q. At least for shape and texture? A. Yeah. *Id.* at 94–97 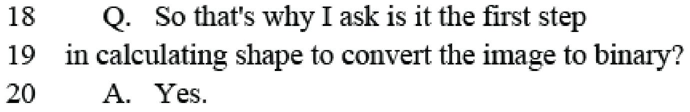 Q. So that's why I ask is it the first step in calculating shape to convert the image to binary? A. Yes. *Id.* at 101–102

Element	'252 Claim Recitation	Exemplary Citations in Reference
		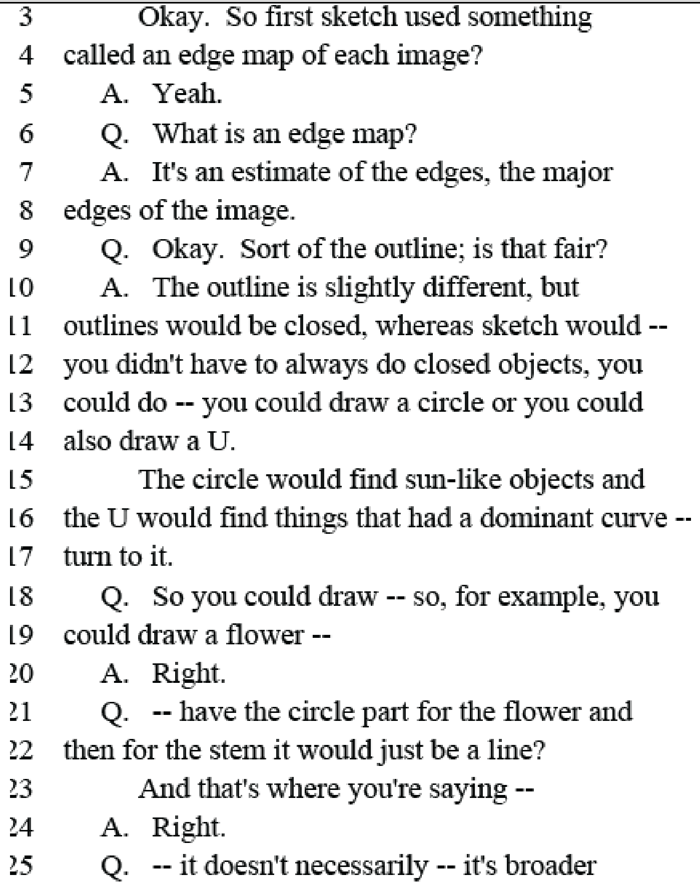 Okay. So first sketch used something called an edge map of each image? A. Yeah. Q. What is an edge map? A. It's an estimate of the edges, the major edges of the image. Q. Okay. Sort of the outline; is that fair? A. The outline is slightly different, but outlines would be closed, whereas sketch would -- you didn't have to always do closed objects, you could do -- you could draw a circle or you could also draw a U. The circle would find sun-like objects and the U would find things that had a dominant curve -- turn to it. Q. So you could draw -- so, for example, you could draw a flower -- A. Right. Q. -- have the circle part for the flower and then for the stem it would just be a line? And that's where you're saying -- A. Right. Q. -- it doesn't necessarily -- it's broader

Element	'252 Claim Recitation	Exemplary Citations in Reference
		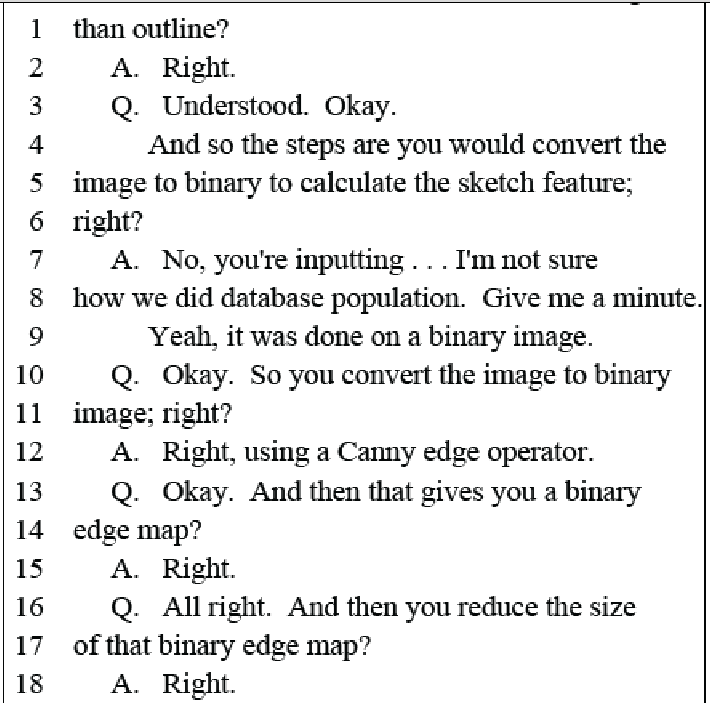 than outline? A. Right. Q. Understood. Okay. And so the steps are you would convert the image to binary to calculate the sketch feature; right? A. No, you're inputting . . . I'm not sure how we did database population. Give me a minute. Yeah, it was done on a binary image. Q. Okay. So you convert the image to binary image; right? A. Right, using a Canny edge operator. Q. Okay. And then that gives you a binary edge map? A. Right. Q. All right. And then you reduce the size of that binary edge map? A. Right. *Id.* at 140–141

Element	'252 Claim Recitation	Exemplary Citations in Reference
		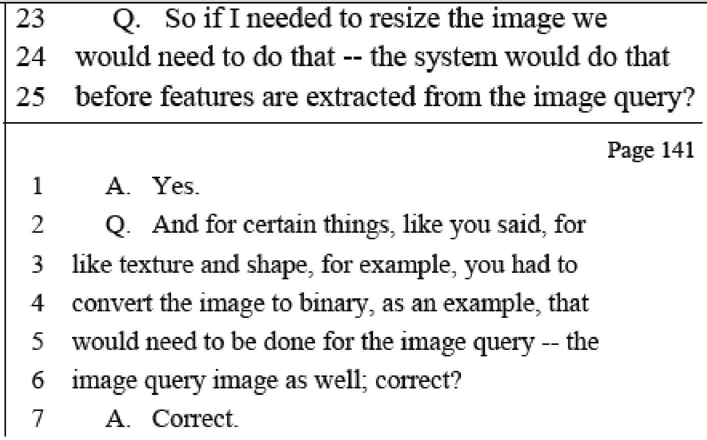 Q. So if I needed to resize the image we would need to do that -- the system would do that before features are extracted from the image query? Page 141 A. Yes. Q. And for certain things, like you said, for like texture and shape, for example, you had to convert the image to binary, as an example, that would need to be done for the image query -- the image query image as well; correct? A. Correct. *Id*. at 164–165

Element	'252 Claim Recitation	Exemplary Citations in Reference
		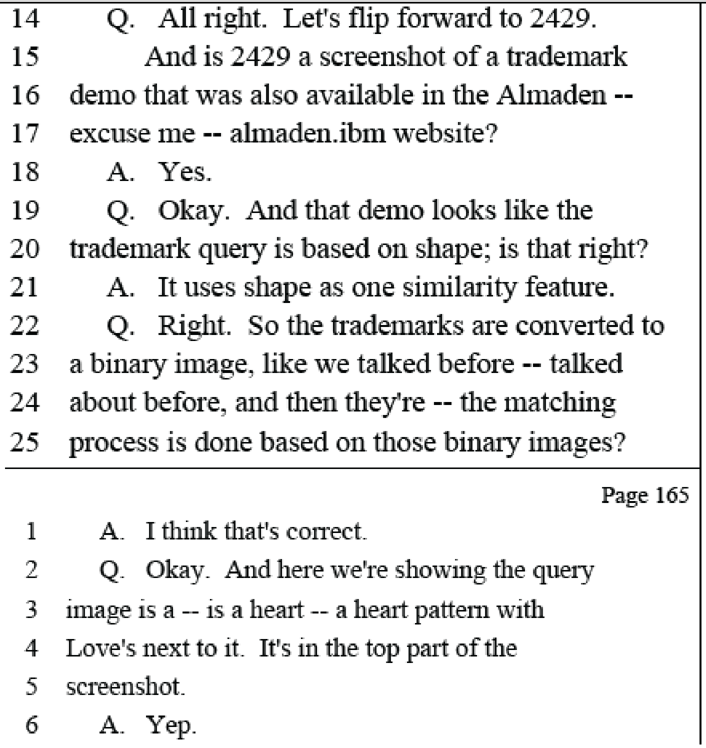 Q. All right. Let's flip forward to 2429. And is 2429 a screenshot of a trademark demo that was also available in the Almaden -- excuse me -- almaden.ibm website? A. Yes. Q. Okay. And that demo looks like the trademark query is based on shape; is that right? A. It uses shape as one similarity feature. Q. Right. So the trademarks are converted to a binary image, like we talked before -- talked about before, and then they're -- the matching process is done based on those binary images? Page 165 A. I think that's correct. Q. Okay. And here we're showing the query image is a -- is a heart -- a heart pattern with Love's next to it. It's in the top part of the screenshot. A. Yep. *Id*. at 82 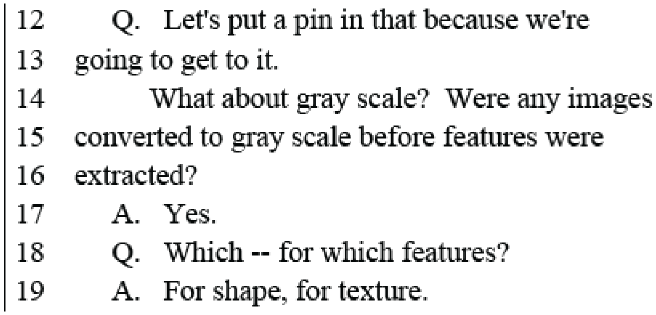 Q. Let's put a pin in that because we're going to get to it. What about gray scale? Were any images converted to gray scale before features were extracted? A. Yes. Q. Which -- for which features? A. For shape, for texture.

Element	'252 Claim Recitation	Exemplary Citations in Reference
		To the extent it is found that QBIC does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that QBIC does not anticipate this claim, QBIC renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine QBIC with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results. For example, the '252 patent acknowledges that acquiring information about objects through symbols, such as barcodes or some other means to encode detectable information in the image or on the object, were traditional methods known in the art and not part the present invention. '252 Patent, 3:7-16. Indeed, the '252 patent acknowledges that algorithms, software, and hardware components were commercially available at the time of the invention that could detect and decode symbols, such as barcodes or text, in an input image. *Id*. at 11:56-61. This limitation was thus known in the art by patentee's own admission, and by asserting this claim, Plaintiff is attempting to reclaim disavowed claim scope and claim subject matter which patentee has acknowledged is not part of the invention.
[18B.i]	operating on the data relating to the image to determine if the image contains one or more recognizable symbols; and	QBIC discloses operating on the data relating to the image to determine if the image contains one or more recognizable symbols.

Element	'252 Claim Recitation	Exemplary Citations in Reference
		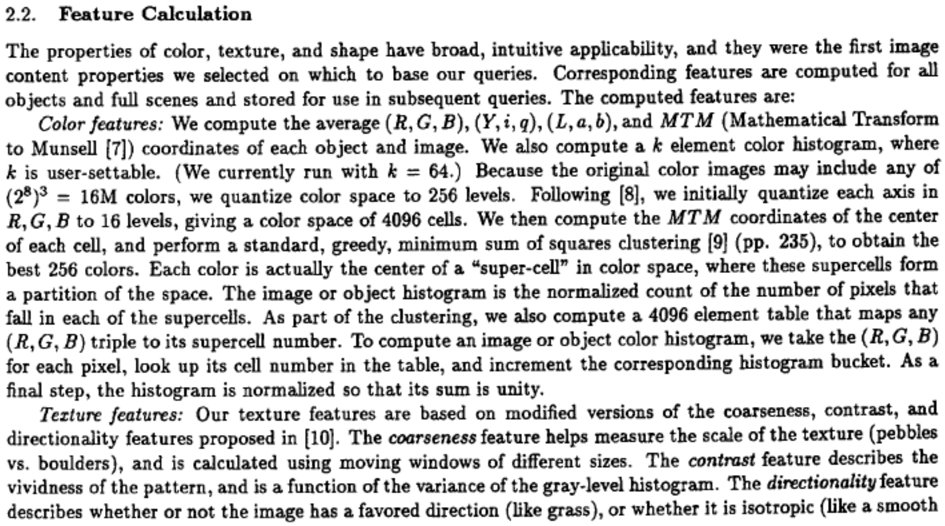 2.2. **Feature Calculation** The properties of color, texture, and shape have broad, intuitive applicability, and they were the first image content properties we selected on which to base our queries. Corresponding features are computed for all objects and full scenes and stored for use in subsequent queries. The computed features are: *Color features:* We compute the average (R, G, B), (Y, i, q), (L, a, b), and MTM (Mathematical Transform to Munsell [7]) coordinates of each object and image. We also compute a k element color histogram, where k is user-settable. (We currently run with $k = 64$.) Because the original color images may include any of $(2^8)^3 = 16$M colors, we quantize color space to 256 levels. Following [8], we initially quantize each axis in R, G, B to 16 levels, giving a color space of 4096 cells. We then compute the MTM coordinates of the center of each cell, and perform a standard, greedy, minimum sum of squares clustering [9] (pp. 235), to obtain the best 256 colors. Each color is actually the center of a "super-cell" in color space, where these supercells form a partition of the space. The image or object histogram is the normalized count of the number of pixels that fall in each of the supercells. As part of the clustering, we also compute a 4096 element table that maps any (R, G, B) triple to its supercell number. To compute an image or object color histogram, we take the (R, G, B) for each pixel, look up its cell number in the table, and increment the corresponding histogram bucket. As a final step, the histogram is normalized so that its sum is unity. *Texture features:* Our texture features are based on modified versions of the coarseness, contrast, and directionality features proposed in [10]. The *coarseness* feature helps measure the scale of the texture (pebbles vs. boulders), and is calculated using moving windows of different sizes. The *contrast* feature describes the vividness of the pattern, and is a function of the variance of the gray-level histogram. The *directionality* feature describes whether or not the image has a favored direction (like grass), or whether it is isotropic (like a smooth

Element	'252 Claim Recitation	Exemplary Citations in Reference
		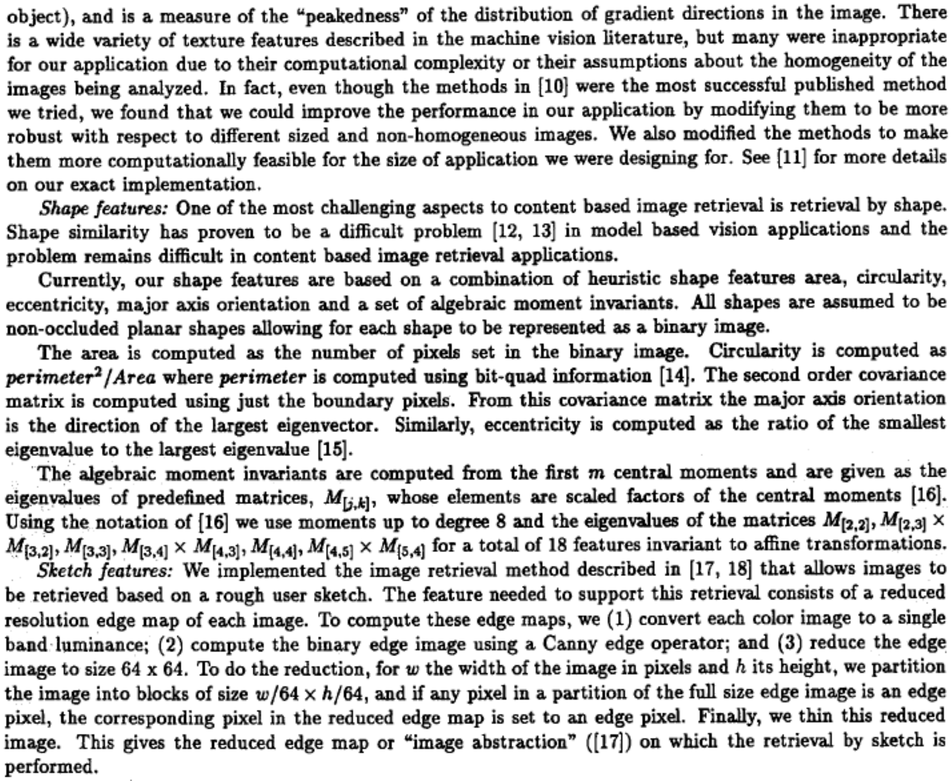 object), and is a measure of the "peakedness" of the distribution of gradient directions in the image. There is a wide variety of texture features described in the machine vision literature, but many were inappropriate for our application due to their computational complexity or their assumptions about the homogeneity of the images being analyzed. In fact, even though the methods in [10] were the most successful published method we tried, we found that we could improve the performance in our application by modifying them to be more robust with respect to different sized and non-homogeneous images. We also modified the methods to make them more computationally feasible for the size of application we were designing for. See [11] for more details on our exact implementation. *Shape features:* One of the most challenging aspects to content based image retrieval is retrieval by shape. Shape similarity has proven to be a difficult problem [12, 13] in model based vision applications and the problem remains difficult in content based image retrieval applications. Currently, our shape features are based on a combination of heuristic shape features area, circularity, eccentricity, major axis orientation and a set of algebraic moment invariants. All shapes are assumed to be non-occluded planar shapes allowing for each shape to be represented as a binary image. The area is computed as the number of pixels set in the binary image. Circularity is computed as $perimeter^2/Area$ where $perimeter$ is computed using bit-quad information [14]. The second order covariance matrix is computed using just the boundary pixels. From this covariance matrix the major axis orientation is the direction of the largest eigenvector. Similarly, eccentricity is computed as the ratio of the smallest eigenvalue to the largest eigenvalue [15]. The algebraic moment invariants are computed from the first m central moments and are given as the eigenvalues of predefined matrices, $M_{[j,k]}$, whose elements are scaled factors of the central moments [16]. Using the notation of [16] we use moments up to degree 8 and the eigenvalues of the matrices $M_{[2,2]}, M_{[2,3]} \times M_{[3,2]}, M_{[3,3]}, M_{[3,4]} \times M_{[4,3]}, M_{[4,4]}, M_{[4,5]} \times M_{[5,4]}$ for a total of 18 features invariant to affine transformations. *Sketch features:* We implemented the image retrieval method described in [17, 18] that allows images to be retrieved based on a rough user sketch. The feature needed to support this retrieval consists of a reduced resolution edge map of each image. To compute these edge maps, we (1) convert each color image to a single band luminance; (2) compute the binary edge image using a Canny edge operator; and (3) reduce the edge image to size 64 x 64. To do the reduction, for w the width of the image in pixels and h its height, we partition the image into blocks of size $w/64 \times h/64$, and if any pixel in a partition of the full size edge image is an edge pixel, the corresponding pixel in the reduced edge map is set to an edge pixel. Finally, we thin this reduced image. This gives the reduced edge map or "image abstraction" ([17]) on which the retrieval by sketch is performed. IBM0002390 at 2392-93

Element	'252 Claim Recitation	Exemplary Citations in Reference
		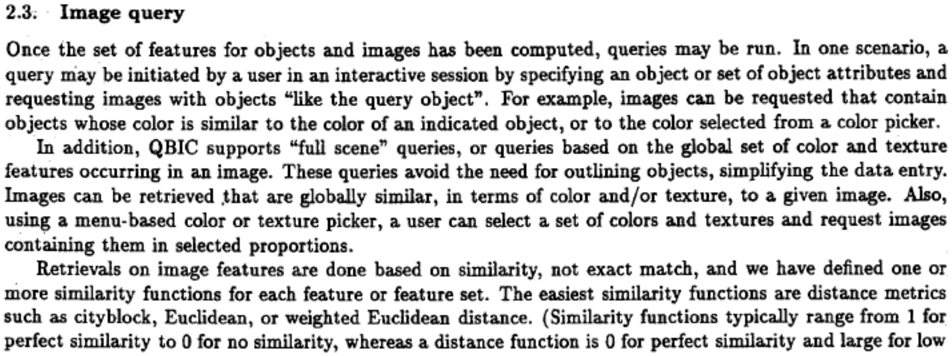 **2.3. Image query** Once the set of features for objects and images has been computed, queries may be run. In one scenario, a query may be initiated by a user in an interactive session by specifying an object or set of object attributes and requesting images with objects "like the query object". For example, images can be requested that contain objects whose color is similar to the color of an indicated object, or to the color selected from a color picker. In addition, QBIC supports "full scene" queries, or queries based on the global set of color and texture features occurring in an image. These queries avoid the need for outlining objects, simplifying the data entry. Images can be retrieved that are globally similar, in terms of color and/or texture, to a given image. Also, using a menu-based color or texture picker, a user can select a set of colors and textures and request images containing them in selected proportions. Retrievals on image features are done based on similarity, not exact match, and we have defined one or more similarity functions for each feature or feature set. The easiest similarity functions are distance metrics such as cityblock, Euclidean, or weighted Euclidean distance. (Similarity functions typically range from 1 for perfect similarity to 0 for no similarity, whereas a distance function is 0 for perfect similarity and large for low

Element	'252 Claim Recitation	Exemplary Citations in Reference
		similarity, so we must convert as necessary to obtain consistent measures.) Different similarity measures may be used, and they may vary from application to application. The matching methods and similarity functions for the different feature types are described below: *Color:* For average color, the distance between a query object and database object is weighted Euclidean distance. The weights are user adjustable so that we may, for example, downweight the L component of an (L, a, b) query. Best results, as judged by our subjective evaluation in on-going experiments, are obtained by representing the average color in the MTM color space, and by inversely weighting each component by its standard deviation over the samples in the database. Improved results are obtained when color matching is done not on average color, but on the distribution of colors that occur in an object or image. In this way, a red and blue object matches other red and blue objects better than to a purple object. To match color histograms, we use a method similar to the one described in [8]. Let X be the query histogram and Y the histogram of an item in the database, both normalized. We compute the element by element difference histogram Z. Then the similarity between X and Y is given by $\|Z\| = Z^T A Z$, where A is a symmetric color similarity matrix with $a(i,j) \equiv 1 - d(c_i, c_j)/d_{max}$ where c_i and c_j are the ith and jth colors in the color histograms, and $d(c_i, c_j)$ is the MTM color distance, and d_{max} is the maximum distance between any two colors ([8] used Luv distance.) This metric gives the (weighted) length of the difference vector between X and Y, weighted by A, which accounts for the perceptual distance between different pairs of colors. The result is a measure that accounts for both the difference between the amounts in any given color (e.g. a particular shade of red), as well as differences between similar colors (e.g. red and orange). Based on our experience, we also found it useful for the user to be able to request images with $x\%$ of color 1, $y\%$ of color 2, etc. where both the percentages and the colors are explicitly specified by the user. This allows searches such as for a beach scene (25% white (for the sand) and 50% blue (for the sky and water)), and for many scenes for which the user has seen before, and for which he/she can remember approximate color percentages. The color picker that allows this type of query is shown in Figure 2. 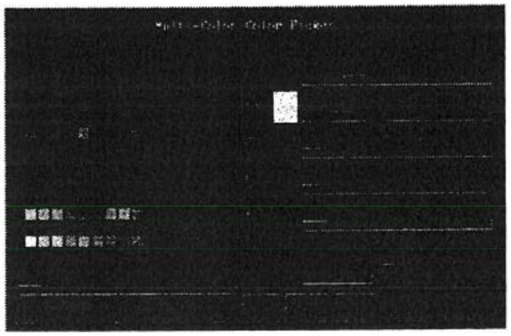 Figure 2: Multi-Color Color Picker with two selected colors. The palette on the left displays the set of selectable colors. Once a color is selected, it appears in a box on the right, and the amount of the color can be specified by the sliders. In this case, 25% light blue and 15% white have been selected. *Texture:* Texture distance is computed as weighted Euclidean distance in the three dimensional texture space. The most common weighting factors are the inverse variances for each component, computed over the

Element	'252 Claim Recitation	Exemplary Citations in Reference
		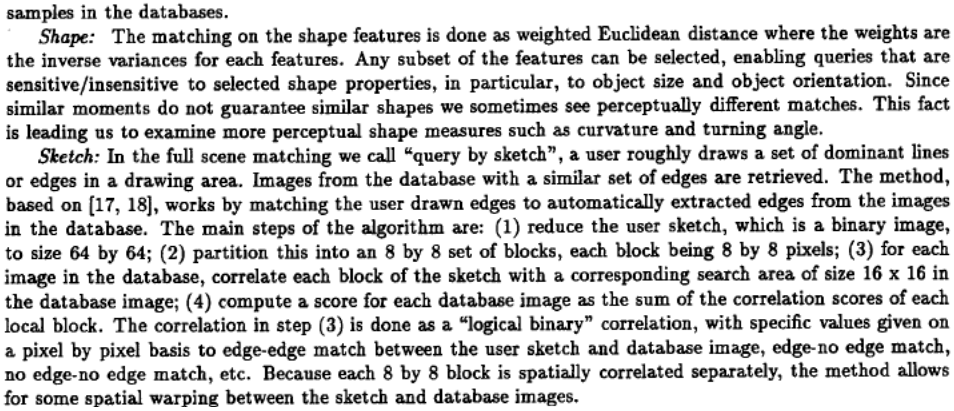 samples in the databases. *Shape:* The matching on the shape features is done as weighted Euclidean distance where the weights are the inverse variances for each features. Any subset of the features can be selected, enabling queries that are sensitive/insensitive to selected shape properties, in particular, to object size and object orientation. Since similar moments do not guarantee similar shapes we sometimes see perceptually different matches. This fact is leading us to examine more perceptual shape measures such as curvature and turning angle. *Sketch:* In the full scene matching we call "query by sketch", a user roughly draws a set of dominant lines or edges in a drawing area. Images from the database with a similar set of edges are retrieved. The method, based on [17, 18], works by matching the user drawn edges to automatically extracted edges from the images in the database. The main steps of the algorithm are: (1) reduce the user sketch, which is a binary image, to size 64 by 64; (2) partition this into an 8 by 8 set of blocks, each block being 8 by 8 pixels; (3) for each image in the database, correlate each block of the sketch with a corresponding search area of size 16 x 16 in the database image; (4) compute a score for each database image as the sum of the correlation scores of each local block. The correlation in step (3) is done as a "logical binary" correlation, with specific values given on a pixel by pixel basis to edge-edge match between the user sketch and database image, edge-no edge match, no edge-no edge match, etc. Because each 8 by 8 block is spatially correlated separately, the method allows for some spatial warping between the sketch and database images. IBM0002390 at 2393-95 Query by Image and Video Content: The QBIC System (IBM 0000893–902): "We have developed the QBIC (Query by Image Content) system to explore content-based retrieval methods. QBIC allows queries on large image and video databases based on • example images, • user-constructed sketches and drawings, • selected color and texture patterns, • camera and object motion, and • other graphical information." *Id*. at pp. 1–3. "To achieve this functionality, QBIC has two main components: database population (the process of creating an image database) and database query. During the population, images and videos are processed to extract features describing their content-colors, textures, shapes, and camera and object motion-and the features are stored in a database. During the query, the user composes a query graphically. Features are generated from the graphical query and then input to a matching engine that finds images or videos from the database with similar features. Figure 2 shows the system architecture." *Id*. at p. 3.

Element	'252 Claim Recitation	Exemplary Citations in Reference
		Image, Audio, and Video Extenders Administration and Programming (IBM 000001–617): "DB2 (DB2) Universal Database (UDB) is a powerful, object-relational database manager. It stores and protects traditional numeric and character data, as well as large, complex objects (LOBs). DB2 extenders help you exploit DB2 's object-relational features. The extenders define distinct data types and special functions for image, audio, video, and text objects. … Because of that, the extenders give your applications a single point of access to any or all of these types of data, along with traditional numeric and character data. In addition, the extenders give your applications new ways to search for information." *Id*. at p. 3. Managing Enterprise Information Portal (IBM 000777–880): "Many paper-intensive enterprises, such as insurance companies and financial institutions, administer large volumes of business-related content. The need for an enterprise solution for managing and accessing business information spans many industries. A content server is a software system that stores multimedia, business forms, documents, and related data, along with metadata that allows employees to process and work with the content." *Id*. at p. 3. "The image search server uses IBM's QBIC (query by image content) technology to help you search for objects by certain visual properties, such as color and texture. The image search server analyzes images and stores the image information in a database. Then users can run image queries, which use the visual properties of images, to match colors, textures, and their positions without describing them in words. You can combine content-based queries with text and keyword searches for more powerful retrieval of image and multimedia data." *Id*. at p. 63. QBIC Demo (IBM000747)

Element	'252 Claim Recitation	Exemplary Citations in Reference
		 Location: http://172.16.85.176:3456/cgi-bin/ibm.tcl Keywords: france Flickner 8/31/23 Depo at Ex. 6 at 1 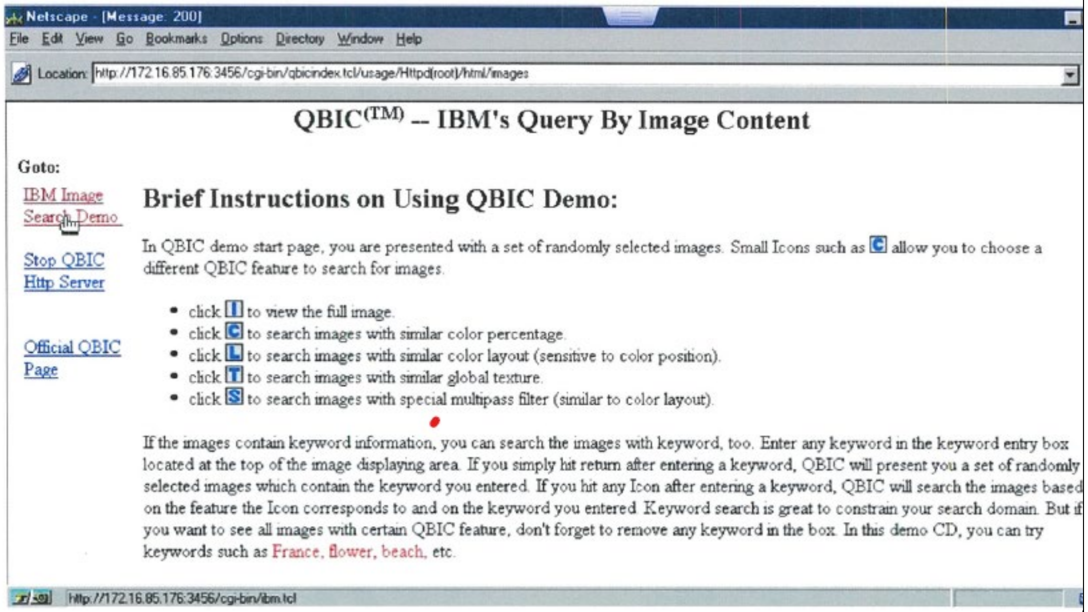 Location: http://172.16.85.176:3456/cgi-bin/qbicindex.tcl/usage/Httpd(root)/html/images QBIC(TM) -- IBM's Query By Image Content Brief Instructions on Using QBIC Demo: In QBIC demo start page, you are presented with a set of randomly selected images. Small Icons such as C allow you to choose a different QBIC feature to search for images. - click I to view the full image. - click C to search images with similar color percentage. - click L to search images with similar color layout (sensitive to color position). - click T to search images with similar global texture. - click S to search images with special multipass filter (similar to color layout). If the images contain keyword information, you can search the images with keyword, too. Enter any keyword in the keyword entry box located at the top of the image displaying area. If you simply hit return after entering a keyword, QBIC will present you a set of randomly selected images which contain the keyword you entered. If you hit any Icon after entering a keyword, QBIC will search the images based on the feature the Icon corresponds to and on the keyword you entered. Keyword search is great to constrain your search domain. But if you want to see all images with certain QBIC feature, don't forget to remove any keyword in the box. In this demo CD, you can try keywords such as France, flower, beach, etc. http://172.16.85.176:3456/cgi-bin/ibm.tcl

Element	'252 Claim Recitation	Exemplary Citations in Reference
		Id. at 2.  Flickner 8/31/23 Depo Ex. 7 at 1

Element	'252 Claim Recitation	Exemplary Citations in Reference
		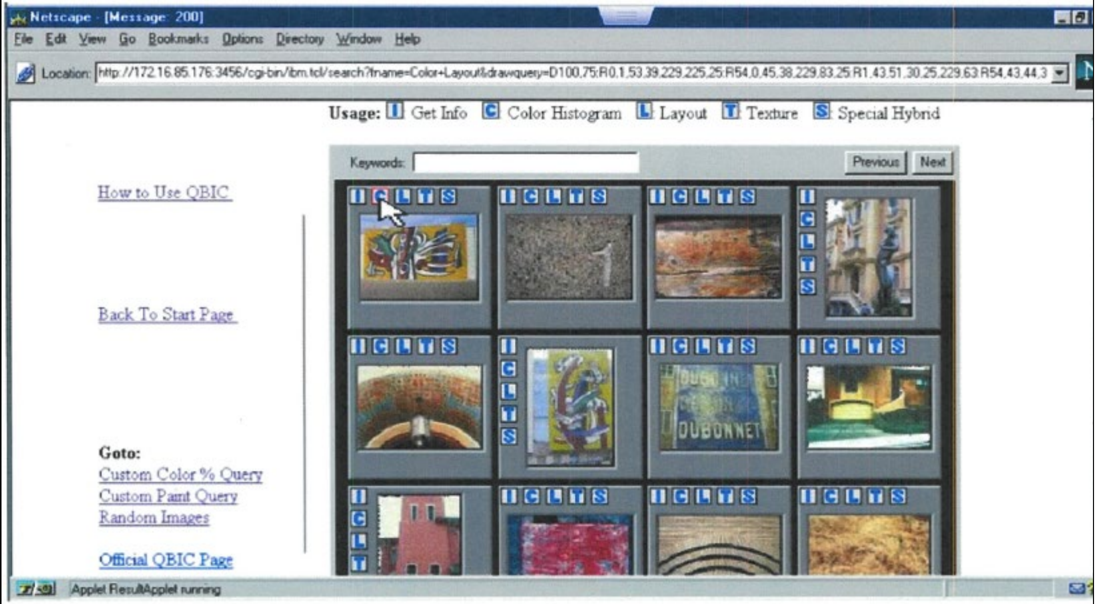 Location: http://172.16.85.176:3456/cgi-bin/ibm.tcl/search?fname=Color+Layout&drawquery=D100,75:R0,1,53,39,229,225,25:R54,0,45,38,229,83,25:R1,43,51,30,25,229,63:R54,43,44,3 *Id.* at 2

Element	'252 Claim Recitation	Exemplary Citations in Reference
		 Flickner 8/31/23 Depo at Ex. 8

Element	'252 Claim Recitation	Exemplary Citations in Reference
		 Location: http://172.16.85.176:3456/cgi-bin/ibm.tcl/search?fname=Color+Histogram&france50%2Ejpg.x=0 *Id.* at 2

Nantworks, LLC v. Bank of America Corporation
Central District of California, No. 2:20-cv-7872
Invalidity Chart for U.S. Patent 7,899,252 based on the Query by Image and Video Content ("QBIC") System

Element	'252 Claim Recitation	Exemplary Citations in Reference
		 Location: http://172.16.85.176:3456/cgi-bin/ibm.tcl/info?france50%2Ejpg Image france50.jpg *Id.* at 3

Nantworks, LLC v. Bank of America Corporation
Central District of California, No. 2:20-cv-7872
Invalidity Chart for U.S. Patent 7,899,252 based on the Query by Image and Video Content ("QBIC") System

Element	'252 Claim Recitation	Exemplary Citations in Reference
		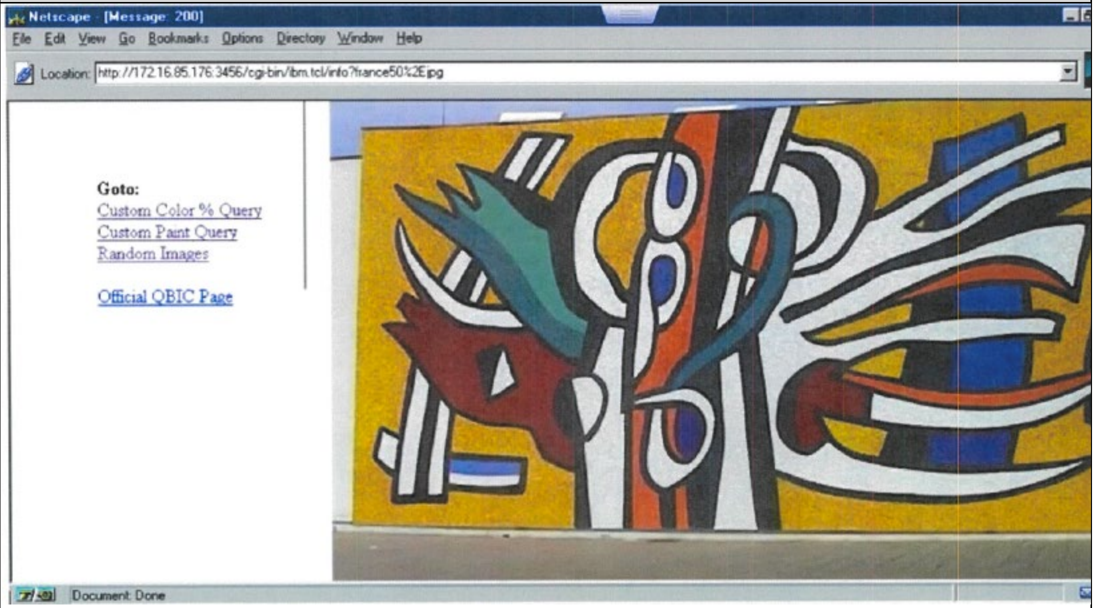 Location: http://172.16.85.176:3456/cgi-bin/ibm.tcl/info?france50%2Ejpg Flickner Depo at 162 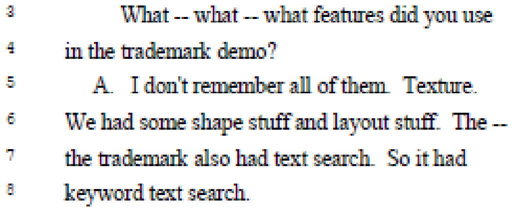 What -- what -- what features did you use in the trademark demo? A. I don't remember all of them. Texture. We had some shape stuff and layout stuff. The -- the trademark also had text search. So it had keyword text search. Flickner Depo at 217

Element	'252 Claim Recitation	Exemplary Citations in Reference
		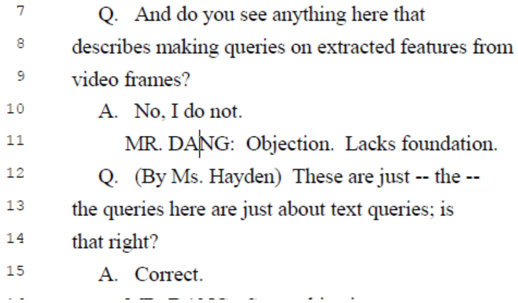 Q. And do you see anything here that describes making queries on extracted features from video frames? A. No, I do not. MR. DANG: Objection. Lacks foundation. Q. (By Ms. Hayden) These are just -- the -- the queries here are just about text queries; is that right? A. Correct. Flickner Depo at 78-79

Element	'252 Claim Recitation	Exemplary Citations in Reference
		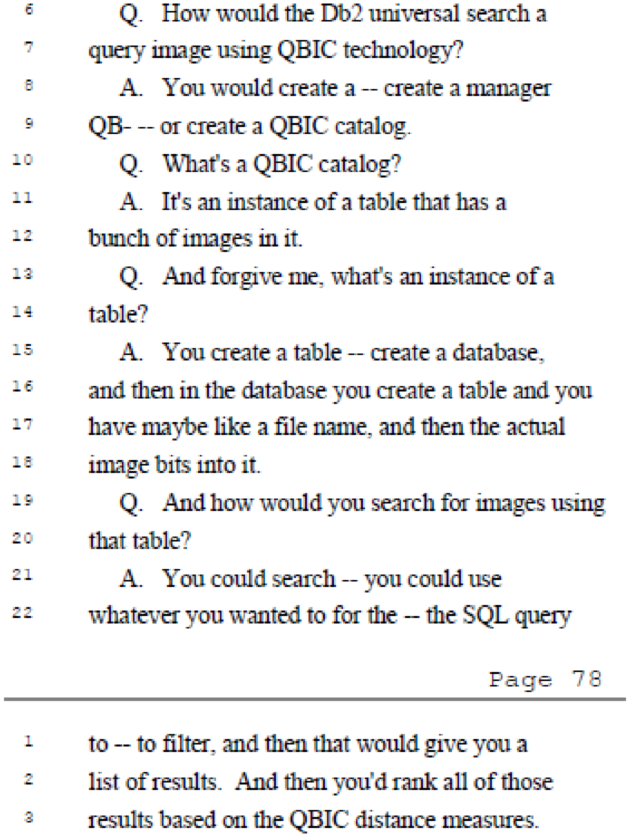 Q. How would the Db2 universal search a query image using QBIC technology? A. You would create a -- create a manager QB- -- or create a QBIC catalog. Q. What's a QBIC catalog? A. It's an instance of a table that has a bunch of images in it. Q. And forgive me, what's an instance of a table? A. You create a table -- create a database, and then in the database you create a table and you have maybe like a file name, and then the actual image bits into it. Q. And how would you search for images using that table? A. You could search -- you could use whatever you wanted to for the -- the SQL query Page 78 to -- to filter, and then that would give you a list of results. And then you'd rank all of those results based on the QBIC distance measures. Flickner Depo at 185

Element	'252 Claim Recitation	Exemplary Citations in Reference
		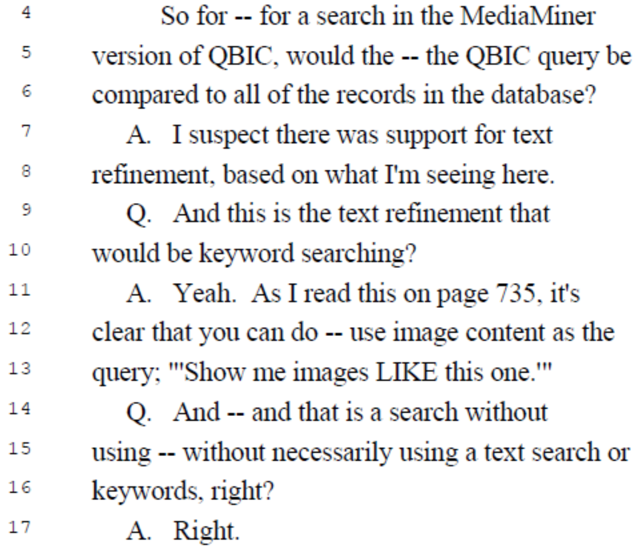 So for -- for a search in the MediaMiner version of QBIC, would the -- the QBIC query be compared to all of the records in the database? A. I suspect there was support for text refinement, based on what I'm seeing here. Q. And this is the text refinement that would be keyword searching? A. Yeah. As I read this on page 735, it's clear that you can do -- use image content as the query; "'Show me images LIKE this one.'" Q. And -- and that is a search without using -- without necessarily using a text search or keywords, right? A. Right.

Element	'252 Claim Recitation	Exemplary Citations in Reference
		Flickner Depo at 214-215 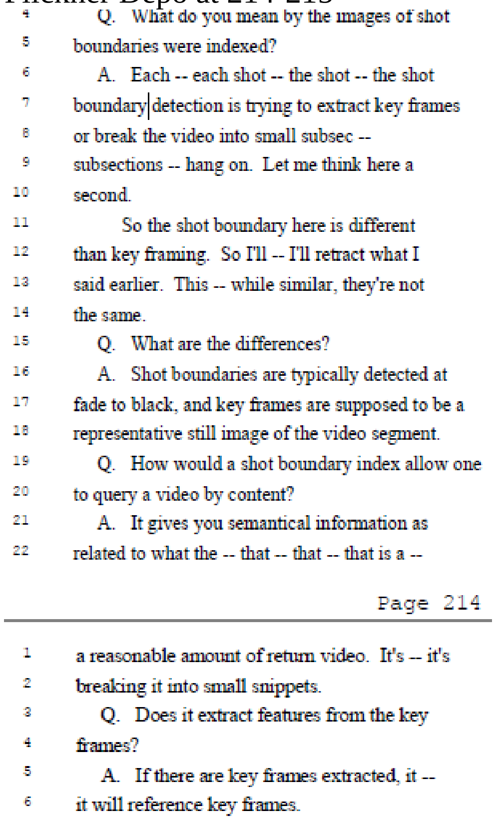 Q. What do you mean by the images of shot boundaries were indexed? A. Each -- each shot -- the shot -- the shot boundary detection is trying to extract key frames or break the video into small subsec -- subsections -- hang on. Let me think here a second. So the shot boundary here is different than key framing. So I'll -- I'll retract what I said earlier. This -- while similar, they're not the same. Q. What are the differences? A. Shot boundaries are typically detected at fade to black, and key frames are supposed to be a representative still image of the video segment. Q. How would a shot boundary index allow one to query a video by content? A. It gives you semantical information as related to what the -- that -- that -- that is a -- Page 214 a reasonable amount of return video. It's -- it's breaking it into small snippets. Q. Does it extract features from the key frames? A. If there are key frames extracted, it -- it will reference key frames. Flickner Depo at 50. Flickner 8/29/23 Depo at 44:

Element	'252 Claim Recitation	Exemplary Citations in Reference
		Q. And how did that problem, going through film stocks being a relatively labor tedious process, how did that play into developing QBIC? A. Well, just like you could query -- you could query text based on court documents or things like that, now you can make visual queries. Q. And querying text was known. People were doing that before the QBIC project started? A. Yes. *Id.* at 45–47 Q. And going back to your testimony about querying text or keywords, does that mean OCR or does that -- does that mean something else? A. You could do OCR. Q. Okay. And OCR -- A. I assume OCR, you mean optical character recognition? Q. Correct. So that could include OCR? A. Yeah. Q. And that was a known technology at the time the QBIC project started? A. Yes. MR. EDWARDS: We'll mark Exhibit 2. (Deposition Exhibit 2 was marked.) BY MR. EDWARDS: Q. And, Mr. Flickner, before we get to Exhibit 2, did -- going back to the OCR, did the QBIC system itself do OCR? A. No. Q. It did not. Did the QBIC system ever

Element	'252 Claim Recitation	Exemplary Citations in Reference
		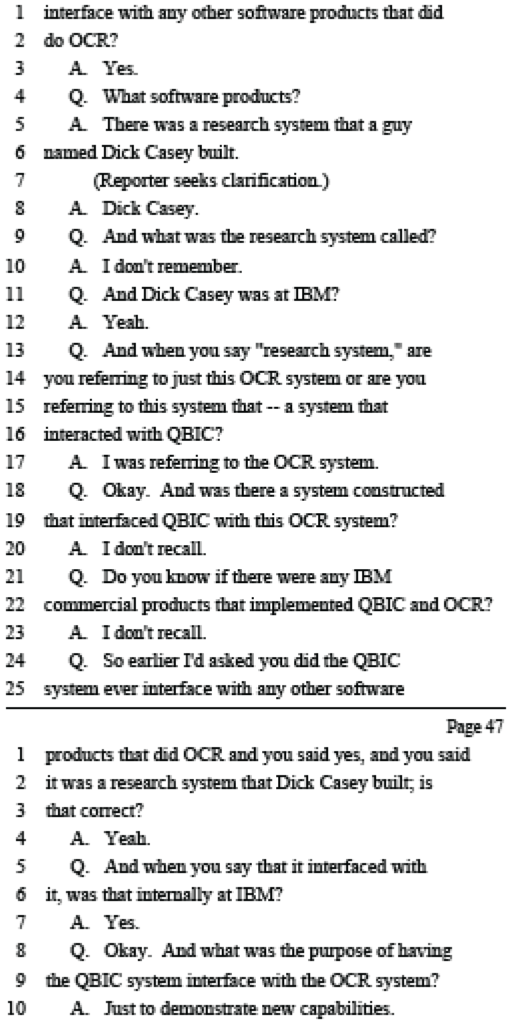 interface with any other software products that did do OCR? A. Yes. Q. What software products? A. There was a research system that a guy named Dick Casey built. (Reporter seeks clarification.) A. Dick Casey. Q. And what was the research system called? A. I don't remember. Q. And Dick Casey was at IBM? A. Yeah. Q. And when you say "research system," are you referring to just this OCR system or are you referring to this system that -- a system that interacted with QBIC? A. I was referring to the OCR system. Q. Okay. And was there a system constructed that interfaced QBIC with this OCR system? A. I don't recall. Q. Do you know if there were any IBM commercial products that implemented QBIC and OCR? A. I don't recall. Q. So earlier I'd asked you did the QBIC system ever interface with any other software Page 47 products that did OCR and you said yes, and you said it was a research system that Dick Casey built; is that correct? A. Yeah. Q. And when you say that it interfaced with it, was that internally at IBM? A. Yes. Q. Okay. And what was the purpose of having the QBIC system interface with the OCR system? A. Just to demonstrate new capabilities. *Id.* at 76-77:

Element	'252 Claim Recitation	Exemplary Citations in Reference
		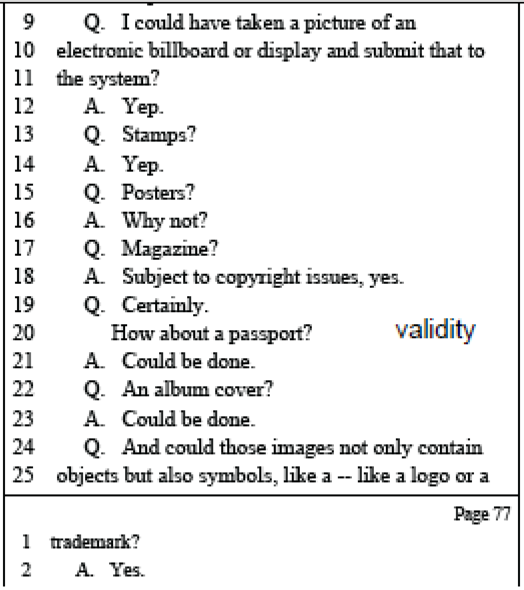 Q. I could have taken a picture of an electronic billboard or display and submit that to the system? A. Yep. Q. Stamps? A. Yep. Q. Posters? A. Why not? Q. Magazine? A. Subject to copyright issues, yes. Q. Certainly. How about a passport? validity A. Could be done. Q. An album cover? A. Could be done. Q. And could those images not only contain objects but also symbols, like a -- like a logo or a Page 77 trademark? A. Yes. *Id*. at 160: 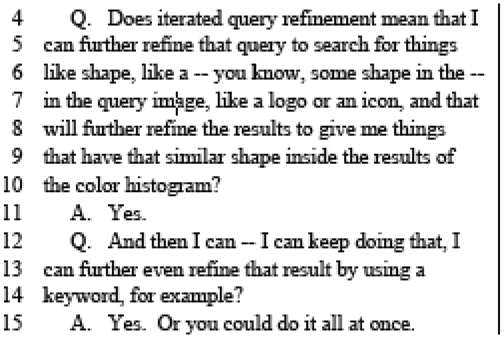 Q. Does iterated query refinement mean that I can further refine that query to search for things like shape, like a -- you know, some shape in the -- in the query image, like a logo or an icon, and that will further refine the results to give me things that have that similar shape inside the results of the color histogram? A. Yes. Q. And then I can -- I can keep doing that, I can further even refine that result by using a keyword, for example? A. Yes. Or you could do it all at once. *Id*. at 163-166:

Element	’252 Claim Recitation	Exemplary Citations in Reference
		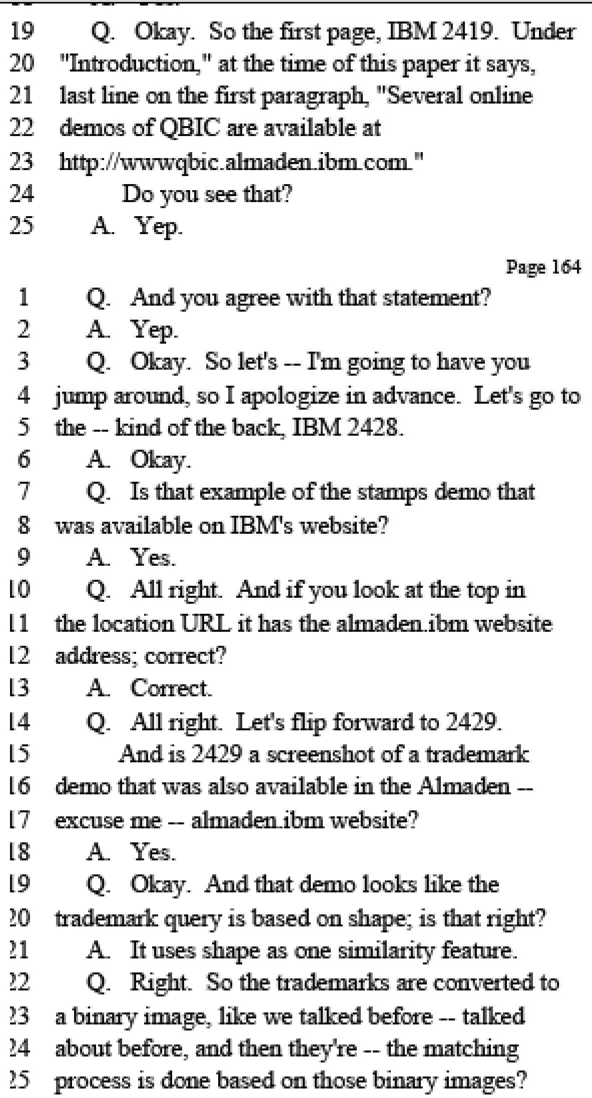 Q. Okay. So the first page, IBM 2419. Under "Introduction," at the time of this paper it says, last line on the first paragraph, "Several online demos of QBIC are available at http://wwwqbic.almaden.ibm.com." Do you see that? A. Yep. Page 164 Q. And you agree with that statement? A. Yep. Q. Okay. So let's -- I'm going to have you jump around, so I apologize in advance. Let's go to the -- kind of the back, IBM 2428. A. Okay. Q. Is that example of the stamps demo that was available on IBM's website? A. Yes. Q. All right. And if you look at the top in the location URL it has the almaden.ibm website address; correct? A. Correct. Q. All right. Let's flip forward to 2429. And is 2429 a screenshot of a trademark demo that was also available in the Almaden -- excuse me -- almaden.ibm website? A. Yes. Q. Okay. And that demo looks like the trademark query is based on shape; is that right? A. It uses shape as one similarity feature. Q. Right. So the trademarks are converted to a binary image, like we talked before -- talked about before, and then they're -- the matching process is done based on those binary images?

Element	'252 Claim Recitation	Exemplary Citations in Reference
		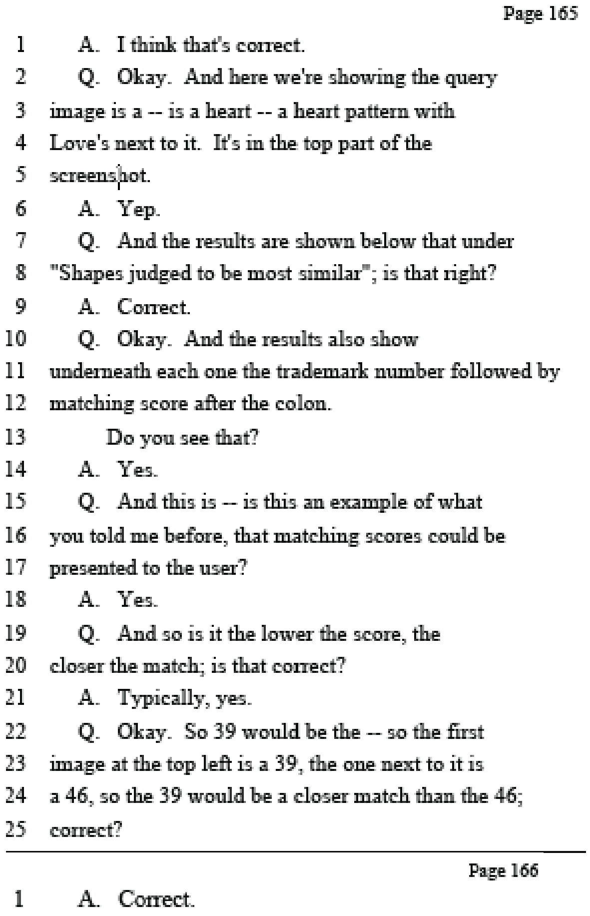 Page 165 A. I think that's correct. Q. Okay. And here we're showing the query image is a -- is a heart -- a heart pattern with Love's next to it. It's in the top part of the screenshot. A. Yep. Q. And the results are shown below that under "Shapes judged to be most similar"; is that right? A. Correct. Q. Okay. And the results also show underneath each one the trademark number followed by matching score after the colon. Do you see that? A. Yes. Q. And this is -- is this an example of what you told me before, that matching scores could be presented to the user? A. Yes. Q. And so is it the lower the score, the closer the match; is that correct? A. Typically, yes. Q. Okay. So 39 would be the -- so the first image at the top left is a 39, the one next to it is a 46, so the 39 would be a closer match than the 46; correct? Page 166 A. Correct. *Id.* at 187–188:

Element	'252 Claim Recitation	Exemplary Citations in Reference
		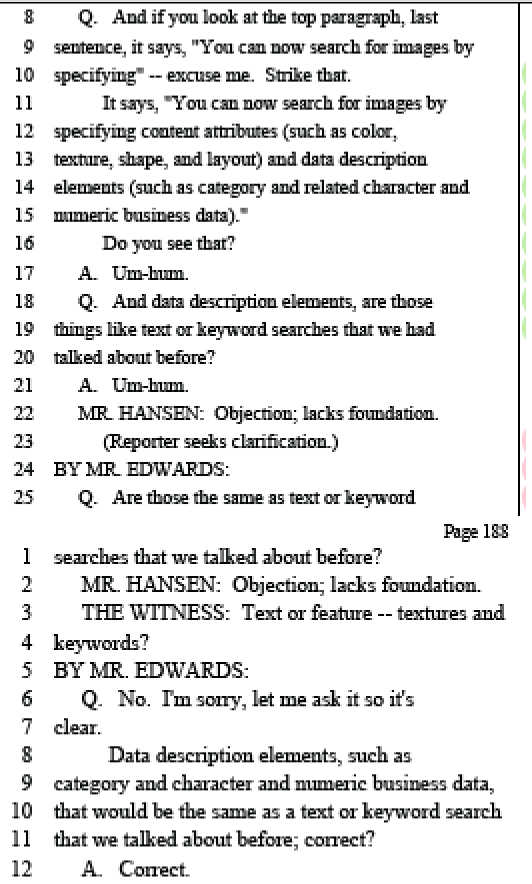 Q. And if you look at the top paragraph, last sentence, it says, "You can now search for images by specifying" -- excuse me. Strike that. It says, "You can now search for images by specifying content attributes (such as color, texture, shape, and layout) and data description elements (such as category and related character and numeric business data)." Do you see that? A. Um-hum. Q. And data description elements, are those things like text or keyword searches that we had talked about before? A. Um-hum. MR. HANSEN: Objection; lacks foundation. (Reporter seeks clarification.) BY MR. EDWARDS: Q. Are those the same as text or keyword Page 188 searches that we talked about before? MR. HANSEN: Objection; lacks foundation. THE WITNESS: Text or feature -- textures and keywords? BY MR. EDWARDS: Q. No. I'm sorry, let me ask it so it's clear. Data description elements, such as category and character and numeric business data, that would be the same as a text or keyword search that we talked about before; correct? A. Correct.

Element	'252 Claim Recitation	Exemplary Citations in Reference
		Flickner 8/29/23 Depo, Ex. 6: 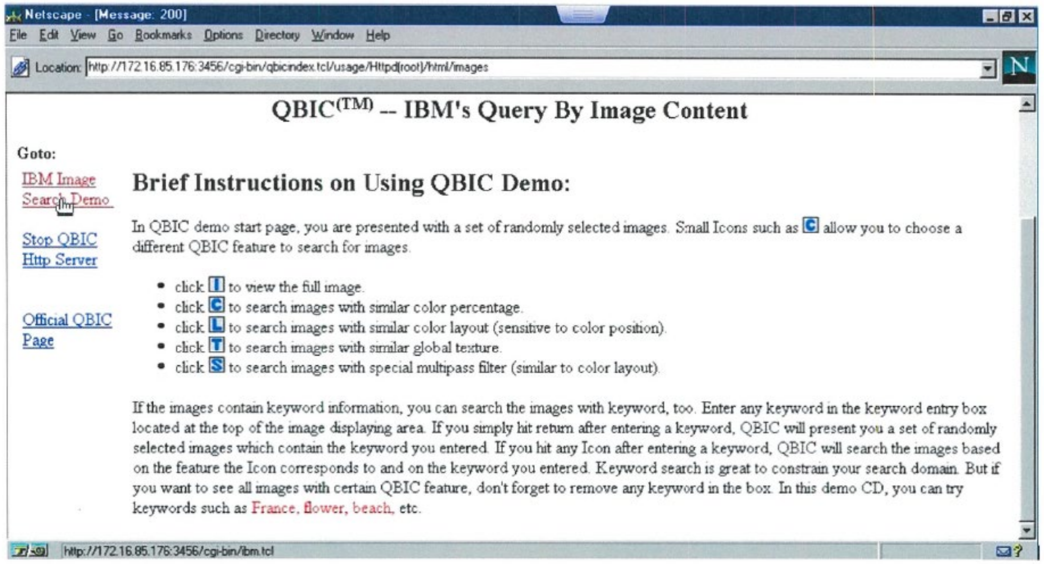 Netscape - [Message: 200] Location: http://172.16.85.176:3456/cgi-bin/qbicindex.tcl/usage/Httpd(root)/html/images QBIC$^{(TM)}$ -- IBM's Query By Image Content Goto: IBM Image Search Demo Stop QBIC Http Server Official QBIC Page **Brief Instructions on Using QBIC Demo:** In QBIC demo start page, you are presented with a set of randomly selected images. Small Icons such as C allow you to choose a different QBIC feature to search for images. - click I to view the full image. - click C to search images with similar color percentage. - click L to search images with similar color layout (sensitive to color position). - click T to search images with similar global texture. - click S to search images with special multipass filter (similar to color layout). If the images contain keyword information, you can search the images with keyword, too. Enter any keyword in the keyword entry box located at the top of the image displaying area. If you simply hit return after entering a keyword, QBIC will present you a set of randomly selected images which contain the keyword you entered. If you hit any Icon after entering a keyword, QBIC will search the images based on the feature the Icon corresponds to and on the keyword you entered. Keyword search is great to constrain your search domain. But if you want to see all images with certain QBIC feature, don't forget to remove any keyword in the box. In this demo CD, you can try keywords such as France, flower, beach, etc. http://172.16.85.176:3456/cgi-bin/ibm.tcl To the extent it is found that QBIC does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that QBIC does not anticipate this claim, QBIC renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine QBIC with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results. For example, the '252 patent acknowledges that acquiring information about objects through symbols, such as barcodes or some other means to encode

Element	'252 Claim Recitation	Exemplary Citations in Reference
		detectable information in the image or on the object, were traditional methods known in the art and not part the present invention. '252 Patent, 3:7-16. Indeed, the '252 patent acknowledges that algorithms, software, and hardware components were commercially available at the time of the invention that could detect and decode symbols, such as barcodes or text, in an input image. *Id.* at 11:56-61. This limitation was thus known in the art by patentee's own admission, and by asserting this claim, Plaintiff is attempting to reclaim disavowed claim scope and claim subject matter which patentee has acknowledged is not part of the invention.
[18B.ii]	decoding the recognizable symbols to extract symbol information by analyzing the recognizable symbols according to type;	QBIC discloses decoding the recognizable symbols to extract symbol information by analyzing the recognizable symbols according to type. 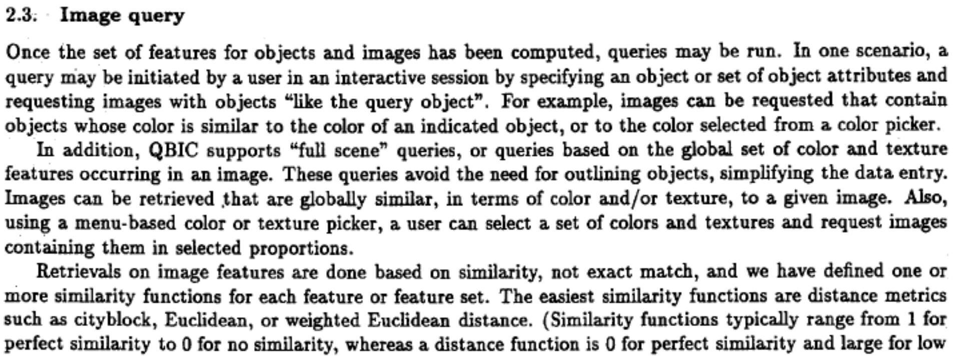 **2.3. Image query** Once the set of features for objects and images has been computed, queries may be run. In one scenario, a query may be initiated by a user in an interactive session by specifying an object or set of object attributes and requesting images with objects "like the query object". For example, images can be requested that contain objects whose color is similar to the color of an indicated object, or to the color selected from a color picker. In addition, QBIC supports "full scene" queries, or queries based on the global set of color and texture features occurring in an image. These queries avoid the need for outlining objects, simplifying the data entry. Images can be retrieved that are globally similar, in terms of color and/or texture, to a given image. Also, using a menu-based color or texture picker, a user can select a set of colors and textures and request images containing them in selected proportions. Retrievals on image features are done based on similarity, not exact match, and we have defined one or more similarity functions for each feature or feature set. The easiest similarity functions are distance metrics such as cityblock, Euclidean, or weighted Euclidean distance. (Similarity functions typically range from 1 for perfect similarity to 0 for no similarity, whereas a distance function is 0 for perfect similarity and large for low

Element	'252 Claim Recitation	Exemplary Citations in Reference
		similarity, so we must convert as necessary to obtain consistent measures.) Different similarity measures may be used, and they may vary from application to application. The matching methods and similarity functions for the different feature types are described below: *Color:* For average color, the distance between a query object and database object is weighted Euclidean distance. The weights are user adjustable so that we may, for example, downweight the L component of an (L, a, b) query. Best results, as judged by our subjective evaluation in on-going experiments, are obtained by representing the average color in the MTM color space, and by inversely weighting each component by its standard deviation over the samples in the database. Improved results are obtained when color matching is done not on average color, but on the distribution of colors that occur in an object or image. In this way, a red and blue object matches other red and blue objects better than to a purple object. To match color histograms, we use a method similar to the one described in [8]. Let X be the query histogram and Y the histogram of an item in the database, both normalized. We compute the element by element difference histogram Z. Then the similarity between X and Y is given by $\|Z\| = Z^T A Z$, where A is a symmetric color similarity matrix with $a(i,j) \equiv 1 - d(c_i, c_j)/d_{max}$ where c_i and c_j are the ith and jth colors in the color histograms, and $d(c_i, c_j)$ is the MTM color distance, and d_{max} is the maximum distance between any two colors ([8] used Luv distance.) This metric gives the (weighted) length of the difference vector between X and Y, weighted by A, which accounts for the perceptual distance between different pairs of colors. The result is a measure that accounts for both the difference between the amounts in any given color (e.g. a particular shade of red), as well as differences between similar colors (e.g. red and orange). Based on our experience, we also found it useful for the user to be able to request images with $x\%$ of color 1, $y\%$ of color 2, etc. where both the percentages and the colors are explicitly specified by the user. This allows searches such as for a beach scene (25% white (for the sand) and 50% blue (for the sky and water)), and for many scenes for which the user has seen before, and for which he/she can remember approximate color percentages. The color picker that allows this type of query is shown in Figure 2. 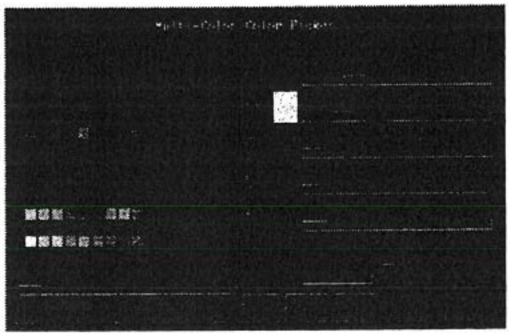 Figure 2: Multi-Color Color Picker with two selected colors. The palette on the left displays the set of selectable colors. Once a color is selected, it appears in a box on the right, and the amount of the color can be specified by the sliders. In this case, 25% light blue and 15% white have been selected. *Texture:* Texture distance is computed as weighted Euclidean distance in the three dimensional texture space. The most common weighting factors are the inverse variances for each component, computed over the

Element	'252 Claim Recitation	Exemplary Citations in Reference
		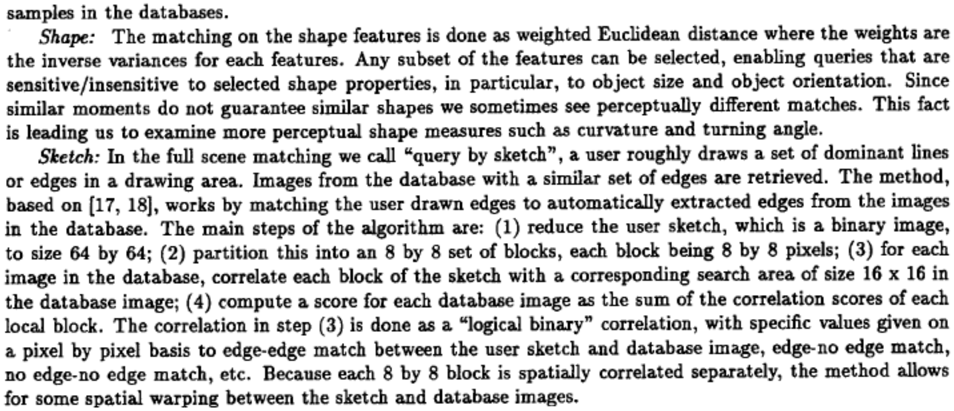 samples in the databases. *Shape:* The matching on the shape features is done as weighted Euclidean distance where the weights are the inverse variances for each features. Any subset of the features can be selected, enabling queries that are sensitive/insensitive to selected shape properties, in particular, to object size and object orientation. Since similar moments do not guarantee similar shapes we sometimes see perceptually different matches. This fact is leading us to examine more perceptual shape measures such as curvature and turning angle. *Sketch:* In the full scene matching we call "query by sketch", a user roughly draws a set of dominant lines or edges in a drawing area. Images from the database with a similar set of edges are retrieved. The method, based on [17, 18], works by matching the user drawn edges to automatically extracted edges from the images in the database. The main steps of the algorithm are: (1) reduce the user sketch, which is a binary image, to size 64 by 64; (2) partition this into an 8 by 8 set of blocks, each block being 8 by 8 pixels; (3) for each image in the database, correlate each block of the sketch with a corresponding search area of size 16 x 16 in the database image; (4) compute a score for each database image as the sum of the correlation scores of each local block. The correlation in step (3) is done as a "logical binary" correlation, with specific values given on a pixel by pixel basis to edge-edge match between the user sketch and database image, edge-no edge match, no edge-no edge match, etc. Because each 8 by 8 block is spatially correlated separately, the method allows for some spatial warping between the sketch and database images. IBM0002390 at 2393-95 Query by Image and Video Content: The QBIC System (IBM 0000893–902): "We have developed the QBIC (Query by Image Content) system to explore content-based retrieval methods. QBIC allows queries on large image and video databases based on • example images, • user-constructed sketches and drawings, • selected color and texture patterns, • camera and object motion, and • other graphical information." *Id.* at pp. 1–3. "To achieve this functionality, QBIC has two main components: database population (the process of creating an image database) and database query. During the population, images and videos are processed to extract features describing their content-colors, textures, shapes, and camera and object motion-and the features are stored in a database. During the query, the user composes a query graphically. Features are generated from the graphical query and then input to a matching engine that finds images or videos from the database with similar features. Figure 2 shows the system architecture." *Id.* at p. 3.

Element	'252 Claim Recitation	Exemplary Citations in Reference
		Image, Audio, and Video Extenders Administration and Programming (IBM 000001–617): "DB2 (DB2) Universal Database (UDB) is a powerful, object-relational database manager. It stores and protects traditional numeric and character data, as well as large, complex objects (LOBs). DB2 extenders help you exploit DB2 's object-relational features. The extenders define distinct data types and special functions for image, audio, video, and text objects. … Because of that, the extenders give your applications a single point of access to any or all of these types of data, along with traditional numeric and character data. In addition, the extenders give your applications new ways to search for information." *Id*. at p. 3. Managing Enterprise Information Portal (IBM 000777–880): "Many paper-intensive enterprises, such as insurance companies and financial institutions, administer large volumes of business-related content. The need for an enterprise solution for managing and accessing business information spans many industries. A content server is a software system that stores multimedia, business forms, documents, and related data, along with metadata that allows employees to process and work with the content." *Id*. at p. 3. "The image search server uses IBM's QBIC (query by image content) technology to help you search for objects by certain visual properties, such as color and texture. The image search server analyzes images and stores the image information in a database. Then users can run image queries, which use the visual properties of images, to match colors, textures, and their positions without describing them in words. You can combine content-based queries with text and keyword searches for more powerful retrieval of image and multimedia data." *Id*. at p. 63. QBIC Demo (IBM000747)

Element	'252 Claim Recitation	Exemplary Citations in Reference
		 Flickner 8/31/23 Depo at Ex. 6 at 1 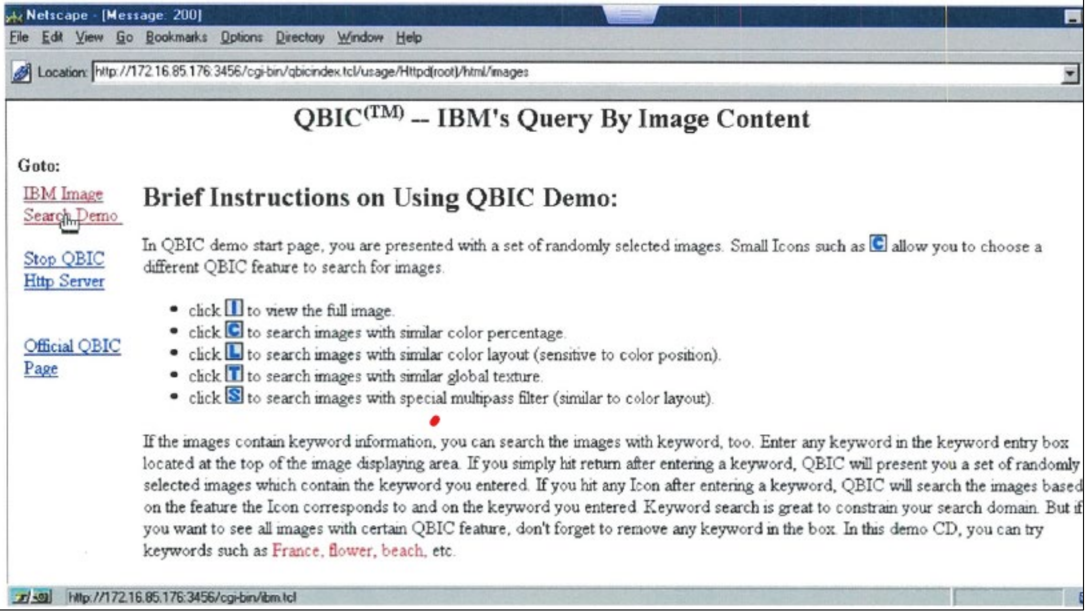 QBIC$^{(TM)}$ -- IBM's Query By Image Content Goto: IBM Image Search Demo Stop QBIC Http Server Official QBIC Page **Brief Instructions on Using QBIC Demo:** In QBIC demo start page, you are presented with a set of randomly selected images. Small Icons such as C allow you to choose a different QBIC feature to search for images. - click I to view the full image. - click C to search images with similar color percentage. - click L to search images with similar color layout (sensitive to color position). - click T to search images with similar global texture. - click S to search images with special multipass filter (similar to color layout). If the images contain keyword information, you can search the images with keyword, too. Enter any keyword in the keyword entry box located at the top of the image displaying area. If you simply hit return after entering a keyword, QBIC will present you a set of randomly selected images which contain the keyword you entered. If you hit any Icon after entering a keyword, QBIC will search the images based on the feature the Icon corresponds to and on the keyword you entered. Keyword search is great to constrain your search domain. But if you want to see all images with certain QBIC feature, don't forget to remove any keyword in the box. In this demo CD, you can try keywords such as France, flower, beach, etc.

Element	'252 Claim Recitation	Exemplary Citations in Reference
		Id. at 2.  Flickner 8/31/23 Depo Ex. 7 at 1

Element	'252 Claim Recitation	Exemplary Citations in Reference
		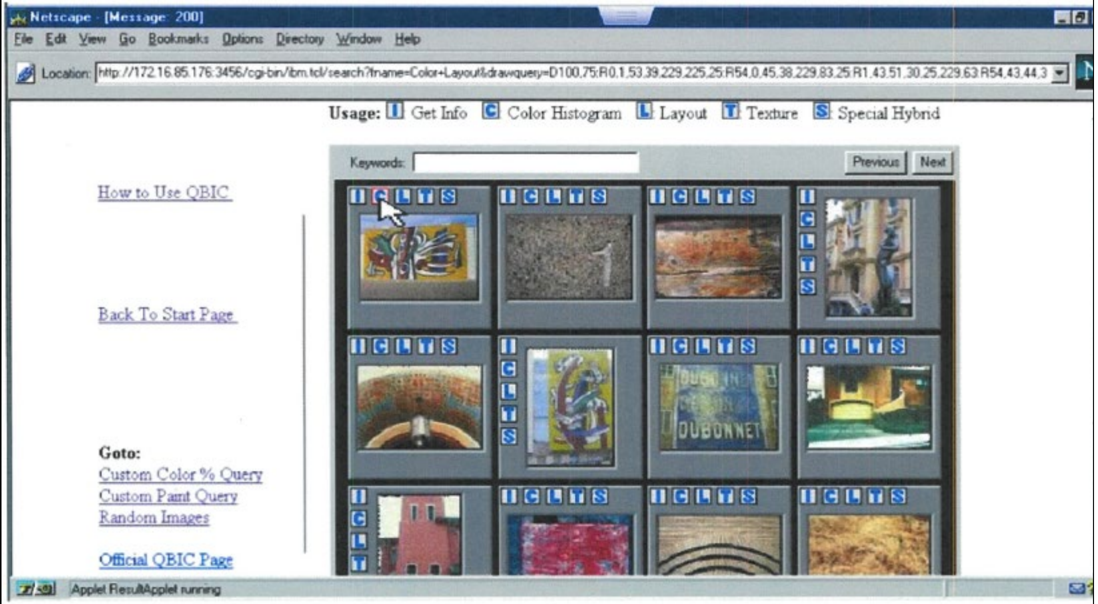 Location: http://172.16.85.176:3456/cgi-bin/ibm.tcl/search?fname=Color+Layout&drawquery=D100,75:R0,1,53,39,229,225,25:R54,0,45,38,229,83,25:R1,43,51,30,25,229,63:R54,43,44,3 *Id.* at 2

Element	'252 Claim Recitation	Exemplary Citations in Reference
		 Flickner 8/31/23 Depo at Ex. 8

Element	'252 Claim Recitation	Exemplary Citations in Reference
		 *Id.* at 2

Element	’252 Claim Recitation	Exemplary Citations in Reference
		 *Id.* at 3

Element	'252 Claim Recitation	Exemplary Citations in Reference
		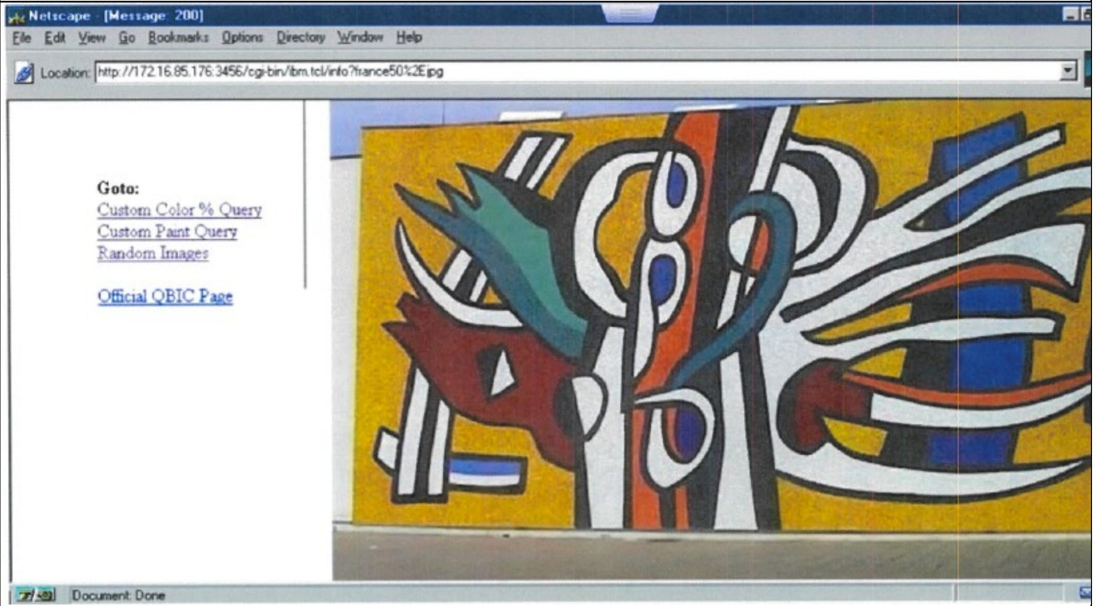 Location: http://172.16.85.176:3456/cgi-bin/ibm.tcl/info?france50%2Ejpg Flickner Depo at 22-25

Element	'252 Claim Recitation	Exemplary Citations in Reference
		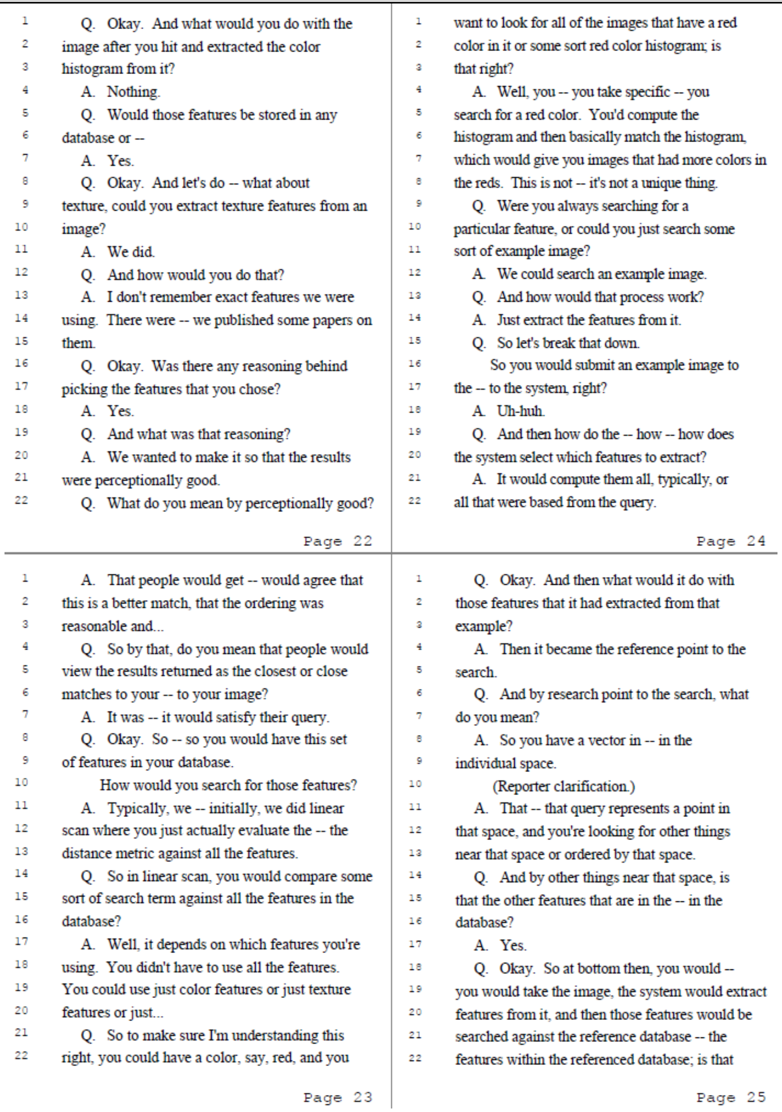

Q. Okay. And what would you do with the image after you hit and extracted the color histogram from it?
A. Nothing.
Q. Would those features be stored in any database or --
A. Yes.
Q. Okay. And let's do -- what about texture, could you extract texture features from an image?
A. We did.
Q. And how would you do that?
A. I don't remember exact features we were using. There were -- we published some papers on them.
Q. Okay. Was there any reasoning behind picking the features that you chose?
A. Yes.
Q. And what was that reasoning?
A. We wanted to make it so that the results were perceptionally good.
Q. What do you mean by perceptionally good?

Page 22

A. That people would get -- would agree that this is a better match, that the ordering was reasonable and...
Q. So by that, do you mean that people would view the results returned as the closest or close matches to your -- to your image?
A. It was -- it would satisfy their query.
Q. Okay. So -- so you would have this set of features in your database.
How would you search for those features?
A. Typically, we -- initially, we did linear scan where you just actually evaluate the -- the distance metric against all the features.
Q. So in linear scan, you would compare some sort of search term against all the features in the database?
A. Well, it depends on which features you're using. You didn't have to use all the features. You could use just color features or just texture features or just...
Q. So to make sure I'm understanding this right, you could have a color, say, red, and you

Page 23

want to look for all of the images that have a red color in it or some sort red color histogram; is that right?
A. Well, you -- you take specific -- you search for a red color. You'd compute the histogram and then basically match the histogram, which would give you images that had more colors in the reds. This is not -- it's not a unique thing.
Q. Were you always searching for a particular feature, or could you just search some sort of example image?
A. We could search an example image.
Q. And how would that process work?
A. Just extract the features from it.
Q. So let's break that down.
So you would submit an example image to the -- to the system, right?
A. Uh-huh.
Q. And then how do the -- how -- how does the system select which features to extract?
A. It would compute them all, typically, or all that were based from the query.

Page 24

Q. Okay. And then what would it do with those features that it had extracted from that example?
A. Then it became the reference point to the search.
Q. And by research point to the search, what do you mean?
A. So you have a vector in -- in the individual space.
(Reporter clarification.)
A. That -- that query represents a point in that space, and you're looking for other things near that space or ordered by that space.
Q. And by other things near that space, is that the other features that are in the -- in the database?
A. Yes.
Q. Okay. So at bottom then, you would -- you would take the image, the system would extract features from it, and then those features would be searched against the reference database -- the features within the referenced database; is that

Page 25

Element	'252 Claim Recitation	Exemplary Citations in Reference
		Flickner Depo at 54-55 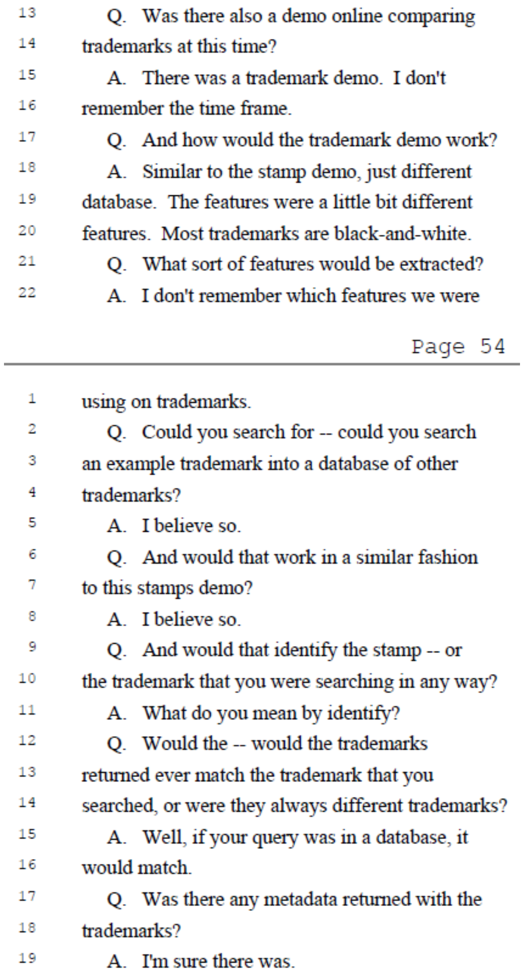 Q. Was there also a demo online comparing trademarks at this time? A. There was a trademark demo. I don't remember the time frame. Q. And how would the trademark demo work? A. Similar to the stamp demo, just different database. The features were a little bit different features. Most trademarks are black-and-white. Q. What sort of features would be extracted? A. I don't remember which features we were Page 54 using on trademarks. Q. Could you search for -- could you search an example trademark into a database of other trademarks? A. I believe so. Q. And would that work in a similar fashion to this stamps demo? A. I believe so. Q. And would that identify the stamp -- or the trademark that you were searching in any way? A. What do you mean by identify? Q. Would the -- would the trademarks returned ever match the trademark that you searched, or were they always different trademarks? A. Well, if your query was in a database, it would match. Q. Was there any metadata returned with the trademarks? A. I'm sure there was. Flickner Depo at 25-26.

Element	'252 Claim Recitation	Exemplary Citations in Reference
		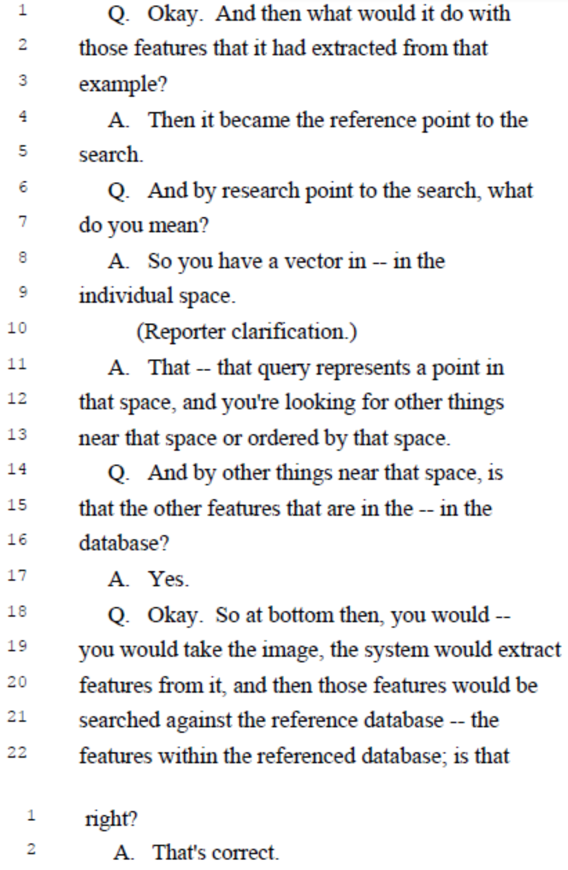 Q. Okay. And then what would it do with those features that it had extracted from that example? A. Then it became the reference point to the search. Q. And by research point to the search, what do you mean? A. So you have a vector in -- in the individual space. (Reporter clarification.) A. That -- that query represents a point in that space, and you're looking for other things near that space or ordered by that space. Q. And by other things near that space, is that the other features that are in the -- in the database? A. Yes. Q. Okay. So at bottom then, you would -- you would take the image, the system would extract features from it, and then those features would be searched against the reference database -- the features within the referenced database; is that right? A. That's correct. *See also* Flickner Depo at 89-90, 162, 185, 214-215, 217 Flickner 8/29/23 Depo at 44

Element	'252 Claim Recitation	Exemplary Citations in Reference
		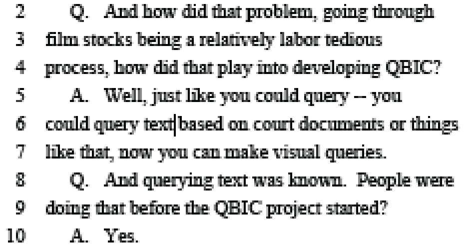 Q. And how did that problem, going through film stocks being a relatively labor tedious process, how did that play into developing QBIC? A. Well, just like you could query -- you could query text based on court documents or things like that, now you can make visual queries. Q. And querying text was known. People were doing that before the QBIC project started? A. Yes. *Id.* at 45–47 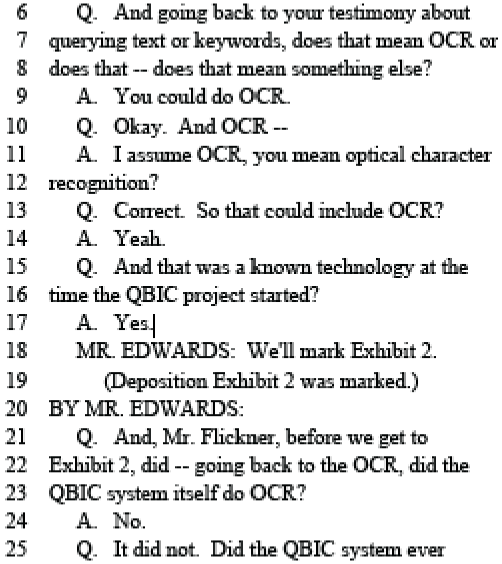 Q. And going back to your testimony about querying text or keywords, does that mean OCR or does that -- does that mean something else? A. You could do OCR. Q. Okay. And OCR -- A. I assume OCR, you mean optical character recognition? Q. Correct. So that could include OCR? A. Yeah. Q. And that was a known technology at the time the QBIC project started? A. Yes. MR. EDWARDS: We'll mark Exhibit 2. (Deposition Exhibit 2 was marked.) BY MR. EDWARDS: Q. And, Mr. Flickner, before we get to Exhibit 2, did -- going back to the OCR, did the QBIC system itself do OCR? A. No. Q. It did not. Did the QBIC system ever

Element	'252 Claim Recitation	Exemplary Citations in Reference
		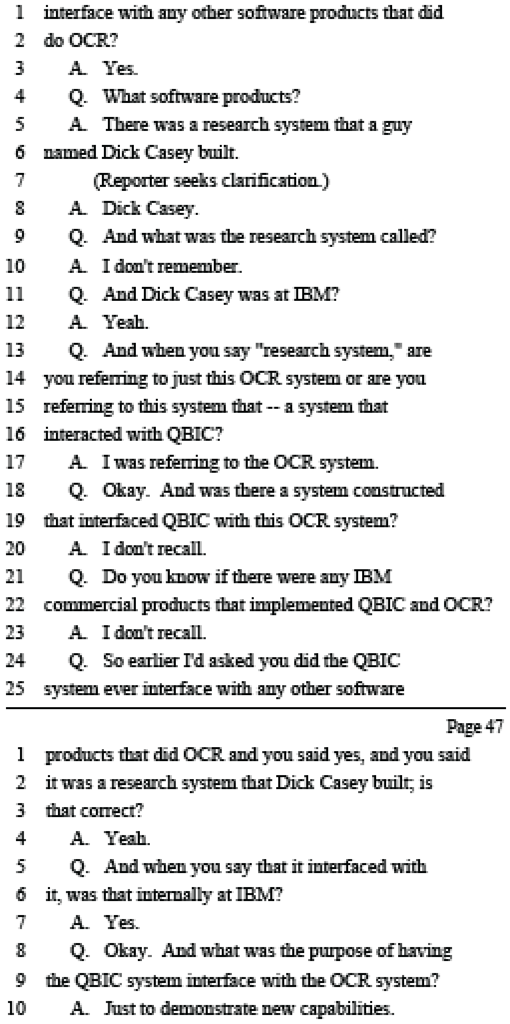 interface with any other software products that did do OCR? A. Yes. Q. What software products? A. There was a research system that a guy named Dick Casey built. (Reporter seeks clarification.) A. Dick Casey. Q. And what was the research system called? A. I don't remember. Q. And Dick Casey was at IBM? A. Yeah. Q. And when you say "research system," are you referring to just this OCR system or are you referring to this system that -- a system that interacted with QBIC? A. I was referring to the OCR system. Q. Okay. And was there a system constructed that interfaced QBIC with this OCR system? A. I don't recall. Q. Do you know if there were any IBM commercial products that implemented QBIC and OCR? A. I don't recall. Q. So earlier I'd asked you did the QBIC system ever interface with any other software Page 47 products that did OCR and you said yes, and you said it was a research system that Dick Casey built; is that correct? A. Yeah. Q. And when you say that it interfaced with it, was that internally at IBM? A. Yes. Q. Okay. And what was the purpose of having the QBIC system interface with the OCR system? A. Just to demonstrate new capabilities. *Id.* at 76

Element	'252 Claim Recitation	Exemplary Citations in Reference
		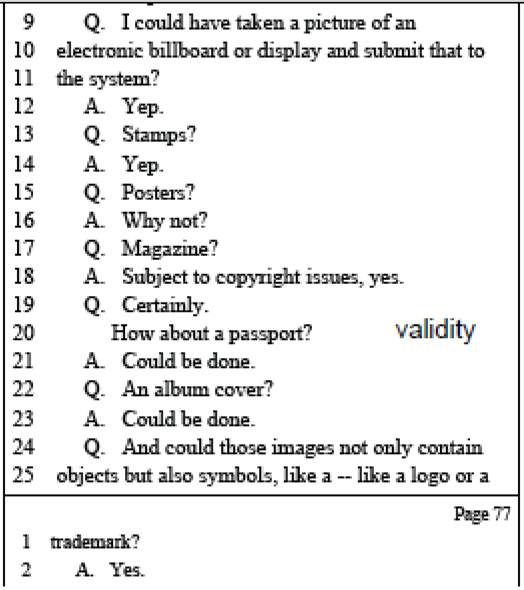 Q. I could have taken a picture of an electronic billboard or display and submit that to the system? A. Yep. Q. Stamps? A. Yep. Q. Posters? A. Why not? Q. Magazine? A. Subject to copyright issues, yes. Q. Certainly. How about a passport? validity A. Could be done. Q. An album cover? A. Could be done. Q. And could those images not only contain objects but also symbols, like a -- like a logo or a Page 77 trademark? A. Yes. *Id.* at 187–188 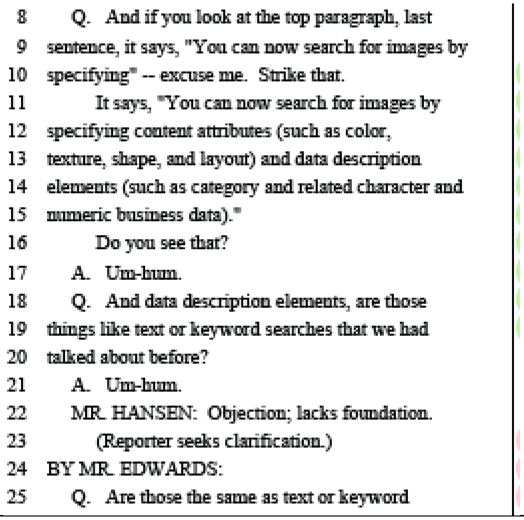 Q. And if you look at the top paragraph, last sentence, it says, "You can now search for images by specifying" -- excuse me. Strike that. It says, "You can now search for images by specifying content attributes (such as color, texture, shape, and layout) and data description elements (such as category and related character and numeric business data)." Do you see that? A. Um-hum. Q. And data description elements, are those things like text or keyword searches that we had talked about before? A. Um-hum. MR. HANSEN: Objection; lacks foundation. (Reporter seeks clarification.) BY MR. EDWARDS: Q. Are those the same as text or keyword

Element	'252 Claim Recitation	Exemplary Citations in Reference
		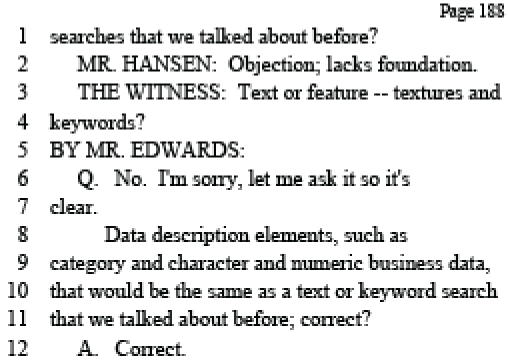 Page 188 searches that we talked about before? MR. HANSEN: Objection; lacks foundation. THE WITNESS: Text or feature -- textures and keywords? BY MR. EDWARDS: Q. No. I'm sorry, let me ask it so it's clear. Data description elements, such as category and character and numeric business data, that would be the same as a text or keyword search that we talked about before; correct? A. Correct. To the extent it is found that QBIC does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that QBIC does not anticipate this claim, QBIC renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine QBIC with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results. For example, the '252 patent acknowledges that acquiring information about objects through symbols, such as barcodes or some other means to encode detectable information in the image or on the object, were traditional methods known in the art and not part the present invention. '252 Patent, 3:7-16. Indeed, the '252 patent acknowledges that algorithms, software, and hardware components were commercially available at the time of the invention that could detect and decode symbols, such as barcodes or text, in an input image. *Id.* at 11:56-61. This limitation was thus known in the art by patentee's own admission, and by asserting

Element	'252 Claim Recitation	Exemplary Citations in Reference
		this claim, Plaintiff is attempting to reclaim disavowed claim scope and claim subject matter which patentee has acknowledged is not part of the invention.
[18C]	providing access to a distal server configured to use a database to identify pertinent information associated with the recognizable symbols based on the symbol information; and	QBIC discloses providing access to a distal server configured to use a database to identify pertinent information associated with the recognizable symbols based on the symbol information. The QBIC Project: Querying Images by Content Using Color, Texture, and Shape (IBM 0002390–404) at 3 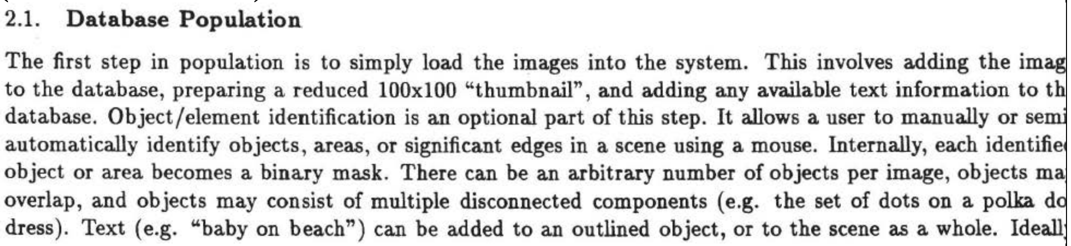 **2.1. Database Population** The first step in population is to simply load the images into the system. This involves adding the imag to the database, preparing a reduced 100x100 "thumbnail", and adding any available text information to th database. Object/element identification is an optional part of this step. It allows a user to manually or semi automatically identify objects, areas, or significant edges in a scene using a mouse. Internally, each identifie object or area becomes a binary mask. There can be an arbitrary number of objects per image, objects ma overlap, and objects may consist of multiple disconnected components (e.g. the set of dots on a polka do dress). Text (e.g. "baby on beach") can be added to an outlined object, or to the scene as a whole. Ideall The QBIC Project: Querying Images by Content Using Color, Texture, and Shape (IBM 0002390–404) at 7 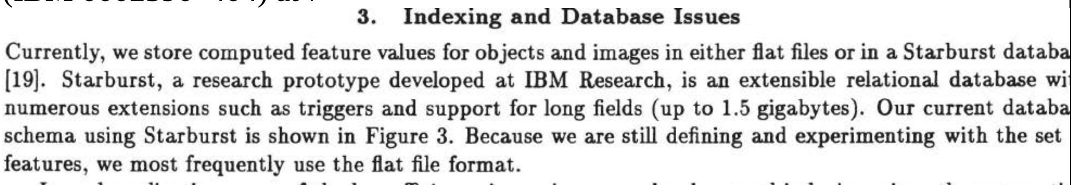 **3. Indexing and Database Issues** Currently, we store computed feature values for objects and images in either flat files or in a Starburst databa [19]. Starburst, a research prototype developed at IBM Research, is an extensible relational database wi numerous extensions such as triggers and support for long fields (up to 1.5 gigabytes). Our current databa schema using Starburst is shown in Figure 3. Because we are still defining and experimenting with the set features, we most frequently use the flat file format. … 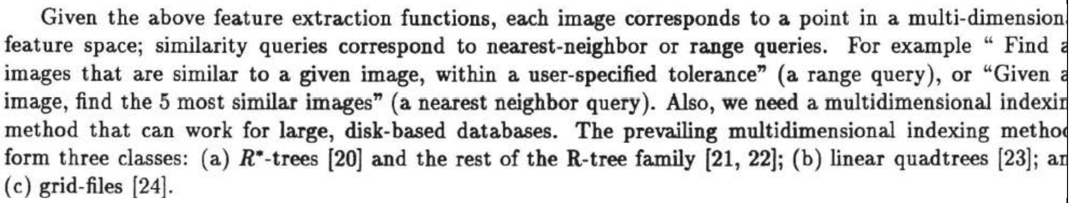 Given the above feature extraction functions, each image corresponds to a point in a multi-dimension feature space; similarity queries correspond to nearest-neighbor or range queries. For example " Find a images that are similar to a given image, within a user-specified tolerance" (a range query), or "Given a image, find the 5 most similar images" (a nearest neighbor query). Also, we need a multidimensional indexir method that can work for large, disk-based databases. The prevailing multidimensional indexing methoc form three classes: (a) R^*-trees [20] and the rest of the R-tree family [21, 22]; (b) linear quadtrees [23]; ar (c) grid-files [24]. Query by Image and Video Content: The QBIC System (IBM 0000893–902):

<table>
<tr><th>Element</th><th>'252 Claim Recitation</th><th>Exemplary Citations in Reference</th></tr>
<tr><td></td><td></td><td>

"QBIC has two main components: database population (the process of creating an image database) and database query. During the population, images and videos are processed to extract features describing their content-colors, textures, shapes, and camera and object motion-and the features are stored in a database. During the query, the user composes a query graphically. Features are generated from the graphical query and then input to a matching engine that finds images or videos from the database with similar features. Figure 2 shows the system architecture." *Id*. at p. 3.

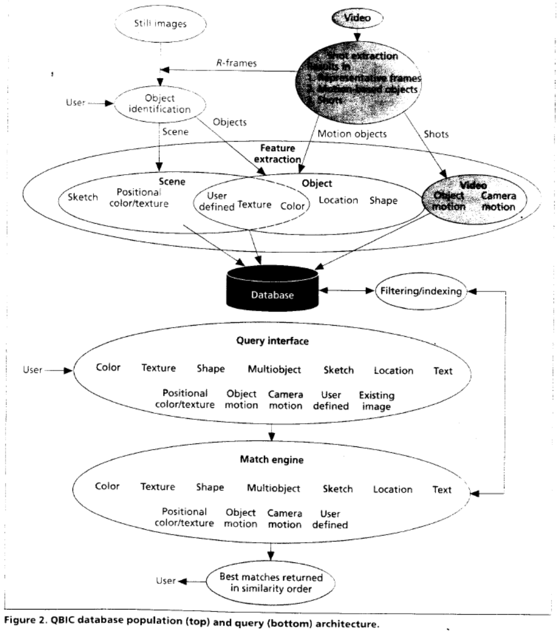

Figure 2. QBIC database population (top) and query (bottom) architecture.

Id. at Fig. 2.

"**SAMPLE QUERIES**

For each full-scene image, identified image object, r-frame, and identified video object resulting from the above processing, a set of features is computed to allow content-based queries. The features are computed and stored during database population. We present a brief description of the features and the associated

</td></tr>
</table>

Element	'252 Claim Recitation	Exemplary Citations in Reference
		queries. Mathematical details on the features and matching methods can be found in Ashley et al. and Niblack et al. Average color queries let users find images or objects that are similar to a selected color, say from a color wheel, or to the color of an object. The feature used in the query is a 3D vector of Munsell color coordinates. Histogram color queries return items with matching color distributions—say, a fabric pattern with approximately 40 percent red and 20 percent blue. For this case, the underlying feature is a 256-element histogram computed over a quantized version of the color space. Figure 7 shows a histogram query on still images and a color query on video r-frames. Note that in the query specification for the histogram query of Figure 7, the user has selected percentages of two colors (blue and white) by adjusting sliders. Using such a query, an advertising agent could, for example, search for a picture of a beach scene, one predominantly blue (for sky and water) and white (for sand and clouds); or find images with similar color spreads for a uniform ad campaign. The average color query demonstrates a query against a video shot database where the user is searching for red r-frames. Again, the query specification is on the left and the best hits are on the right. Figure 8 shows an example texture query. In this case, the query is specified by selecting from a sampler—a set of prestored example images. The underlying texture features are mathematical representations of coarseness, contrast, and directionality features. Coarseness measures the scale of a texture (pebbles versus boulders), contrast describes its vividness, and directionality describes whether it has a favored direction (like grass) or not (like a smooth object). An object shape query is shown in Figure 8. In this case, the query specification is the drawn shape on the left. Area, circularity, eccentricity, major-axis direction, features derived from the object moments, and a set of tangent angles around the object perimeter are the features used to characterize and match shapes. Figure 9 illustrates query by sketch. In this case, the query specification is a freehand drawing of the dominant lines and edges in the image. The sketch feature is an automatically extracted reduced-resolution 'edge map.' Matching is done by using a template-matching technique. A multiobject query asking for images that contain both a red round object and a

Element	'252 Claim Recitation	Exemplary Citations in Reference
		green textured object is shown in the bottom of Figure 9. The features are standard color and texture. The matching is done by combining the color and texture distances. Combining distances is applied to arbitrary sets of objects and features to implement logical And semantics. . . . . **ANNOTATION AND DATABASE POPULATION TOOLS**. Automatic methods (such as our Positional Color query) that don't rely on object identification, methods that identify objects automatically as in the museum image example, fast and easy-to-use semiautomatic outlining tools, and motion-based segmentation algorithms will enable additional application areas. **FEATURE EXTRACTION AND MATCHING METHODS**. New mathematical representations of video, image, and object attributes that capture 'interesting' features for retrieval are needed. Features that describe new image properties such as alternate texture measures or that are based on fractals or wavelet representations, for example, may offer advantages of representation, indexability, and ease of similarity matching." *Id*. at pp. 7–8. Image, Audio, and Video Extenders Administration and Programming (IBM 000001–617): "DB2 (DB2) Universal Database (UDB) is a powerful, object-relational database manager. It stores and protects traditional numeric and character data, as well as large, complex objects (LOBs). DB2 extenders help you exploit DB2 's object-relational features. The extenders define distinct data types and special functions for image, audio, video, and text objects. … Because of that, the extenders give your applications a single point of access to any or all of these types of data, along with traditional numeric and character data. In addition, the extenders give your applications new ways to search for information." *Id*. at p. 3.

Element	'252 Claim Recitation	Exemplary Citations in Reference
		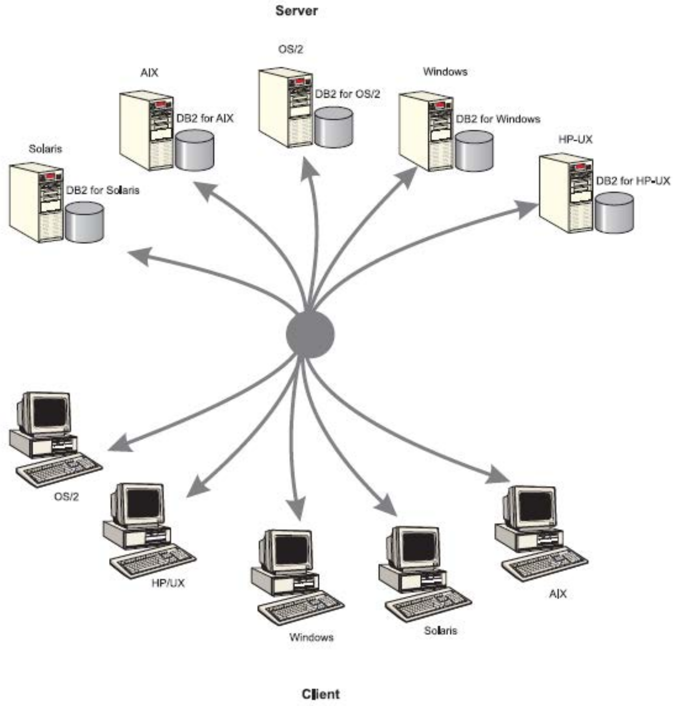 *Figure 6. DB2 extender platforms* *Id.* at p. 12 (Fig. 6). "The DB2 Extenders run in the DB2 client/server environment. This environment consists of a database server and one or more remote database clients. The DB2 extender services run on the server. Before you can access them, you have to start them." *Id.* at p. 47.

Element	'252 Claim Recitation	Exemplary Citations in Reference
		QBIC Demo (IBM000747) 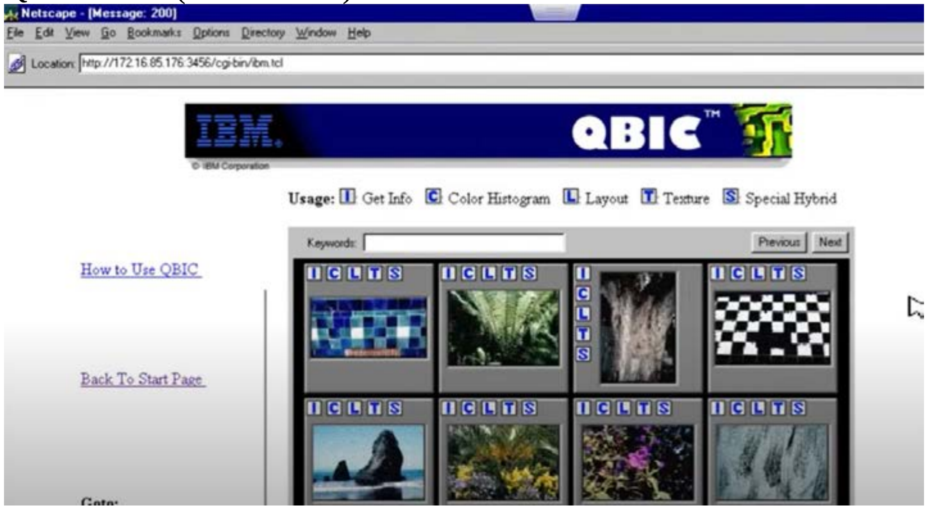 QBIC Demo (IBM000747) 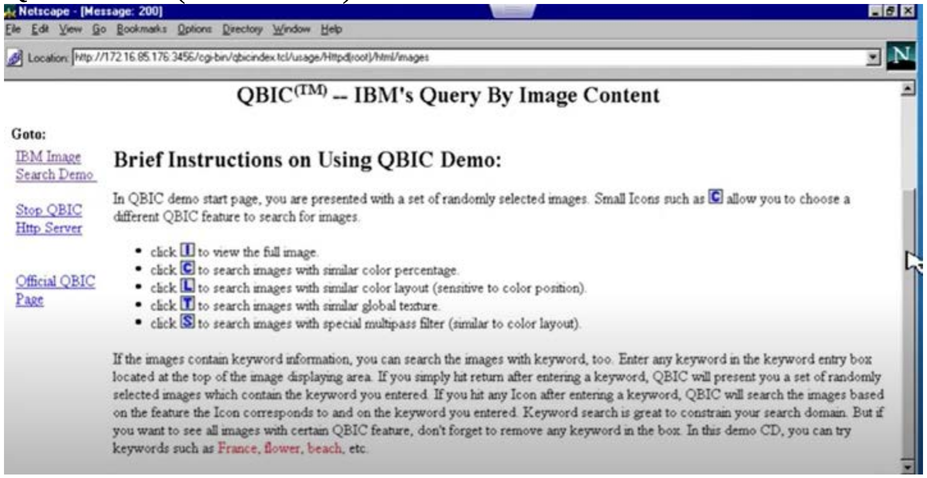 QBIC(TM) -- IBM's Query By Image Content Goto: IBM Image Search Demo Stop QBIC Http Server Official QBIC Page **Brief Instructions on Using QBIC Demo:** In QBIC demo start page, you are presented with a set of randomly selected images. Small Icons such as C allow you to choose a different QBIC feature to search for images. • click I to view the full image. • click C to search images with similar color percentage. • click L to search images with similar color layout (sensitive to color position). • click T to search images with similar global texture. • click S to search images with special multipass filter (similar to color layout). If the images contain keyword information, you can search the images with keyword, too. Enter any keyword in the keyword entry box located at the top of the image displaying area. If you simply hit return after entering a keyword, QBIC will present you a set of randomly selected images which contain the keyword you entered. If you hit any Icon after entering a keyword, QBIC will search the images based on the feature the Icon corresponds to and on the keyword you entered. Keyword search is great to constrain your search domain. But if you want to see all images with certain QBIC feature, don't forget to remove any keyword in the box. In this demo CD, you can try keywords such as France, flower, beach, etc.

Element	'252 Claim Recitation	Exemplary Citations in Reference
		Flickner Depo at 20 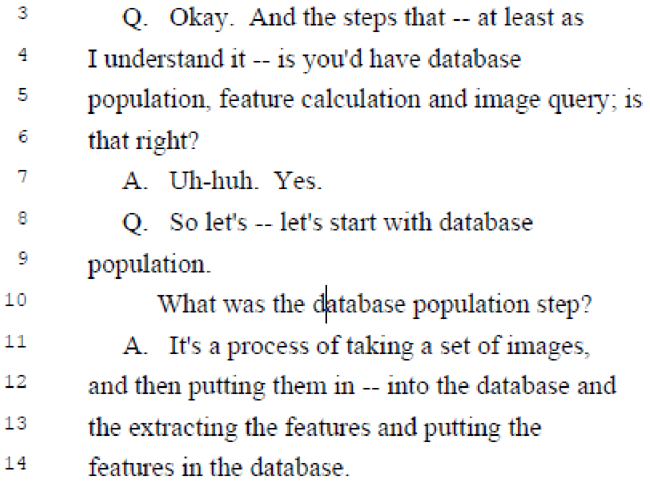 Q. Okay. And the steps that -- at least as I understand it -- is you'd have database population, feature calculation and image query; is that right? A. Uh-huh. Yes. Q. So let's -- let's start with database population. What was the database population step? A. It's a process of taking a set of images, and then putting them in -- into the database and the extracting the features and putting the features in the database. 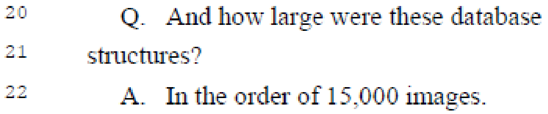 Q. And how large were these database structures? A. In the order of 15,000 images. Flickner Depo at 18 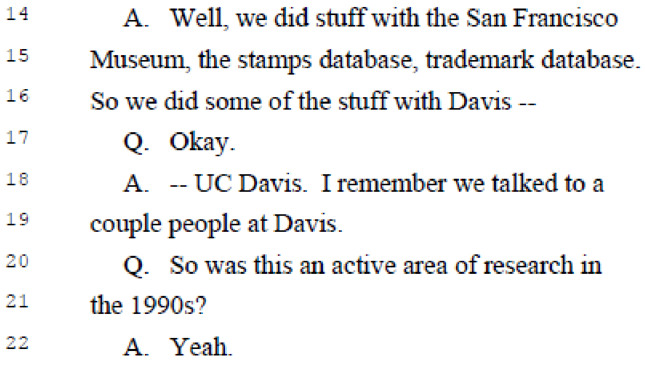 A. Well, we did stuff with the San Francisco Museum, the stamps database, trademark database. So we did some of the stuff with Davis -- Q. Okay. A. -- UC Davis. I remember we talked to a couple people at Davis. Q. So was this an active area of research in the 1990s? A. Yeah. Flickner Depo at 36

Element	'252 Claim Recitation	Exemplary Citations in Reference
		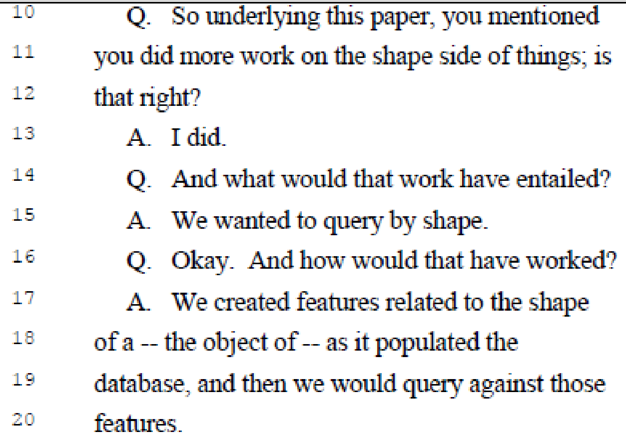 Q. So underlying this paper, you mentioned you did more work on the shape side of things; is that right? A. I did. Q. And what would that work have entailed? A. We wanted to query by shape. Q. Okay. And how would that have worked? A. We created features related to the shape of a -- the object of -- as it populated the database, and then we would query against those features. Flickner Depo at 39-40 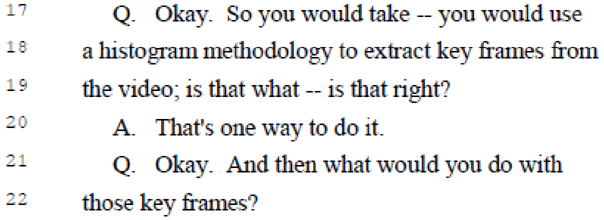 Q. Okay. So you would take -- you would use a histogram methodology to extract key frames from the video; is that what -- is that right? A. That's one way to do it. Q. Okay. And then what would you do with those key frames?

Element	'252 Claim Recitation	Exemplary Citations in Reference
		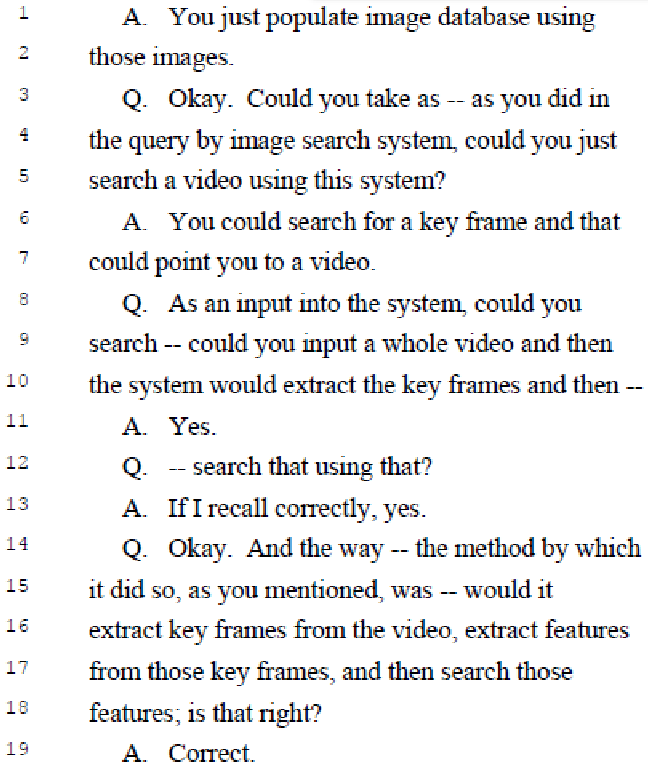 A. You just populate image database using those images. Q. Okay. Could you take as -- as you did in the query by image search system, could you just search a video using this system? A. You could search for a key frame and that could point you to a video. Q. As an input into the system, could you search -- could you input a whole video and then the system would extract the key frames and then -- A. Yes. Q. -- search that using that? A. If I recall correctly, yes. Q. Okay. And the way -- the method by which it did so, as you mentioned, was -- would it extract key frames from the video, extract features from those key frames, and then search those features; is that right? A. Correct. Flickner Depo at 88-89

Element	'252 Claim Recitation	Exemplary Citations in Reference
		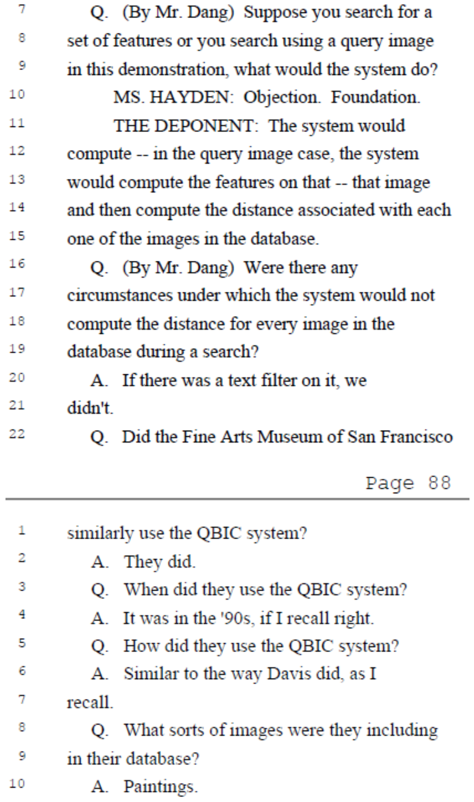 Q. (By Mr. Dang) Suppose you search for a set of features or you search using a query image in this demonstration, what would the system do? MS. HAYDEN: Objection. Foundation. THE DEPONENT: The system would compute -- in the query image case, the system would compute the features on that -- that image and then compute the distance associated with each one of the images in the database. Q. (By Mr. Dang) Were there any circumstances under which the system would not compute the distance for every image in the database during a search? A. If there was a text filter on it, we didn't. Q. Did the Fine Arts Museum of San Francisco Page 88 similarly use the QBIC system? A. They did. Q. When did they use the QBIC system? A. It was in the '90s, if I recall right. Q. How did they use the QBIC system? A. Similar to the way Davis did, as I recall. Q. What sorts of images were they including in their database? A. Paintings. Flickner Depo at 142

Element	'252 Claim Recitation	Exemplary Citations in Reference
		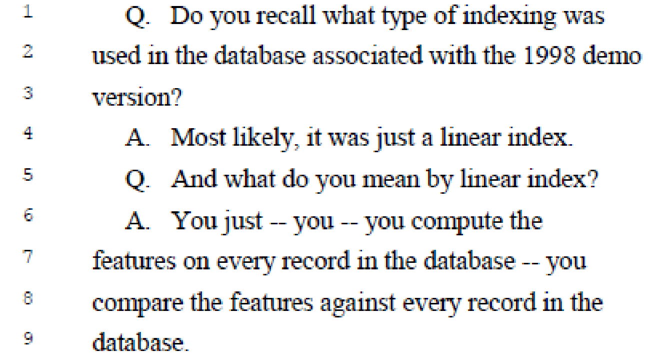 Q. Do you recall what type of indexing was used in the database associated with the 1998 demo version? A. Most likely, it was just a linear index. Q. And what do you mean by linear index? A. You just -- you -- you compute the features on every record in the database -- you compare the features against every record in the database. Flickner Depo at 168-169 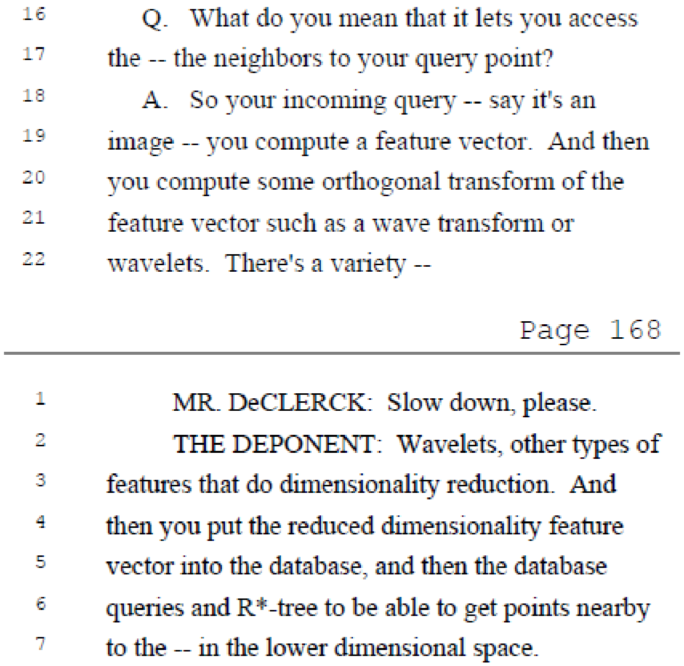 Q. What do you mean that it lets you access the -- the neighbors to your query point? A. So your incoming query -- say it's an image -- you compute a feature vector. And then you compute some orthogonal transform of the feature vector such as a wave transform or wavelets. There's a variety -- Page 168 MR. DeCLERCK: Slow down, please. THE DEPONENT: Wavelets, other types of features that do dimensionality reduction. And then you put the reduced dimensionality feature vector into the database, and then the database queries and R*-tree to be able to get points nearby to the -- in the lower dimensional space.

Element	'252 Claim Recitation	Exemplary Citations in Reference
		Flickner Depo at 181 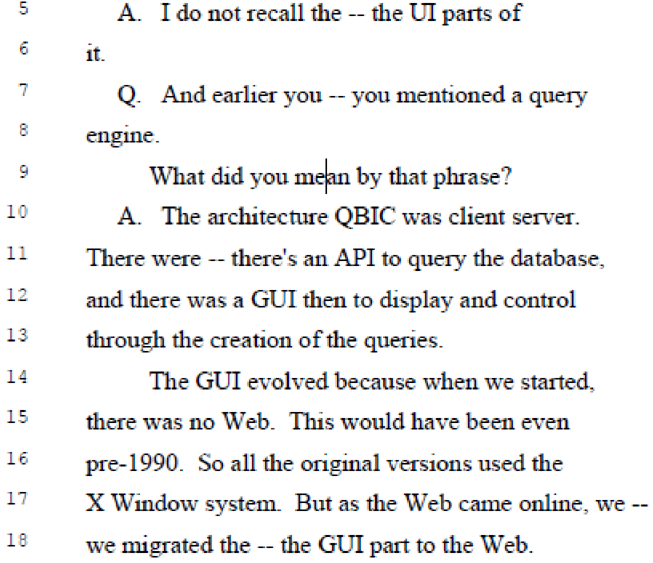 A. I do not recall the -- the UI parts of it. Q. And earlier you -- you mentioned a query engine. What did you mean by that phrase? A. The architecture QBIC was client server. There were -- there's an API to query the database, and there was a GUI then to display and control through the creation of the queries. The GUI evolved because when we started, there was no Web. This would have been even pre-1990. So all the original versions used the X Window system. But as the Web came online, we -- we migrated the -- the GUI part to the Web. Flickner Depo at 201-202

Element	'252 Claim Recitation	Exemplary Citations in Reference
		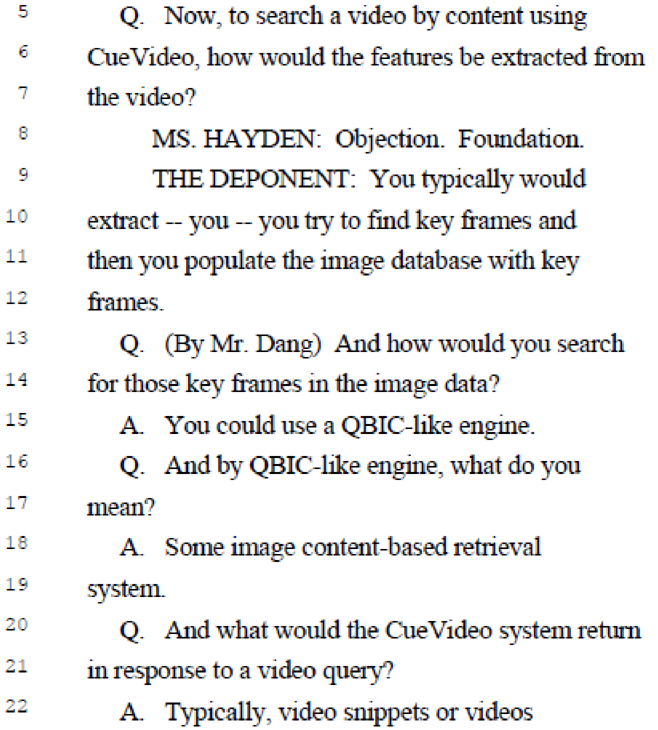 Q. Now, to search a video by content using CueVideo, how would the features be extracted from the video? MS. HAYDEN: Objection. Foundation. THE DEPONENT: You typically would extract -- you -- you try to find key frames and then you populate the image database with key frames. Q. (By Mr. Dang) And how would you search for those key frames in the image data? A. You could use a QBIC-like engine. Q. And by QBIC-like engine, what do you mean? A. Some image content-based retrieval system. Q. And what would the CueVideo system return in response to a video query? A. Typically, video snippets or videos

Element	'252 Claim Recitation	Exemplary Citations in Reference
		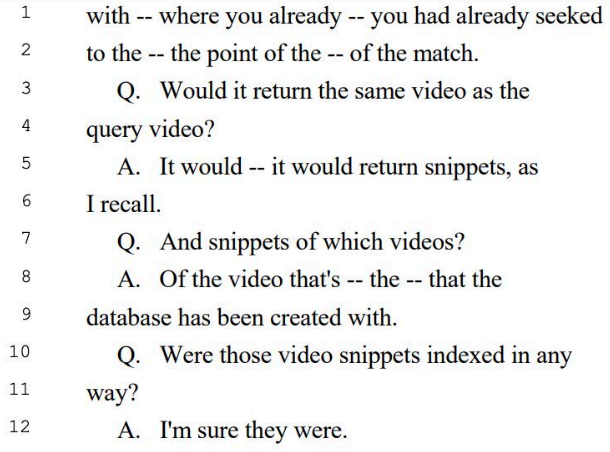 with -- where you already -- you had already seeked to the -- the point of the -- of the match. Q. Would it return the same video as the query video? A. It would -- it would return snippets, as I recall. Q. And snippets of which videos? A. Of the video that's -- the -- that the database has been created with. Q. Were those video snippets indexed in any way? A. I'm sure they were. Flickner Depo at 91-92 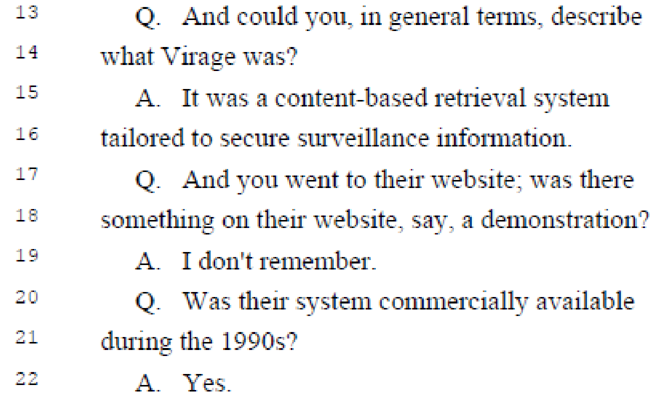 Q. And could you, in general terms, describe what Virage was? A. It was a content-based retrieval system tailored to secure surveillance information. Q. And you went to their website; was there something on their website, say, a demonstration? A. I don't remember. Q. Was their system commercially available during the 1990s? A. Yes.

Element	'252 Claim Recitation	Exemplary Citations in Reference
		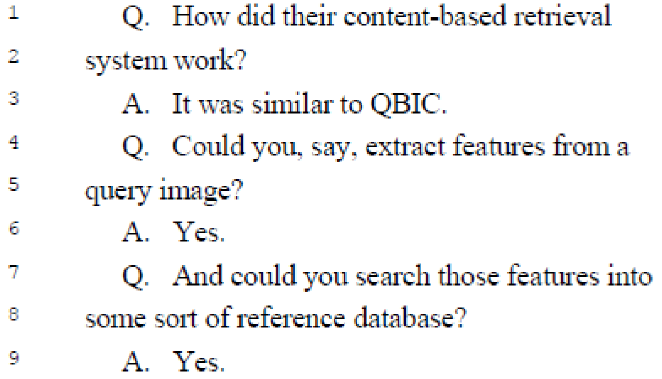 Q. How did their content-based retrieval system work? A. It was similar to QBIC. Q. Could you, say, extract features from a query image? A. Yes. Q. And could you search those features into some sort of reference database? A. Yes. Flickner Depo at 48-49 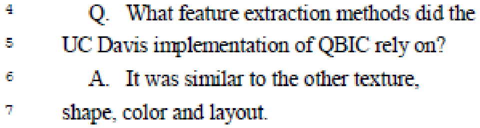 Q. What feature extraction methods did the UC Davis implementation of QBIC rely on? A. It was similar to the other texture, shape, color and layout. Flickner Depo at 159 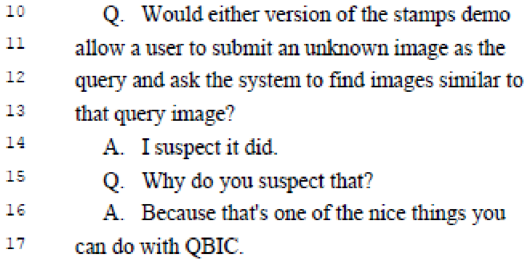 Q. Would either version of the stamps demo allow a user to submit an unknown image as the query and ask the system to find images similar to that query image? A. I suspect it did. Q. Why do you suspect that? A. Because that's one of the nice things you can do with QBIC. Flickner Depo at 209-210

Element	'252 Claim Recitation	Exemplary Citations in Reference
		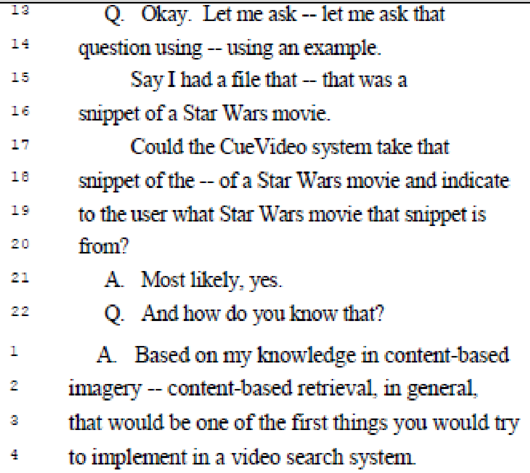 Q. Okay. Let me ask -- let me ask that question using -- using an example. Say I had a file that -- that was a snippet of a Star Wars movie. Could the CueVideo system take that snippet of the -- of a Star Wars movie and indicate to the user what Star Wars movie that snippet is from? A. Most likely, yes. Q. And how do you know that? A. Based on my knowledge in content-based imagery -- content-based retrieval, in general, that would be one of the first things you would try to implement in a video search system. *See also* Flickner Depo at 21, 22-25, 25-26, 54-55, 76, 78-79, 185, 217 Flickner 8/29/23 Depo at 53 – 56 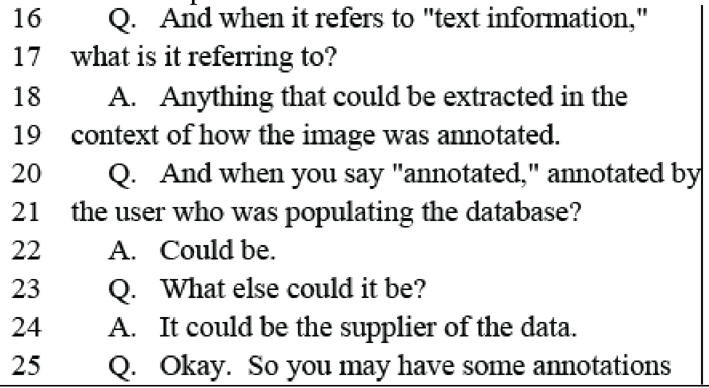 Q. And when it refers to "text information," what is it referring to? A. Anything that could be extracted in the context of how the image was annotated. Q. And when you say "annotated," annotated by the user who was populating the database? A. Could be. Q. What else could it be? A. It could be the supplier of the data. Q. Okay. So you may have some annotations

Element	'252 Claim Recitation	Exemplary Citations in Reference
		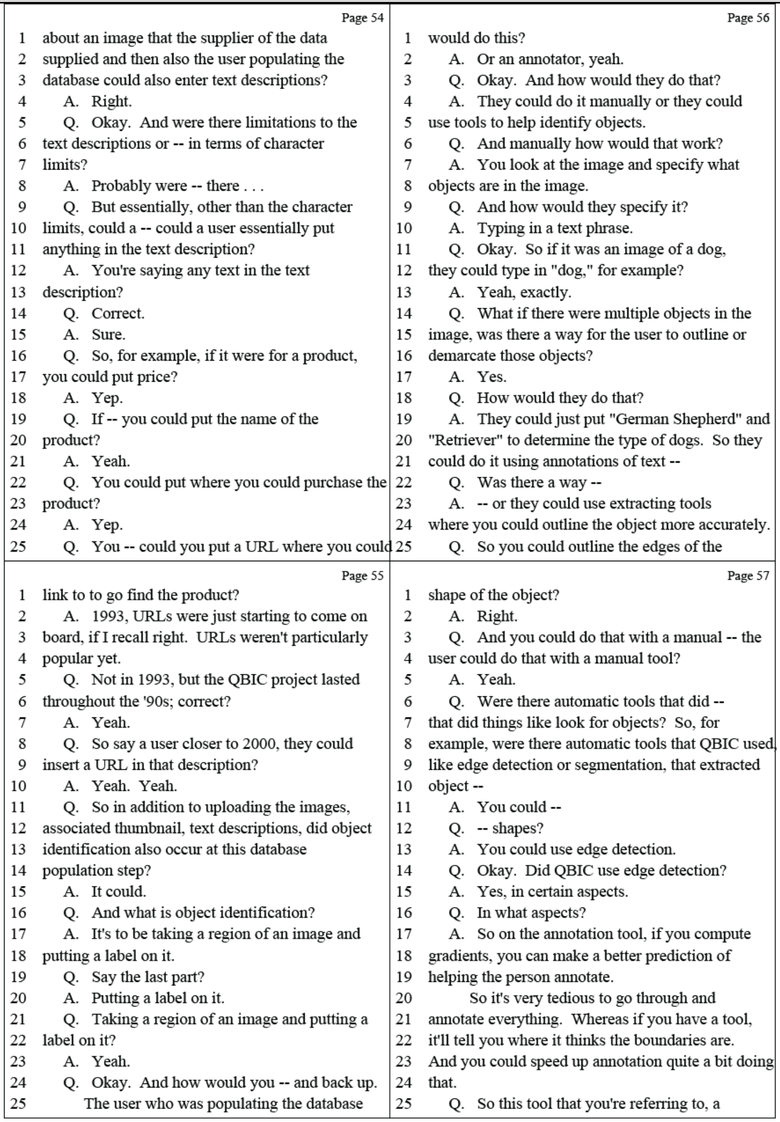

Page 54

about an image that the supplier of the data supplied and then also the user populating the database could also enter text descriptions?
A. Right.
Q. Okay. And were there limitations to the text descriptions or -- in terms of character limits?
A. Probably were -- there . . .
Q. But essentially, other than the character limits, could a -- could a user essentially put anything in the text description?
A. You're saying any text in the text description?
Q. Correct.
A. Sure.
Q. So, for example, if it were for a product, you could put price?
A. Yep.
Q. If -- you could put the name of the product?
A. Yeah.
Q. You could put where you could purchase the product?
A. Yep.
Q. You -- could you put a URL where you could

Page 55

link to to go find the product?
A. 1993, URLs were just starting to come on board, if I recall right. URLs weren't particularly popular yet.
Q. Not in 1993, but the QBIC project lasted throughout the '90s; correct?
A. Yeah.
Q. So say a user closer to 2000, they could insert a URL in that description?
A. Yeah. Yeah.
Q. So in addition to uploading the images, associated thumbnail, text descriptions, did object identification also occur at this database population step?
A. It could.
Q. And what is object identification?
A. It's to be taking a region of an image and putting a label on it.
Q. Say the last part?
A. Putting a label on it.
Q. Taking a region of an image and putting a label on it?
A. Yeah.
Q. Okay. And how would you -- and back up.
The user who was populating the database

Page 56

would do this?
A. Or an annotator, yeah.
Q. Okay. And how would they do that?
A. They could do it manually or they could use tools to help identify objects.
Q. And manually how would that work?
A. You look at the image and specify what objects are in the image.
Q. And how would they specify it?
A. Typing in a text phrase.
Q. Okay. So if it was an image of a dog, they could type in "dog," for example?
A. Yeah, exactly.
Q. What if there were multiple objects in the image, was there a way for the user to outline or demarcate those objects?
A. Yes.
Q. How would they do that?
A. They could just put "German Shepherd" and "Retriever" to determine the type of dogs. So they could do it using annotations of text --
Q. Was there a way --
A. -- or they could use extracting tools where you could outline the object more accurately.
Q. So you could outline the edges of the

Page 57

shape of the object?
A. Right.
Q. And you could do that with a manual -- the user could do that with a manual tool?
A. Yeah.
Q. Were there automatic tools that did -- that did things like look for objects? So, for example, were there automatic tools that QBIC used, like edge detection or segmentation, that extracted object --
A. You could --
Q. -- shapes?
A. You could use edge detection.
Q. Okay. Did QBIC use edge detection?
A. Yes, in certain aspects.
Q. In what aspects?
A. So on the annotation tool, if you compute gradients, you can make a better prediction of helping the person annotate.
So it's very tedious to go through and annotate everything. Whereas if you have a tool, it'll tell you where it thinks the boundaries are. And you could speed up annotation quite a bit doing that.
Q. So this tool that you're referring to, a

Element	'252 Claim Recitation	Exemplary Citations in Reference
		Id. at 60 – 62 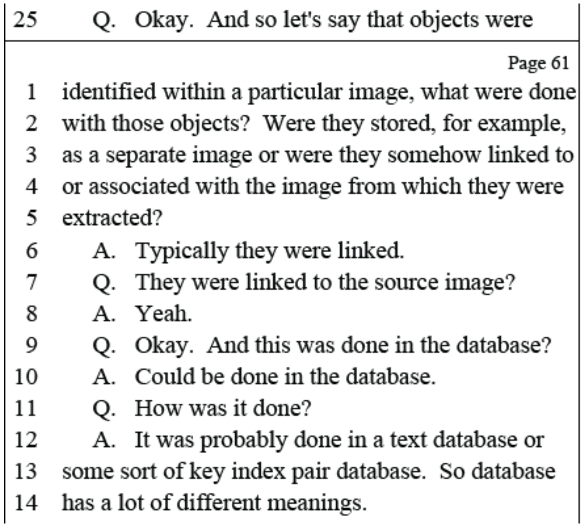 Q. Okay. And so let's say that objects were Page 61 identified within a particular image, what were done with those objects? Were they stored, for example, as a separate image or were they somehow linked to or associated with the image from which they were extracted? A. Typically they were linked. Q. They were linked to the source image? A. Yeah. Q. Okay. And this was done in the database? A. Could be done in the database. Q. How was it done? A. It was probably done in a text database or some sort of key index pair database. So database has a lot of different meanings. 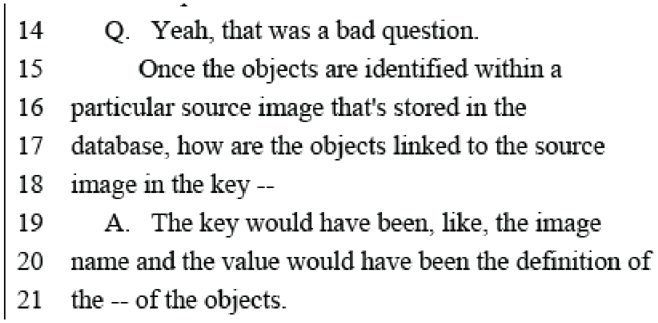 Q. Yeah, that was a bad question. Once the objects are identified within a particular source image that's stored in the database, how are the objects linked to the source image in the key -- A. The key would have been, like, the image name and the value would have been the definition of the -- of the objects. *Id.* at 103

Element	'252 Claim Recitation	Exemplary Citations in Reference
		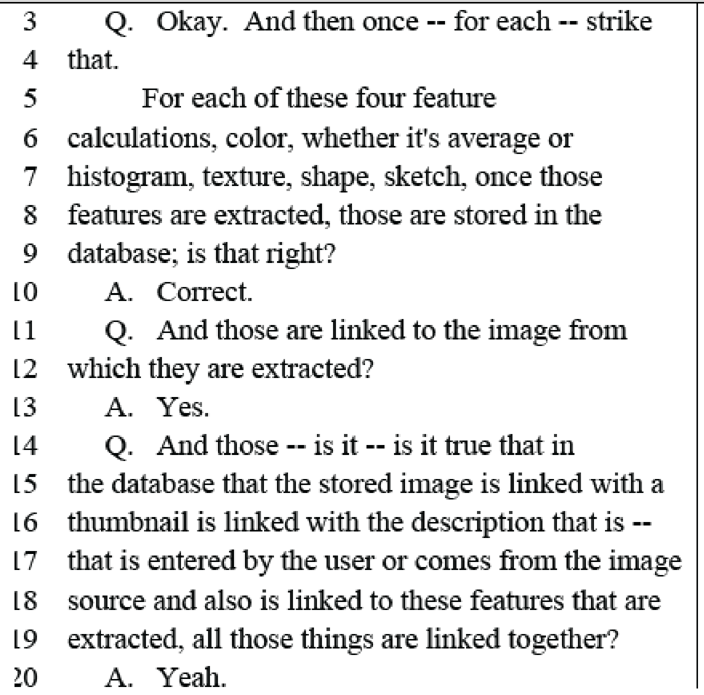 Q. Okay. And then once -- for each -- strike that. For each of these four feature calculations, color, whether it's average or histogram, texture, shape, sketch, once those features are extracted, those are stored in the database; is that right? A. Correct. Q. And those are linked to the image from which they are extracted? A. Yes. Q. And those -- is it -- is it true that in the database that the stored image is linked with a thumbnail is linked with the description that is -- that is entered by the user or comes from the image source and also is linked to these features that are extracted, all those things are linked together? A. Yeah. *Id.* at 129 – 131

Element	'252 Claim Recitation	Exemplary Citations in Reference
		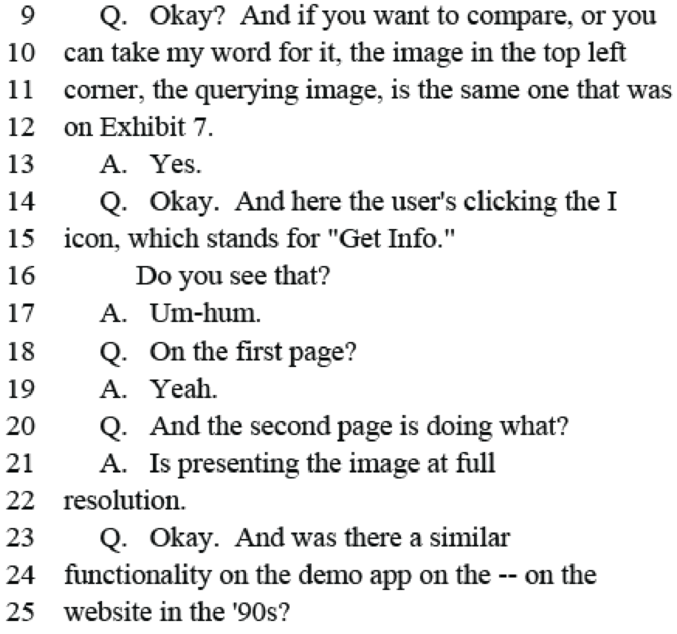 Q. Okay? And if you want to compare, or you can take my word for it, the image in the top left corner, the querying image, is the same one that was on Exhibit 7. A. Yes. Q. Okay. And here the user's clicking the I icon, which stands for "Get Info." Do you see that? A. Um-hum. Q. On the first page? A. Yeah. Q. And the second page is doing what? A. Is presenting the image at full resolution. Q. Okay. And was there a similar functionality on the demo app on the -- on the website in the '90s?

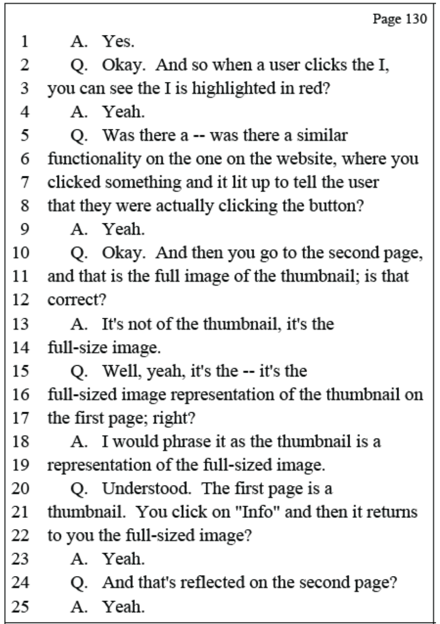

Page 130

A. Yes.
Q. Okay. And so when a user clicks the I, you can see the I is highlighted in red?
A. Yeah.
Q. Was there a -- was there a similar functionality on the one on the website, where you clicked something and it lit up to tell the user that they were actually clicking the button?
A. Yeah.
Q. Okay. And then you go to the second page, and that is the full image of the thumbnail; is that correct?
A. It's not of the thumbnail, it's the full-size image.
Q. Well, yeah, it's the -- it's the full-sized image representation of the thumbnail on the first page; right?
A. I would phrase it as the thumbnail is a representation of the full-sized image.
Q. Understood. The first page is a thumbnail. You click on "Info" and then it returns to you the full-sized image?
A. Yeah.
Q. And that's reflected on the second page?
A. Yeah.

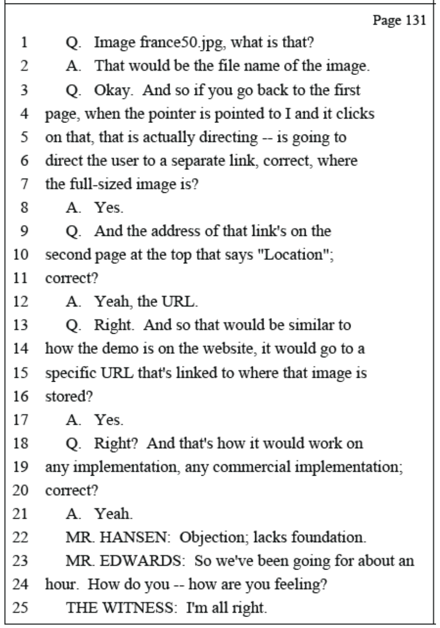

Page 131

Q. Image france50.jpg, what is that?
A. That would be the file name of the image.
Q. Okay. And so if you go back to the first page, when the pointer is pointed to I and it clicks on that, that is actually directing -- is going to direct the user to a separate link, correct, where the full-sized image is?
A. Yes.
Q. And the address of that link's on the second page at the top that says "Location"; correct?
A. Yeah, the URL.
Q. Right. And so that would be similar to how the demo is on the website, it would go to a specific URL that's linked to where that image is stored?
A. Yes.
Q. Right? And that's how it would work on any implementation, any commercial implementation; correct?
A. Yeah.
MR. HANSEN: Objection; lacks foundation.
MR. EDWARDS: So we've been going for about an hour. How do you -- how are you feeling?
THE WITNESS: I'm all right.

Element	'252 Claim Recitation	Exemplary Citations in Reference
		Id. at 161 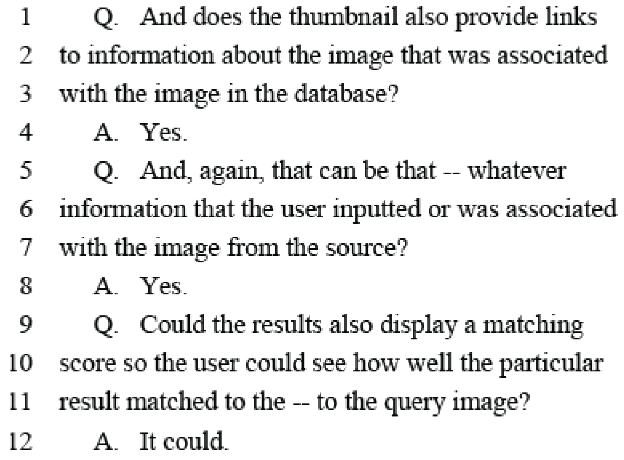 Q. And does the thumbnail also provide links to information about the image that was associated with the image in the database? A. Yes. Q. And, again, that can be that -- whatever information that the user inputted or was associated with the image from the source? A. Yes. Q. Could the results also display a matching score so the user could see how well the particular result matched to the -- to the query image? A. It could. *Id.* at 321 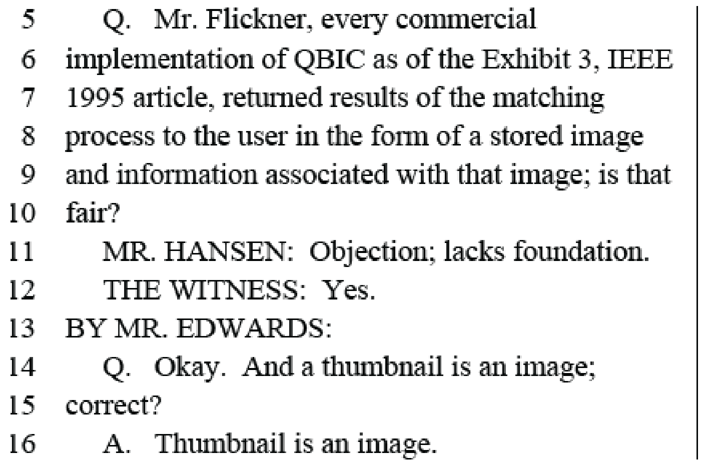 Q. Mr. Flickner, every commercial implementation of QBIC as of the Exhibit 3, IEEE 1995 article, returned results of the matching process to the user in the form of a stored image and information associated with that image; is that fair? MR. HANSEN: Objection; lacks foundation. THE WITNESS: Yes. BY MR. EDWARDS: Q. Okay. And a thumbnail is an image; correct? A. Thumbnail is an image.

Element	'252 Claim Recitation	Exemplary Citations in Reference
		To the extent it is found that QBIC does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that QBIC does not anticipate this claim, QBIC renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine QBIC with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results. For example, the '252 patent acknowledges that acquiring information about objects through symbols, such as barcodes or some other means to encode detectable information in the image or on the object, were traditional methods known in the art and not part the present invention. '252 Patent, 3:7-16. Indeed, the '252 patent acknowledges that algorithms, software, and hardware components were commercially available at the time of the invention that could detect and decode symbols, such as barcodes or text, in an input image. *Id*. at 11:56-61. This limitation was thus known in the art by patentee's own admission, and by asserting this claim, Plaintiff is attempting to reclaim disavowed claim scope and claim subject matter which patentee has acknowledged is not part of the invention.
[18D]	allowing the mobile device to receive the pertinent information over a network.	QBIC discloses allowing the mobile device to receive the pertinent information over a network. The QBIC Project: Querying Images by Content Using Color, Texture, and Shape (IBM 0002390–404) at 8 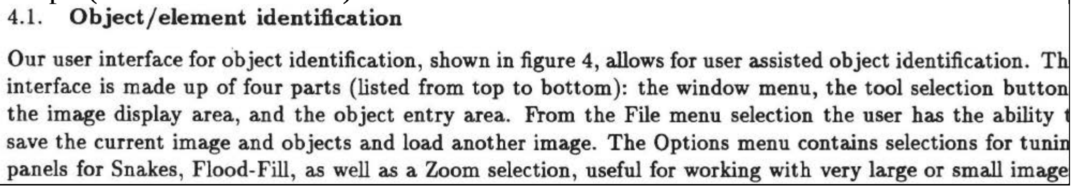 4.1. **Object/element identification** Our user interface for object identification, shown in figure 4, allows for user assisted object identification. Th interface is made up of four parts (listed from top to bottom): the window menu, the tool selection button the image display area, and the object entry area. From the File menu selection the user has the ability t save the current image and objects and load another image. The Options menu contains selections for tunin panels for Snakes, Flood-Fill, as well as a Zoom selection, useful for working with very large or small image

Element	'252 Claim Recitation	Exemplary Citations in Reference
		The QBIC Project: Querying Images by Content Using Color, Texture, and Shape (IBM 0002390–404) at Fig. 4 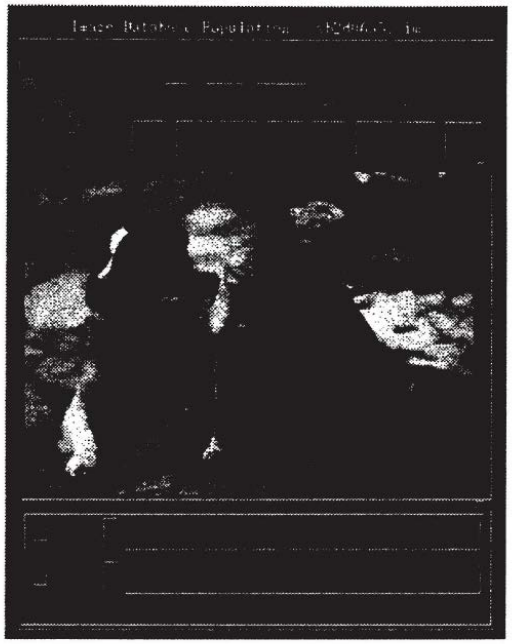 Figure 4: Object Identification Tool. Note for example, the right dog's ear, which has been outlined. The QBIC Project: Querying Images by Content Using Color, Texture, and Shape (IBM 0002390–404) at 8

Element	'252 Claim Recitation	Exemplary Citations in Reference
		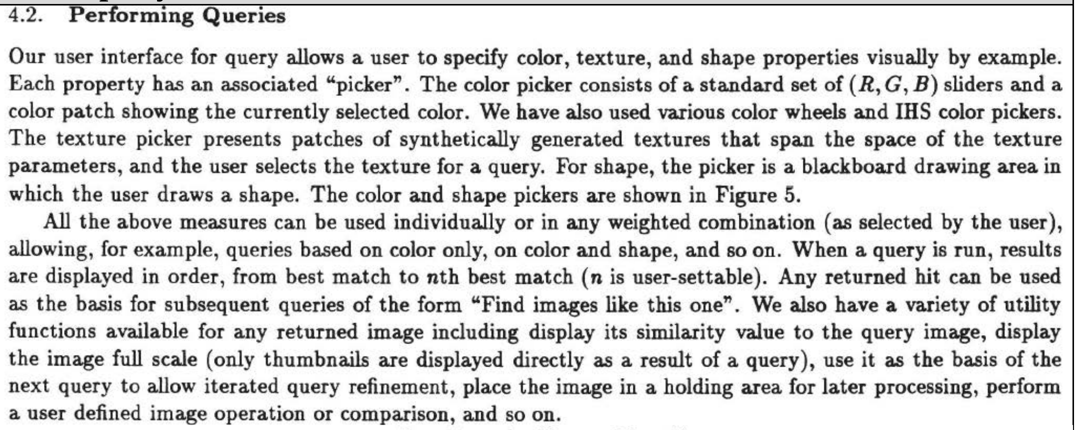 4.2. **Performing Queries** Our user interface for query allows a user to specify color, texture, and shape properties visually by example. Each property has an associated "picker". The color picker consists of a standard set of (R, G, B) sliders and a color patch showing the currently selected color. We have also used various color wheels and IHS color pickers. The texture picker presents patches of synthetically generated textures that span the space of the texture parameters, and the user selects the texture for a query. For shape, the picker is a blackboard drawing area in which the user draws a shape. The color and shape pickers are shown in Figure 5. All the above measures can be used individually or in any weighted combination (as selected by the user), allowing, for example, queries based on color only, on color and shape, and so on. When a query is run, results are displayed in order, from best match to nth best match (n is user-settable). Any returned hit can be used as the basis for subsequent queries of the form "Find images like this one". We also have a variety of utility functions available for any returned image including display its similarity value to the query image, display the image full scale (only thumbnails are displayed directly as a result of a query), use it as the basis of the next query to allow iterated query refinement, place the image in a holding area for later processing, perform a user defined image operation or comparison, and so on. 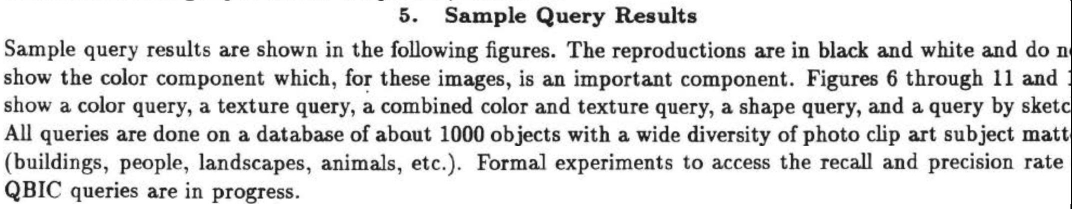 5. **Sample Query Results** Sample query results are shown in the following figures. The reproductions are in black and white and do n show the color component which, for these images, is an important component. Figures 6 through 11 and show a color query, a texture query, a combined color and texture query, a shape query, and a query by sketc All queries are done on a database of about 1000 objects with a wide diversity of photo clip art subject matt (buildings, people, landscapes, animals, etc.). Formal experiments to access the recall and precision rate QBIC queries are in progress. The QBIC Project: Querying Images by Content Using Color, Texture, and Shape (IBM 0002390–404) at Figs. 7–9

Element	'252 Claim Recitation	Exemplary Citations in Reference
		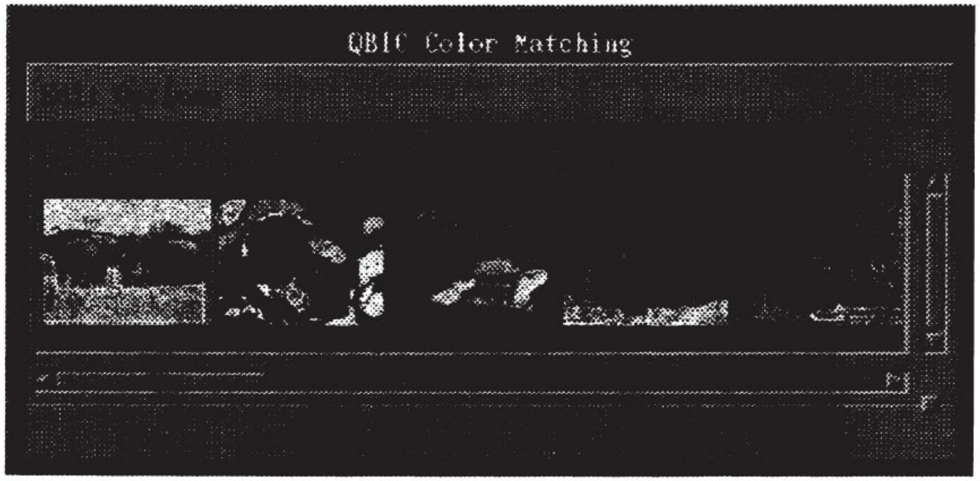 Figure 7: Result of query by example using color only.

Element	'252 Claim Recitation	Exemplary Citations in Reference
		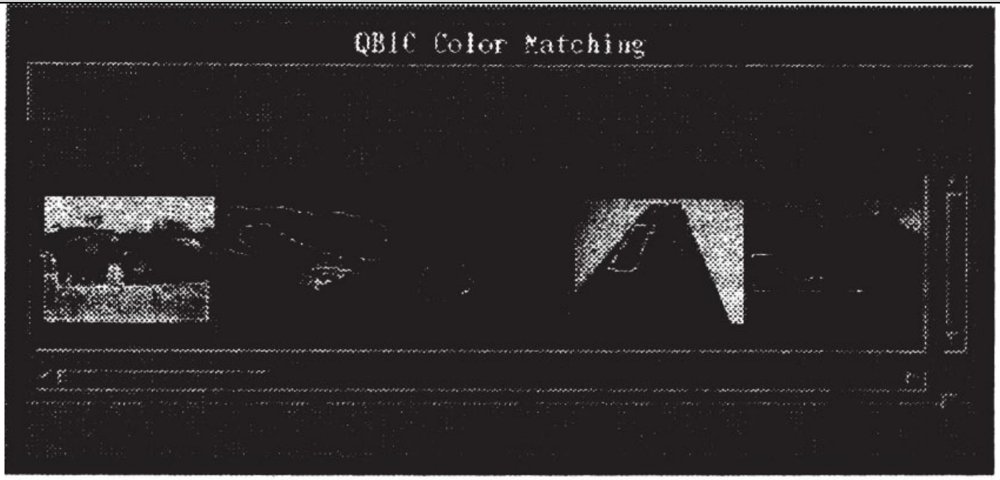 Figure 8: Result of query by example using texture only.

Element	'252 Claim Recitation	Exemplary Citations in Reference
		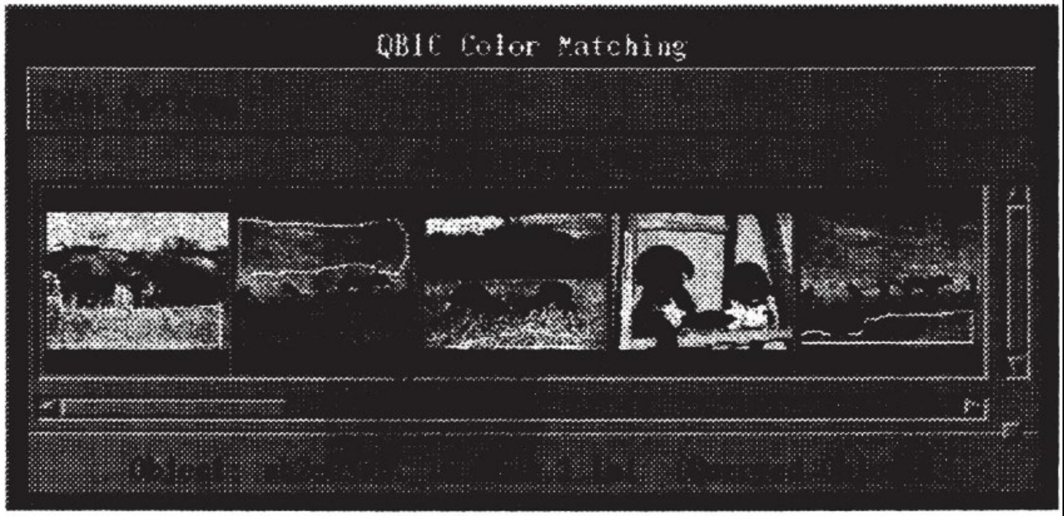 Figure 9: Result of query by example using color and texture. QBIC (Query by Image Content) at 1 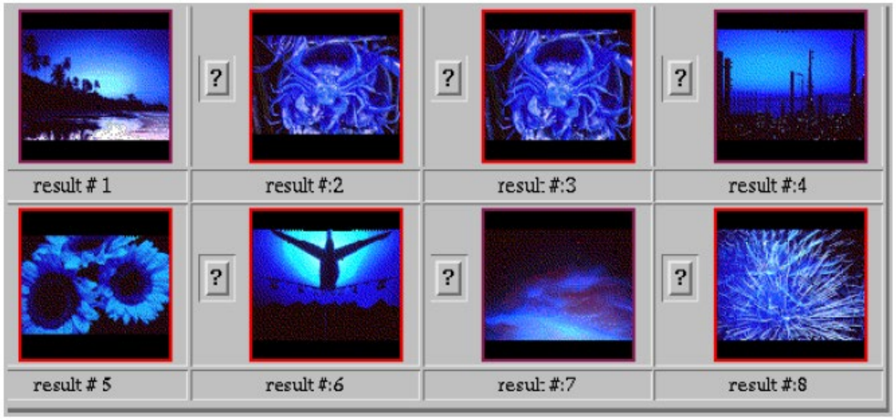 **Result of QBIC Search**

Element	'252 Claim Recitation	Exemplary Citations in Reference
		Query by Image and Video Content: The QBIC System (IBM 0000893–902): "QBIC has two main components: database population (the process of creating an image database) and database query. During the population, images and videos are processed to extract features describing their content-colors, textures, shapes, and camera and object motion-and the features are stored in a database. During the query, the user composes a query graphically. Features are generated from the graphical query and then input to a matching engine that finds images or videos from the database with similar features. Figure 2 shows the system architecture." *Id.* at p. 3. 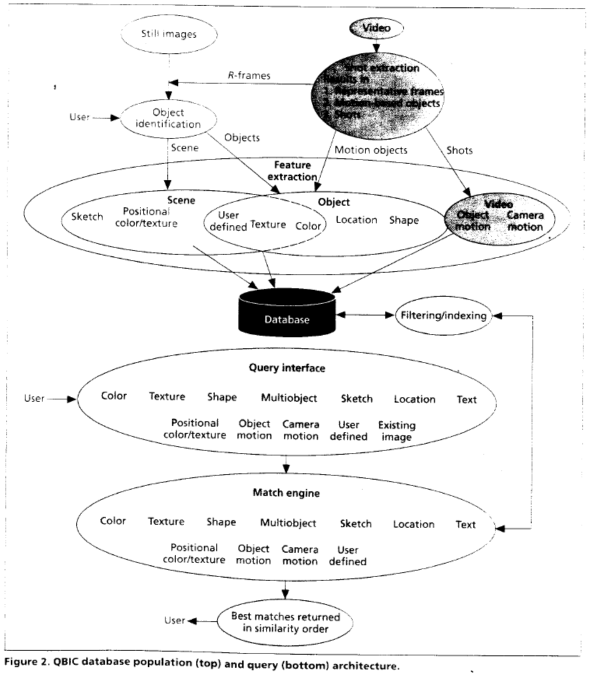 Figure 2. QBIC database population (top) and query (bottom) architecture. *Id.* at Fig. 2. Image, Audio, and Video Extenders Administration and Programming (IBM 000001–617):

Element	'252 Claim Recitation	Exemplary Citations in Reference
		"To make an image available for searching by content, you catalog the image. When you catalog an image, the Image Extender analyzes the image, by computing the feature values for the image, and stores the values in a QBIC catalog. When you search for an image by content, your query identifies one or more features for the search (such as average color), a source for each feature (such as an example image), and a target set of cataloged images. The Image Extender computes the feature value of the source and compares it to the cataloged feature values for the target images. It then computes a score that indicates how similar the feature values of the target images are to the source. You can have the Image Extender return the images whose features are most similar to the source. The Image Extender will return the handle of each image and the image score. You can also have the Image Extender return only the score of a single image." *Id.* at pp. 20–21. Managing Enterprise Information Portal (IBM 000777–880): "Many paper-intensive enterprises, such as insurance companies and financial institutions, administer large volumes of business-related content. The need for an enterprise solution for managing and accessing business information spans many industries. A content server is a software system that stores multimedia, business forms, documents, and related data, along with metadata that allows employees to process and work with the content." *Id.* at p. 3. "The image search server uses IBM's QBIC (query by image content) technology to help you search for objects by certain visual properties, such as color and texture. The image search server analyzes images and stores the image information in a database. Then users can run image queries, which use the visual properties of images, to match colors, textures, and their positions without describing them in words. You can combine content-based queries with text and keyword searches for more powerful retrieval of image and multimedia data." *Id.* at p. 63.

Element	'252 Claim Recitation	Exemplary Citations in Reference
		QBIC Demo (IBM000747) 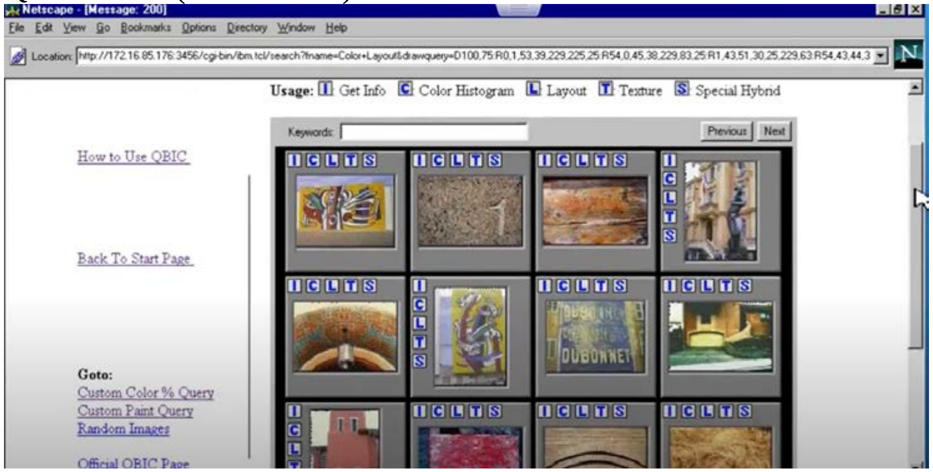 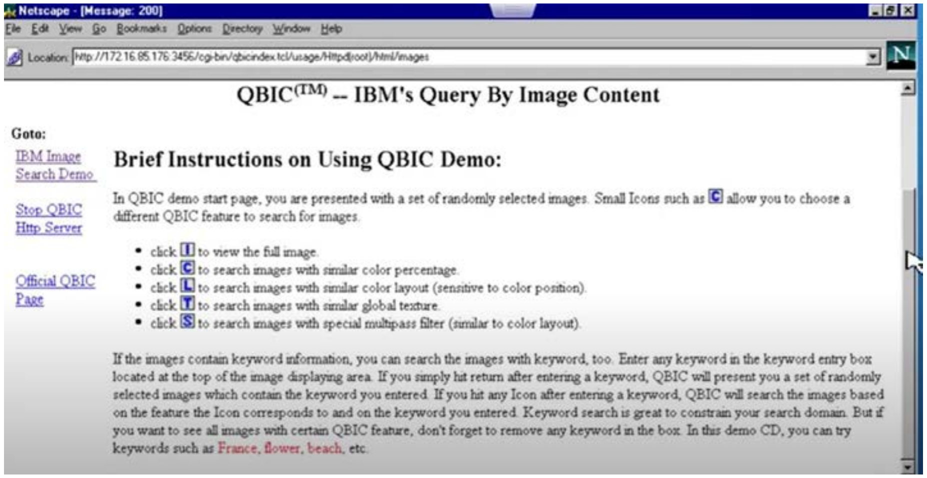 QBIC(TM) -- IBM's Query By Image Content **Brief Instructions on Using QBIC Demo:** In QBIC demo start page, you are presented with a set of randomly selected images. Small Icons such as C allow you to choose a different QBIC feature to search for images. - click I to view the full image. - click C to search images with similar color percentage. - click L to search images with similar color layout (sensitive to color position). - click T to search images with similar global texture. - click S to search images with special multipass filter (similar to color layout). If the images contain keyword information, you can search the images with keyword, too. Enter any keyword in the keyword entry box located at the top of the image displaying area. If you simply hit return after entering a keyword, QBIC will present you a set of randomly selected images which contain the keyword you entered. If you hit any Icon after entering a keyword, QBIC will search the images based on the feature the Icon corresponds to and on the keyword you entered. Keyword search is great to constrain your search domain. But if you want to see all images with certain QBIC feature, don't forget to remove any keyword in the box. In this demo CD, you can try keywords such as France, flower, beach, etc.

Element	'252 Claim Recitation	Exemplary Citations in Reference
		 Flickner Depo at 41-42 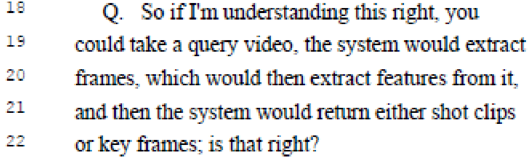 Q. So if I'm understanding this right, you could take a query video, the system would extract frames, which would then extract features from it, and then the system would return either shot clips or key frames; is that right? 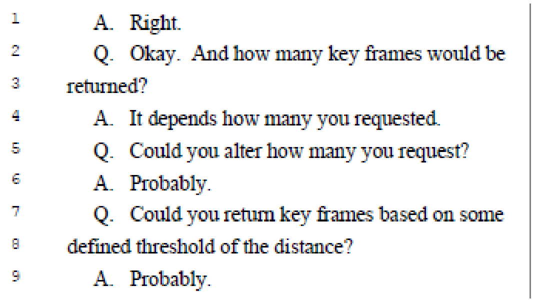 A. Right. Q. Okay. And how many key frames would be returned? A. It depends how many you requested. Q. Could you alter how many you request? A. Probably. Q. Could you return key frames based on some defined threshold of the distance? A. Probably. Flickner Depo at 50

Element	'252 Claim Recitation	Exemplary Citations in Reference
		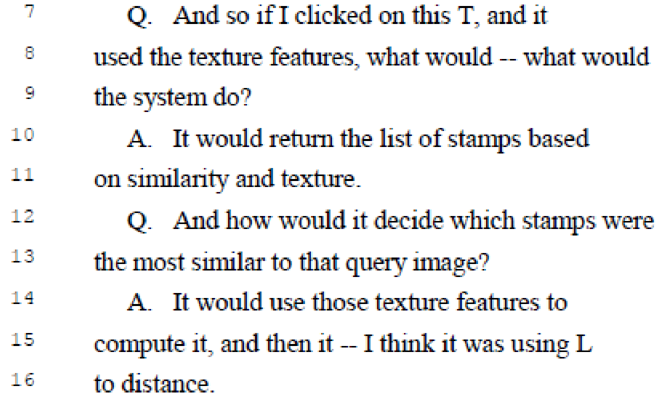 Q. And so if I clicked on this T, and it used the texture features, what would -- what would the system do? A. It would return the list of stamps based on similarity and texture. Q. And how would it decide which stamps were the most similar to that query image? A. It would use those texture features to compute it, and then it -- I think it was using L to distance. Flickner Depo at 131-132 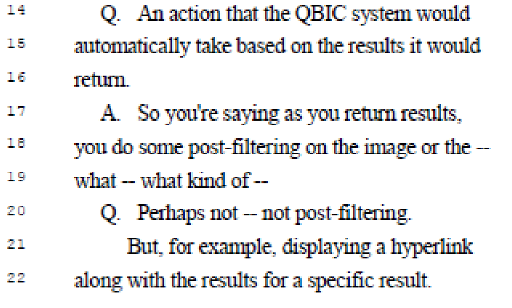 Q. An action that the QBIC system would automatically take based on the results it would return. A. So you're saying as you return results, you do some post-filtering on the image or the -- what -- what kind of -- Q. Perhaps not -- not post-filtering. But, for example, displaying a hyperlink along with the results for a specific result.

Element	'252 Claim Recitation	Exemplary Citations in Reference
		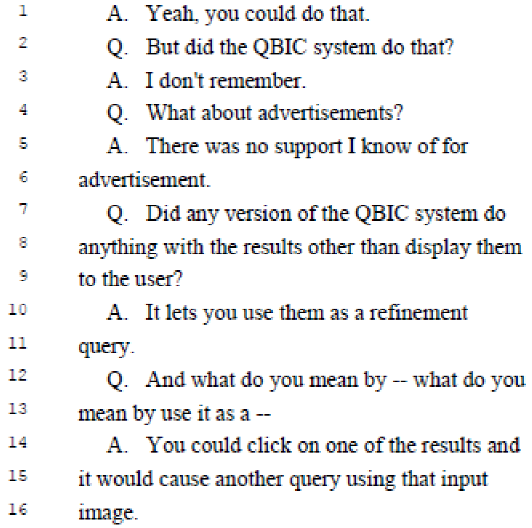 A. Yeah, you could do that. Q. But did the QBIC system do that? A. I don't remember. Q. What about advertisements? A. There was no support I know of for advertisement. Q. Did any version of the QBIC system do anything with the results other than display them to the user? A. It lets you use them as a refinement query. Q. And what do you mean by -- what do you mean by use it as a -- A. You could click on one of the results and it would cause another query using that input image. *See also* Flickner Depo at 22-25, 25-26, 40-42, 54-55, 76, 78-79, 87, 88, 89-90, 159, 162-163, 201-202, 209-210 Flickner 8/31/23 Depo at 68–72

Element	'252 Claim Recitation	Exemplary Citations in Reference
		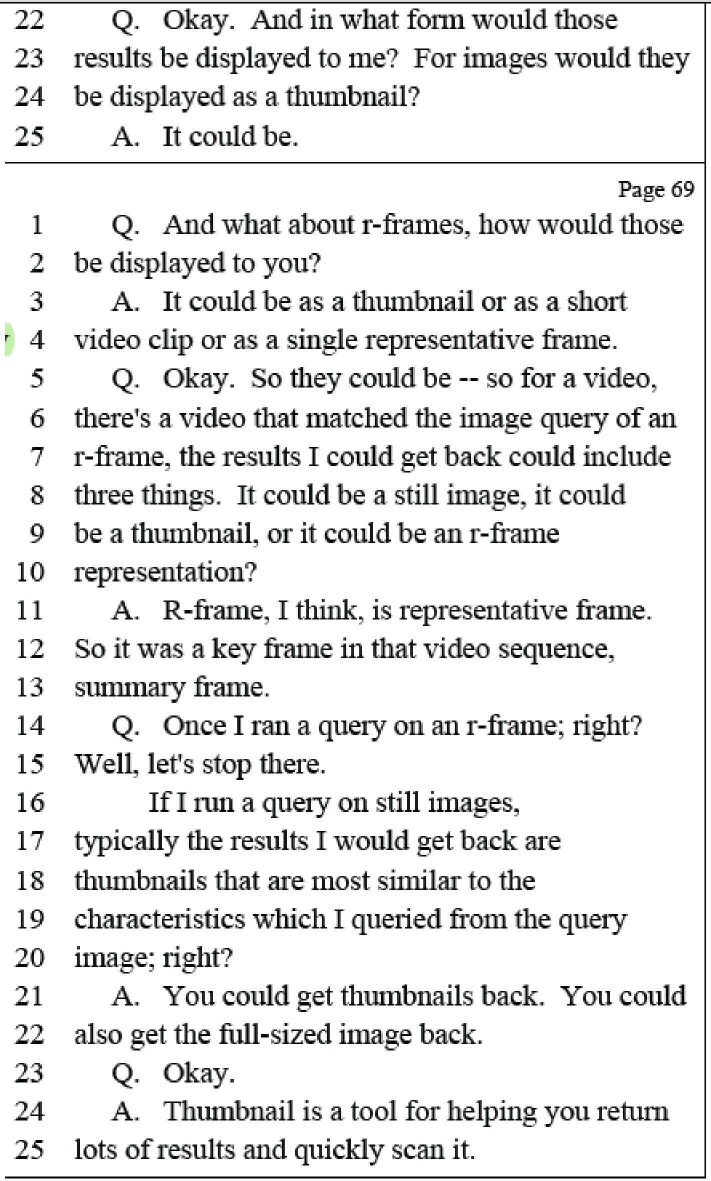 Q. Okay. And in what form would those results be displayed to me? For images would they be displayed as a thumbnail? A. It could be. Page 69 Q. And what about r-frames, how would those be displayed to you? A. It could be as a thumbnail or as a short video clip or as a single representative frame. Q. Okay. So they could be -- so for a video, there's a video that matched the image query of an r-frame, the results I could get back could include three things. It could be a still image, it could be a thumbnail, or it could be an r-frame representation? A. R-frame, I think, is representative frame. So it was a key frame in that video sequence, summary frame. Q. Once I ran a query on an r-frame; right? Well, let's stop there. If I run a query on still images, typically the results I would get back are thumbnails that are most similar to the characteristics which I queried from the query image; right? A. You could get thumbnails back. You could also get the full-sized image back. Q. Okay. A. Thumbnail is a tool for helping you return lots of results and quickly scan it.

Element	'252 Claim Recitation	Exemplary Citations in Reference
		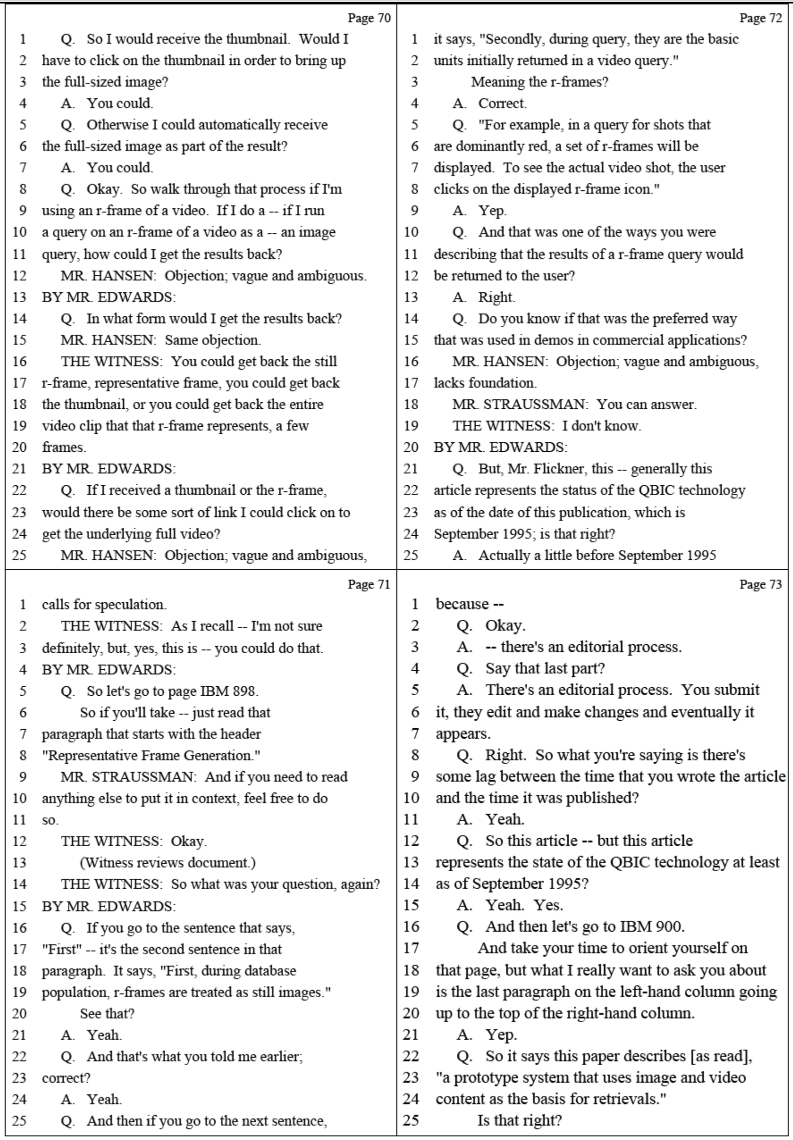

Page 70

Q. So I would receive the thumbnail. Would I have to click on the thumbnail in order to bring up the full-sized image?

A. You could.

Q. Otherwise I could automatically receive the full-sized image as part of the result?

A. You could.

Q. Okay. So walk through that process if I'm using an r-frame of a video. If I do a -- if I run a query on an r-frame of a video as a -- an image query, how could I get the results back?

MR. HANSEN: Objection; vague and ambiguous.

BY MR. EDWARDS:

Q. In what form would I get the results back?

MR. HANSEN: Same objection.

THE WITNESS: You could get back the still r-frame, representative frame, you could get back the thumbnail, or you could get back the entire video clip that that r-frame represents, a few frames.

BY MR. EDWARDS:

Q. If I received a thumbnail or the r-frame, would there be some sort of link I could click on to get the underlying full video?

MR. HANSEN: Objection; vague and ambiguous,

Page 71

calls for speculation.

THE WITNESS: As I recall -- I'm not sure definitely, but, yes, this is -- you could do that.

BY MR. EDWARDS:

Q. So let's go to page IBM 898.

So if you'll take -- just read that paragraph that starts with the header "Representative Frame Generation."

MR. STRAUSSMAN: And if you need to read anything else to put it in context, feel free to do so.

THE WITNESS: Okay.

(Witness reviews document.)

THE WITNESS: So what was your question, again?

BY MR. EDWARDS:

Q. If you go to the sentence that says, "First" -- it's the second sentence in that paragraph. It says, "First, during database population, r-frames are treated as still images."

See that?

A. Yeah.

Q. And that's what you told me earlier; correct?

A. Yeah.

Q. And then if you go to the next sentence,

Page 72

it says, "Secondly, during query, they are the basic units initially returned in a video query."

Meaning the r-frames?

A. Correct.

Q. "For example, in a query for shots that are dominantly red, a set of r-frames will be displayed. To see the actual video shot, the user clicks on the displayed r-frame icon."

A. Yep.

Q. And that was one of the ways you were describing that the results of a r-frame query would be returned to the user?

A. Right.

Q. Do you know if that was the preferred way that was used in demos in commercial applications?

MR. HANSEN: Objection; vague and ambiguous, lacks foundation.

MR. STRAUSSMAN: You can answer.

THE WITNESS: I don't know.

BY MR. EDWARDS:

Q. But, Mr. Flickner, this -- generally this article represents the status of the QBIC technology as of the date of this publication, which is September 1995; is that right?

A. Actually a little before September 1995

Page 73

because --

Q. Okay.

A. -- there's an editorial process.

Q. Say that last part?

A. There's an editorial process. You submit it, they edit and make changes and eventually it appears.

Q. Right. So what you're saying is there's some lag between the time that you wrote the article and the time it was published?

A. Yeah.

Q. So this article -- but this article represents the state of the QBIC technology at least as of September 1995?

A. Yeah. Yes.

Q. And then let's go to IBM 900.

And take your time to orient yourself on that page, but what I really want to ask you about is the last paragraph on the left-hand column going up to the top of the right-hand column.

A. Yep.

Q. So it says this paper describes [as read], "a prototype system that uses image and video content as the basis for retrievals."

Is that right?

Element	'252 Claim Recitation	Exemplary Citations in Reference
		Id. at 129 – 131 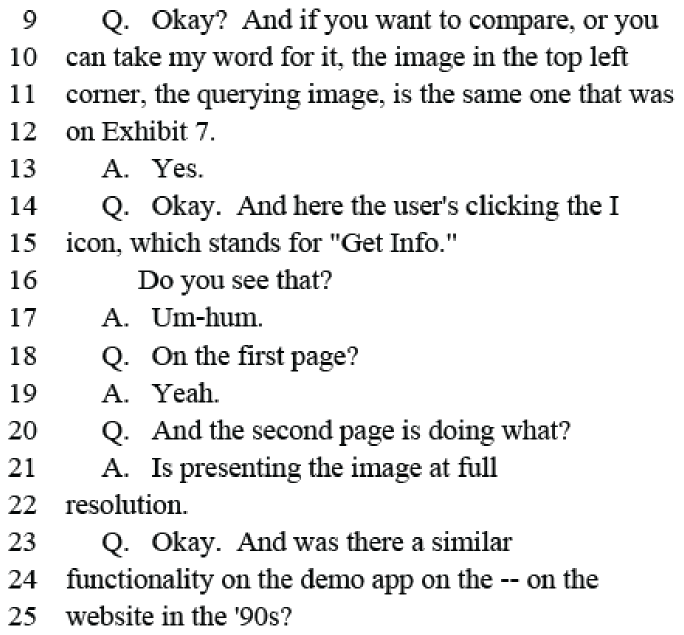 Q. Okay? And if you want to compare, or you can take my word for it, the image in the top left corner, the querying image, is the same one that was on Exhibit 7. A. Yes. Q. Okay. And here the user's clicking the I icon, which stands for "Get Info." Do you see that? A. Um-hum. Q. On the first page? A. Yeah. Q. And the second page is doing what? A. Is presenting the image at full resolution. Q. Okay. And was there a similar functionality on the demo app on the -- on the website in the '90s?

		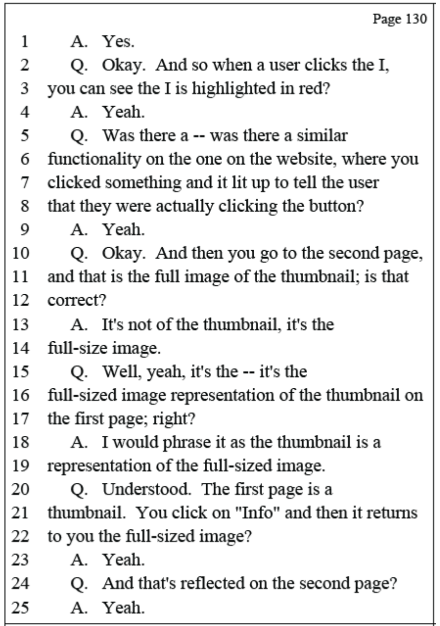 Page 130 A. Yes. Q. Okay. And so when a user clicks the I, you can see the I is highlighted in red? A. Yeah. Q. Was there a -- was there a similar functionality on the one on the website, where you clicked something and it lit up to tell the user that they were actually clicking the button? A. Yeah. Q. Okay. And then you go to the second page, and that is the full image of the thumbnail; is that correct? A. It's not of the thumbnail, it's the full-size image. Q. Well, yeah, it's the -- it's the full-sized image representation of the thumbnail on the first page; right? A. I would phrase it as the thumbnail is a representation of the full-sized image. Q. Understood. The first page is a thumbnail. You click on "Info" and then it returns to you the full-sized image? A. Yeah. Q. And that's reflected on the second page? A. Yeah. 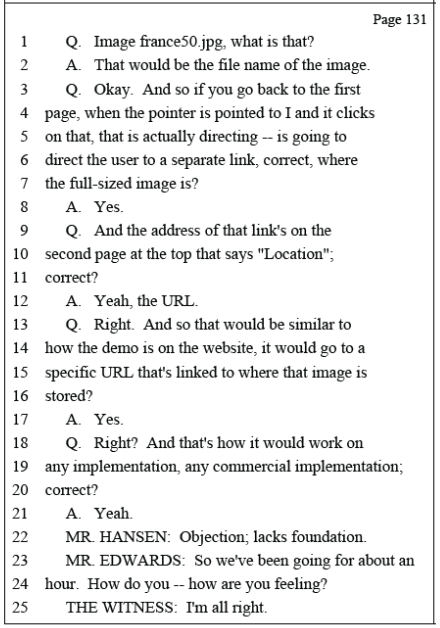 Page 131 Q. Image france50.jpg, what is that? A. That would be the file name of the image. Q. Okay. And so if you go back to the first page, when the pointer is pointed to I and it clicks on that, that is actually directing -- is going to direct the user to a separate link, correct, where the full-sized image is? A. Yes. Q. And the address of that link's on the second page at the top that says "Location"; correct? A. Yeah, the URL. Q. Right. And so that would be similar to how the demo is on the website, it would go to a specific URL that's linked to where that image is stored? A. Yes. Q. Right? And that's how it would work on any implementation, any commercial implementation; correct? A. Yeah. MR. HANSEN: Objection; lacks foundation. MR. EDWARDS: So we've been going for about an hour. How do you -- how are you feeling? THE WITNESS: I'm all right.

Element	'252 Claim Recitation	Exemplary Citations in Reference
		Id. at 161 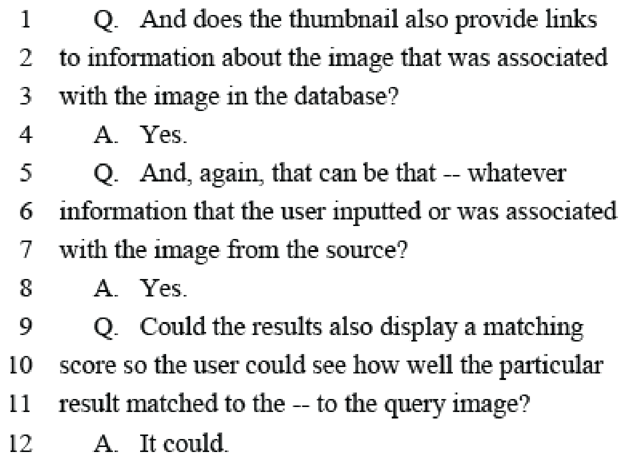 Q. And does the thumbnail also provide links to information about the image that was associated with the image in the database? A. Yes. Q. And, again, that can be that -- whatever information that the user inputted or was associated with the image from the source? A. Yes. Q. Could the results also display a matching score so the user could see how well the particular result matched to the -- to the query image? A. It could. *Id.* at 321 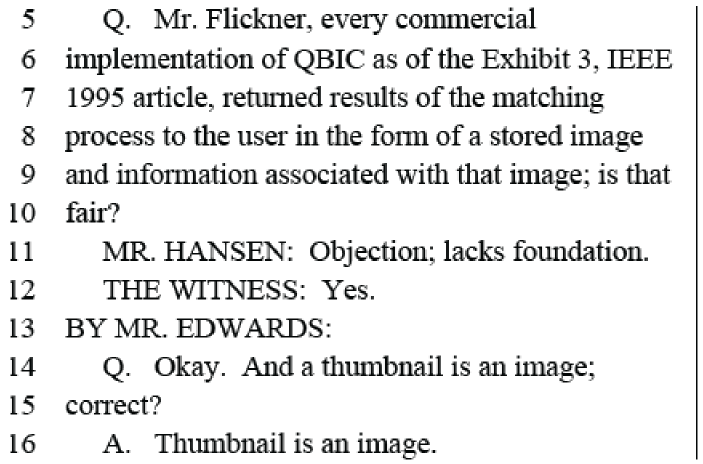 Q. Mr. Flickner, every commercial implementation of QBIC as of the Exhibit 3, IEEE 1995 article, returned results of the matching process to the user in the form of a stored image and information associated with that image; is that fair? MR. HANSEN: Objection; lacks foundation. THE WITNESS: Yes. BY MR. EDWARDS: Q. Okay. And a thumbnail is an image; correct? A. Thumbnail is an image.

Element	'252 Claim Recitation	Exemplary Citations in Reference
		To the extent it is found that QBIC does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that QBIC does not anticipate this claim, QBIC renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine QBIC with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results. For example, the '252 patent acknowledges that acquiring information about objects through symbols, such as barcodes or some other means to encode detectable information in the image or on the object, were traditional methods known in the art and not part the present invention. '252 Patent, 3:7-16. Indeed, the '252 patent acknowledges that algorithms, software, and hardware components were commercially available at the time of the invention that could detect and decode symbols, such as barcodes or text, in an input image. *Id*. at 11:56-61. This limitation was thus known in the art by patentee's own admission, and by asserting this claim, Plaintiff is attempting to reclaim disavowed claim scope and claim subject matter which patentee has acknowledged is not part of the invention.

2. Claim 26

Claim	'252 Claim Recitation	Exemplary Citations in Reference
26	The method of claim 18, wherein the data relating to the image comprises the image.	QBIC discloses the method of claim 18, wherein the data relating to the image comprises the image. The QBIC Project: Querying Images by Content Using Color, Texture, and Shape (IBM 0002390–404) at 8 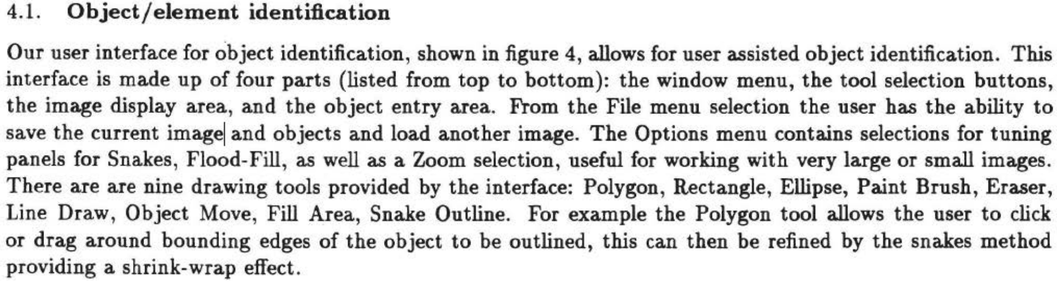 4.1. **Object/element identification** Our user interface for object identification, shown in figure 4, allows for user assisted object identification. This interface is made up of four parts (listed from top to bottom): the window menu, the tool selection buttons, the image display area, and the object entry area. From the File menu selection the user has the ability to save the current image and objects and load another image. The Options menu contains selections for tuning panels for Snakes, Flood-Fill, as well as a Zoom selection, useful for working with very large or small images. There are are nine drawing tools provided by the interface: Polygon, Rectangle, Ellipse, Paint Brush, Eraser, Line Draw, Object Move, Fill Area, Snake Outline. For example the Polygon tool allows the user to click or drag around bounding edges of the object to be outlined, this can then be refined by the snakes method providing a shrink-wrap effect. QBIC Demo (IBM000747) 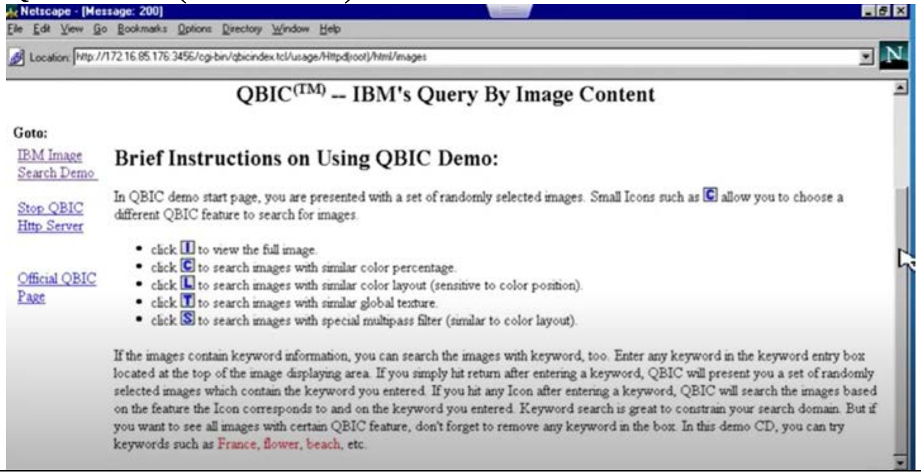 QBIC(TM) -- IBM's Query By Image Content **Brief Instructions on Using QBIC Demo:** In QBIC demo start page, you are presented with a set of randomly selected images. Small Icons such as C allow you to choose a different QBIC feature to search for images. - click I to view the full image. - click C to search images with similar color percentage. - click L to search images with similar color layout (sensitive to color position). - click T to search images with similar global texture. - click S to search images with special multipass filter (similar to color layout). If the images contain keyword information, you can search the images with keyword, too. Enter any keyword in the keyword entry box located at the top of the image displaying area. If you simply hit return after entering a keyword, QBIC will present you a set of randomly selected images which contain the keyword you entered. If you hit any Icon after entering a keyword, QBIC will search the images based on the feature the Icon corresponds to and on the keyword you entered. Keyword search is great to constrain your search domain. But if you want to see all images with certain QBIC feature, don't forget to remove any keyword in the box. In this demo CD, you can try keywords such as France, flower, beach, etc.

Location: http://172.16.85.176:3456/cgi-bin/ibm.tcl/search?fname=HybridColor&france10%2Ejpg.x=0

Goto:
Custom Color % Query
Custom Paint Query
Random Images

Official QBIC Page

Query was:

Example:	=france10.jpg
Query Type:	QbHybridColorFeatureClass

Flickner Depo at 131-133

Q. An action that the QBIC system would automatically take based on the results it would return.

A. So you're saying as you return results, you do some post-filtering on the image or the -- what -- what kind of --

Q. Perhaps not -- not post-filtering.

But, for example, displaying a hyperlink along with the results for a specific result.

		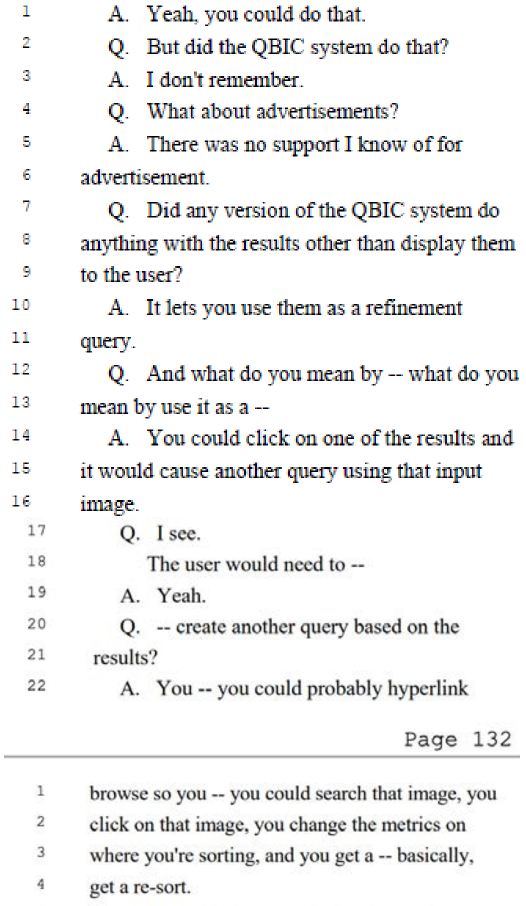 A. Yeah, you could do that. Q. But did the QBIC system do that? A. I don't remember. Q. What about advertisements? A. There was no support I know of for advertisement. Q. Did any version of the QBIC system do anything with the results other than display them to the user? A. It lets you use them as a refinement query. Q. And what do you mean by -- what do you mean by use it as a -- A. You could click on one of the results and it would cause another query using that input image. Q. I see. The user would need to -- A. Yeah. Q. -- create another query based on the results? A. You -- you could probably hyperlink Page 132 browse so you -- you could search that image, you click on that image, you change the metrics on where you're sorting, and you get a -- basically, get a re-sort. *Id.* at 32–33

		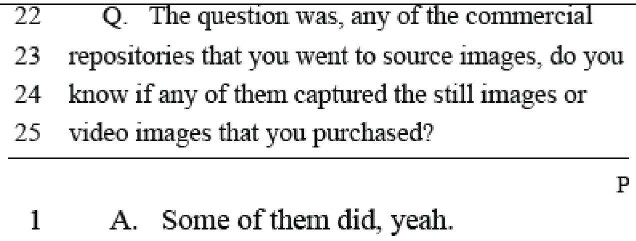 Q. The question was, any of the commercial repositories that you went to source images, do you know if any of them captured the still images or video images that you purchased? P A. Some of them did, yeah. *Id.* at 74–75

			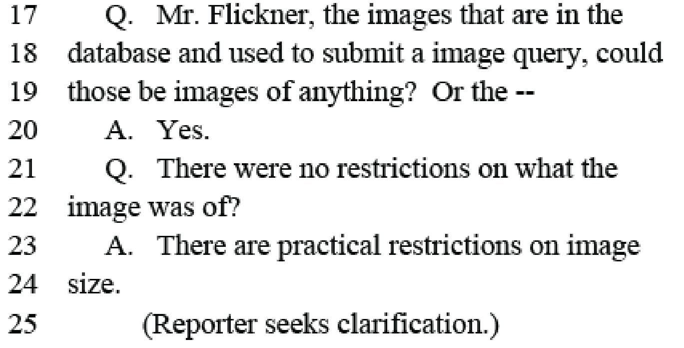 Q. Mr. Flickner, the images that are in the database and used to submit a image query, could those be images of anything? Or the -- A. Yes. Q. There were no restrictions on what the image was of? A. There are practical restrictions on image size. (Reporter seeks clarification.) 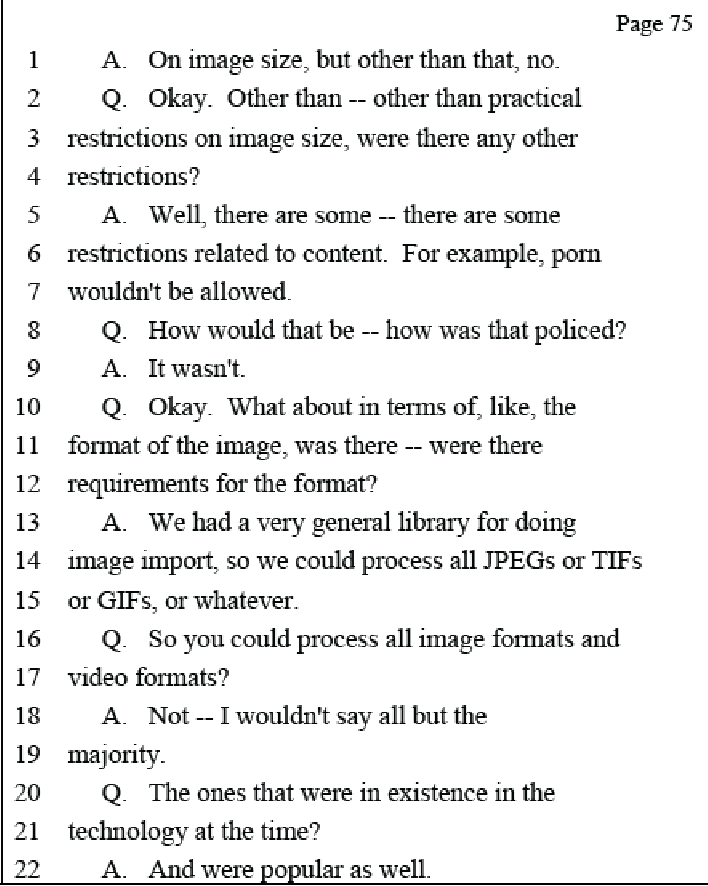 Page 75 A. On image size, but other than that, no. Q. Okay. Other than -- other than practical restrictions on image size, were there any other restrictions? A. Well, there are some -- there are some restrictions related to content. For example, porn wouldn't be allowed. Q. How would that be -- how was that policed? A. It wasn't. Q. Okay. What about in terms of, like, the format of the image, was there -- were there requirements for the format? A. We had a very general library for doing image import, so we could process all JPEGs or TIFs or GIFs, or whatever. Q. So you could process all image formats and video formats? A. Not -- I wouldn't say all but the majority. Q. The ones that were in existence in the technology at the time? A. And were popular as well.

		Id. at 76–77 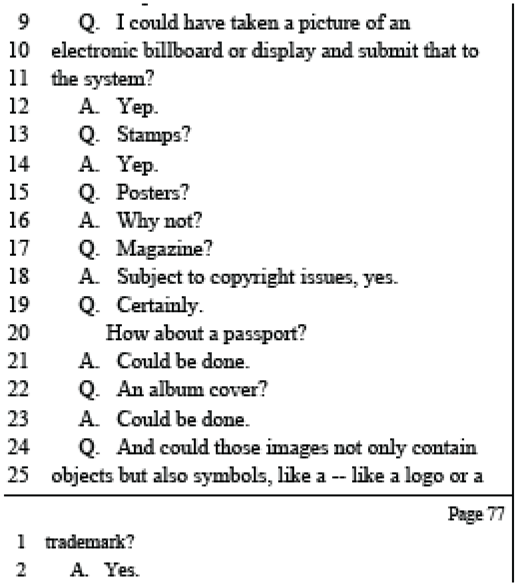 Q. I could have taken a picture of an electronic billboard or display and submit that to the system? A. Yep. Q. Stamps? A. Yep. Q. Posters? A. Why not? Q. Magazine? A. Subject to copyright issues, yes. Q. Certainly. How about a passport? A. Could be done. Q. An album cover? A. Could be done. Q. And could those images not only contain objects but also symbols, like a -- like a logo or a Page 77 trademark? A. Yes. *Id.* at 114

		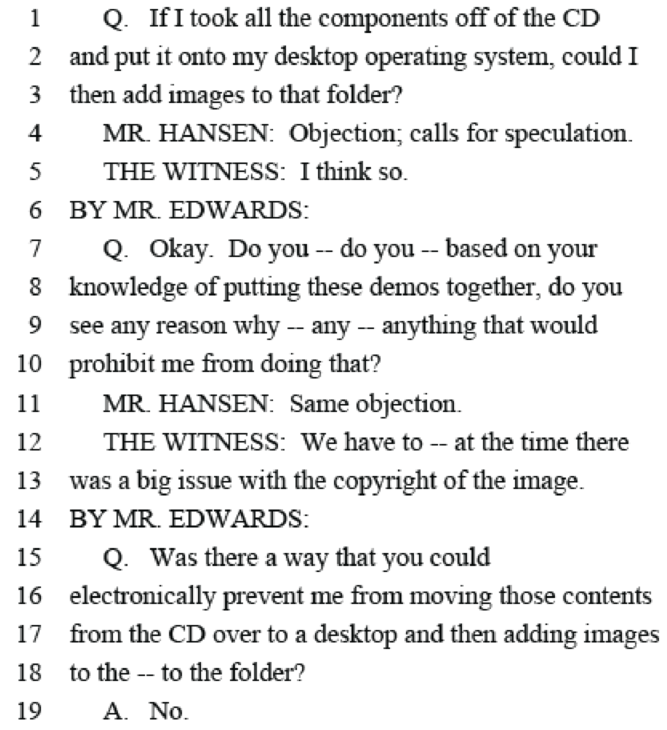 Q. If I took all the components off of the CD and put it onto my desktop operating system, could I then add images to that folder? MR. HANSEN: Objection; calls for speculation. THE WITNESS: I think so. BY MR. EDWARDS: Q. Okay. Do you -- do you -- based on your knowledge of putting these demos together, do you see any reason why -- any -- anything that would prohibit me from doing that? MR. HANSEN: Same objection. THE WITNESS: We have to -- at the time there was a big issue with the copyright of the image. BY MR. EDWARDS: Q. Was there a way that you could electronically prevent me from moving those contents from the CD over to a desktop and then adding images to the -- to the folder? A. No. *Id.* at 202

			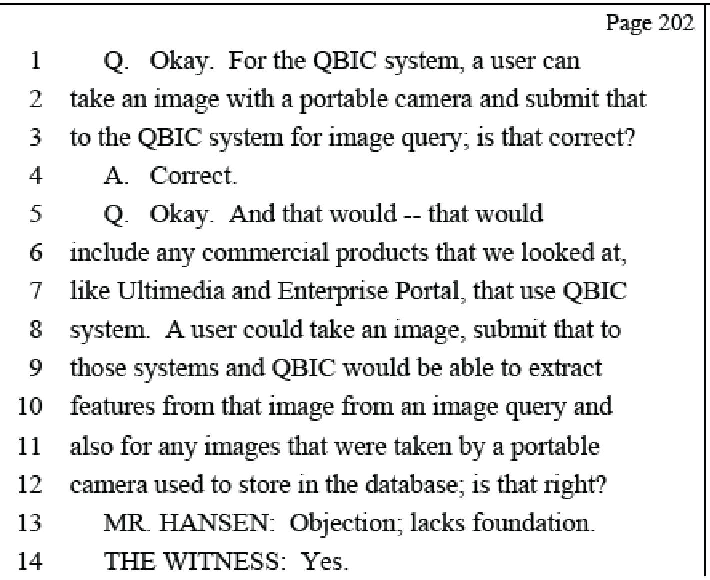 Page 202 Q. Okay. For the QBIC system, a user can take an image with a portable camera and submit that to the QBIC system for image query; is that correct? A. Correct. Q. Okay. And that would -- that would include any commercial products that we looked at, like Ultimedia and Enterprise Portal, that use QBIC system. A user could take an image, submit that to those systems and QBIC would be able to extract features from that image from an image query and also for any images that were taken by a portable camera used to store in the database; is that right? MR. HANSEN: Objection; lacks foundation. THE WITNESS: Yes. *Id.* at 300 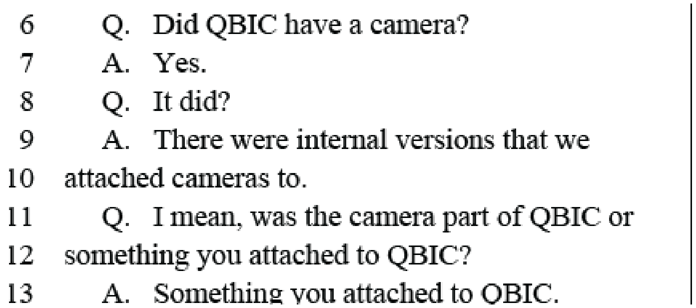 Q. Did QBIC have a camera? A. Yes. Q. It did? A. There were internal versions that we attached cameras to. Q. I mean, was the camera part of QBIC or something you attached to QBIC? A. Something you attached to QBIC. *Id.* at 312

		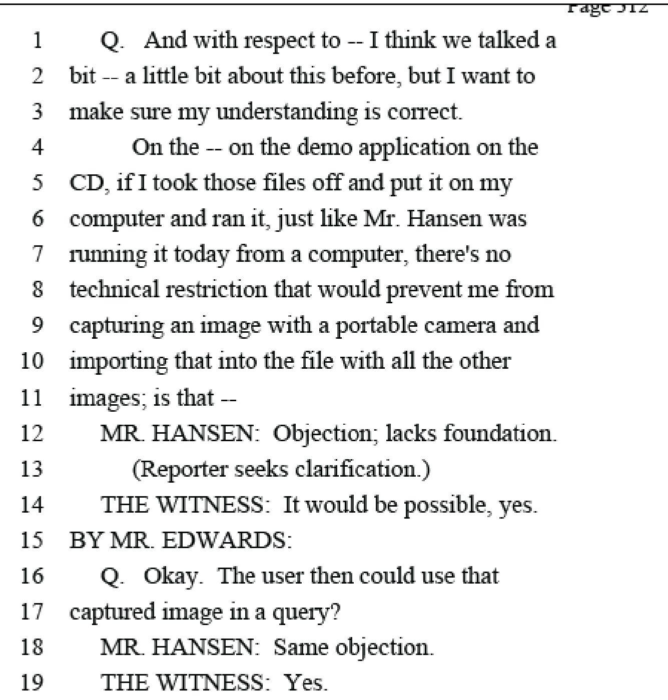 Page 312 Q. And with respect to -- I think we talked a bit -- a little bit about this before, but I want to make sure my understanding is correct. On the -- on the demo application on the CD, if I took those files off and put it on my computer and ran it, just like Mr. Hansen was running it today from a computer, there's no technical restriction that would prevent me from capturing an image with a portable camera and importing that into the file with all the other images; is that -- MR. HANSEN: Objection; lacks foundation. (Reporter seeks clarification.) THE WITNESS: It would be possible, yes. BY MR. EDWARDS: Q. Okay. The user then could use that captured image in a query? MR. HANSEN: Same objection. THE WITNESS: Yes. To the extent it is found that QBIC does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that QBIC does not anticipate this claim, QBIC renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine QBIC with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results. For example, the '252 patent acknowledges that acquiring information about objects through symbols, such as barcodes or some other means to encode detectable information in

		the image or on the object, were traditional methods known in the art and not part the present invention. '252 Patent, 3:7-16. Indeed, the '252 patent acknowledges that algorithms, software, and hardware components were commercially available at the time of the invention that could detect and decode symbols, such as barcodes or text, in an input image. *Id*. at 11:56-61. This limitation was thus known in the art by patentee's own admission, and by asserting this claim, Plaintiff is attempting to reclaim disavowed claim scope and claim subject matter which patentee has acknowledged is not part of the invention.

3. Claim 27

Claim	'252 Claim Recitation	Exemplary Citations in Reference
27	The method of claim 26, wherein image processing further comprises compressing the image to form the data relating to the image.	QBIC discloses the method of claim 26, wherein image processing further comprises compressing the image to form the data relating to the image. Query by Image and Video Content: The QBIC System (IBM 0000893–902): "QBIC has two main components: database population (the process of creating an image database) and database query. During the population, images and videos are processed to extract features describing their content-colors, textures, shapes, and camera and object motion-and the features are stored in a database." *Id*. at p. 3. Image, Audio, and Video Extenders Administration and Programming (IBM 000001–617): "In addition, you can specify conversion options that identify additional changes, such as rotation or compression, that you want to apply to the stored image." *Id*. at p. 95. QBIC Demo (IBM000747)  To specify your color histogram query, "paint" a graphical representation of your pattern of colors and approximate percentages that you want. First select a *percentage* by clicking in the painting area to block off a rectangle of the desired size (percentage of the full area). All new regions are painted white, initially. Then select a color from the

		QBIC Demo (IBM000747) 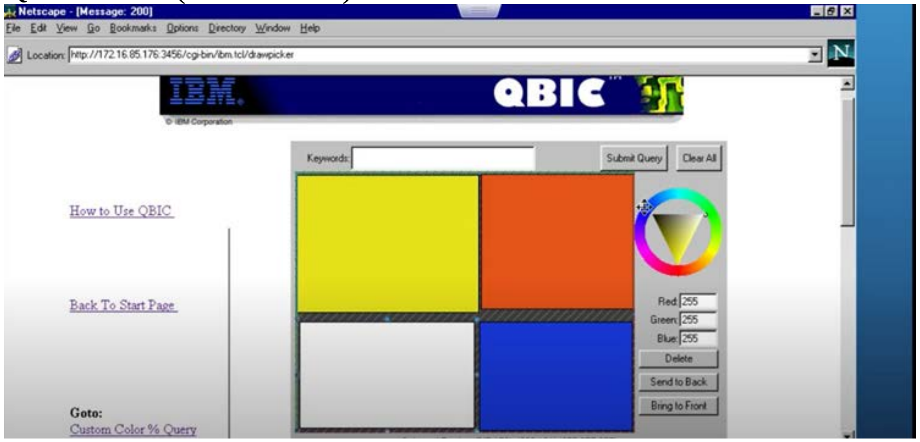 Location: http://172.16.85.176:3456/cgi-bin/ibm.tcl/drawpicker Red: 255 Green: 255 Blue: 255  Location: http://172.16.85.176:3456/cgi-bin/ibm.tcl/search?fname=HybridColor&france10%2Ejpg.x=0 Query was: Example: =france10.jpg Query Type: QbHybridColorFeatureClass Flickner Depo at 49-50

		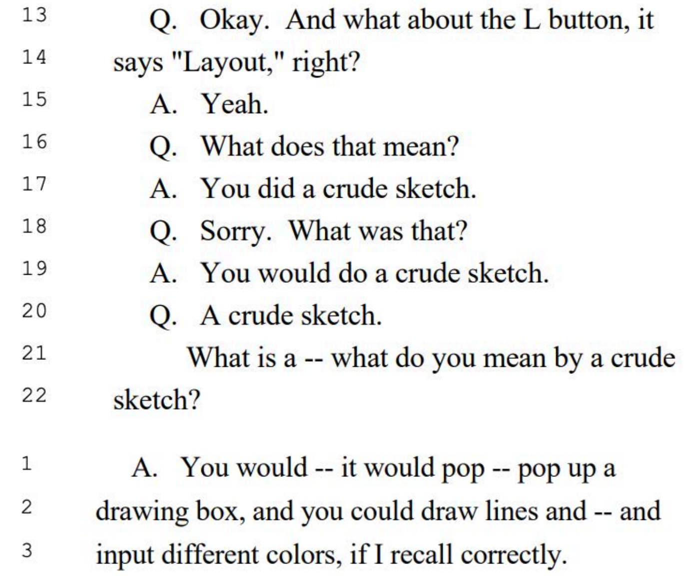 Q. Okay. And what about the L button, it says "Layout," right? A. Yeah. Q. What does that mean? A. You did a crude sketch. Q. Sorry. What was that? A. You would do a crude sketch. Q. A crude sketch. What is a -- what do you mean by a crude sketch? A. You would -- it would pop -- pop up a drawing box, and you could draw lines and -- and input different colors, if I recall correctly. Flickner Depo at 214-215 Flickner 8/29/23 Depo at 79–80

		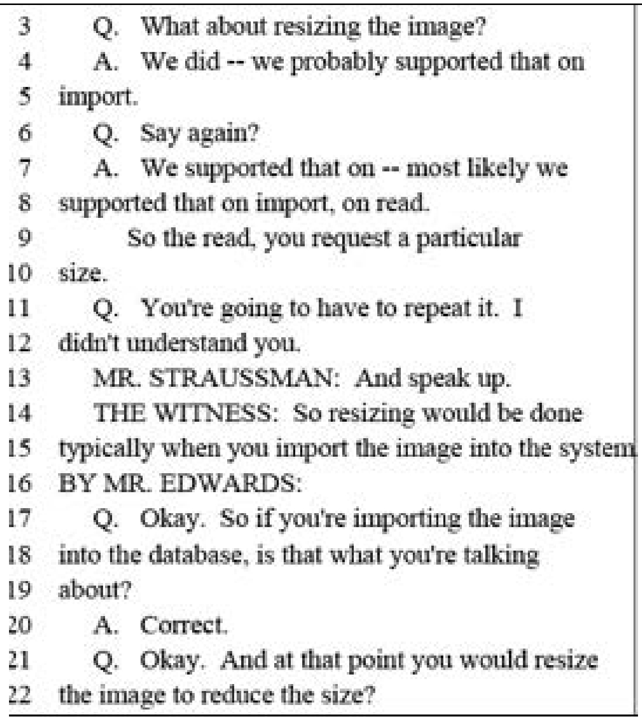 Q. What about resizing the image? A. We did -- we probably supported that on import. Q. Say again? A. We supported that on -- most likely we supported that on import, on read. So the read, you request a particular size. Q. You're going to have to repeat it. I didn't understand you. MR. STRAUSSMAN: And speak up. THE WITNESS: So resizing would be done typically when you import the image into the system BY MR. EDWARDS: Q. Okay. So if you're importing the image into the database, is that what you're talking about? A. Correct. Q. Okay. And at that point you would resize the image to reduce the size?	

		How would you resize the image on importation? A. You would request a particular size of image and then use that to populate the database. Q. So if the source images that you were acquiring were say too big for purposes of QBIC, on importation you could resize those to get them to the size that are -- meet the -- meet the requirements of the QBIC system? A. Yes. Q. Was that done as a matter of course in the QBIC system? MR. HANSEN: Objection; vague and ambiguous. THE WITNESS: As I recall, yes. BY MR. EDWARDS: Q. Was it an automatic process? A. As I recall, yes. Q. So if a user were to import source images into the database, when the user did that, the QBIC system would automatically resize it to fit the requirement of the QBIC system? A. Make sure it's in the range that the QBIC system could handle. Q. That's a "yes"? A. That's a "yes." *Id.* at 141 – 142

		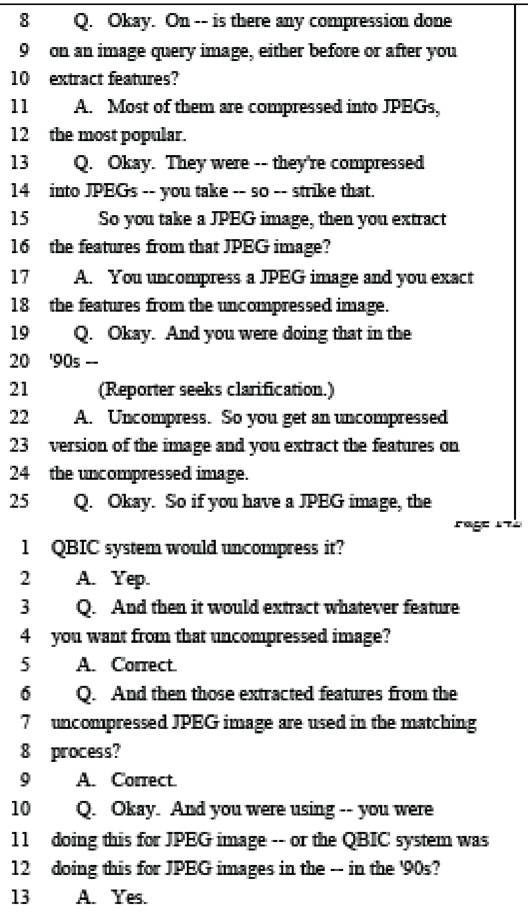 Q. Okay. On -- is there any compression done on an image query image, either before or after you extract features? A. Most of them are compressed into JPEGs, the most popular. Q. Okay. They were -- they're compressed into JPEGs -- you take -- so -- strike that. So you take a JPEG image, then you extract the features from that JPEG image? A. You uncompress a JPEG image and you exact the features from the uncompressed image. Q. Okay. And you were doing that in the '90s -- (Reporter seeks clarification.) A. Uncompress. So you get an uncompressed version of the image and you extract the features on the uncompressed image. Q. Okay. So if you have a JPEG image, the QBIC system would uncompress it? A. Yep. Q. And then it would extract whatever feature you want from that uncompressed image? A. Correct. Q. And then those extracted features from the uncompressed JPEG image are used in the matching process? A. Correct. Q. Okay. And you were using -- you were doing this for JPEG image -- or the QBIC system was doing this for JPEG images in the -- in the '90s? A. Yes. *Id.* at 201

		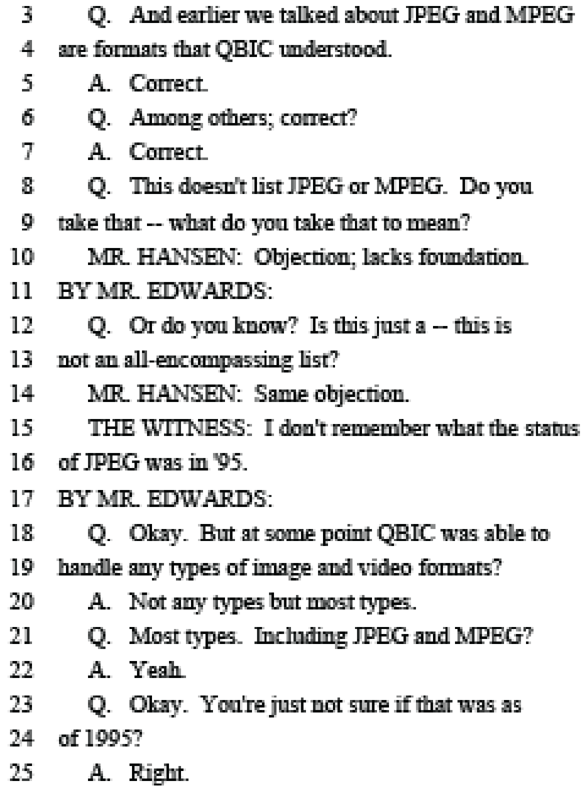 Q. And earlier we talked about JPEG and MPEG are formats that QBIC understood. A. Correct. Q. Among others; correct? A. Correct. Q. This doesn't list JPEG or MPEG. Do you take that -- what do you take that to mean? MR. HANSEN: Objection; lacks foundation. BY MR. EDWARDS: Q. Or do you know? Is this just a -- this is not an all-encompassing list? MR. HANSEN: Same objection. THE WITNESS: I don't remember what the status of JPEG was in '95. BY MR. EDWARDS: Q. Okay. But at some point QBIC was able to handle any types of image and video formats? A. Not any types but most types. Q. Most types. Including JPEG and MPEG? A. Yeah. Q. Okay. You're just not sure if that was as of 1995? A. Right. *Id.* at 300 – 301 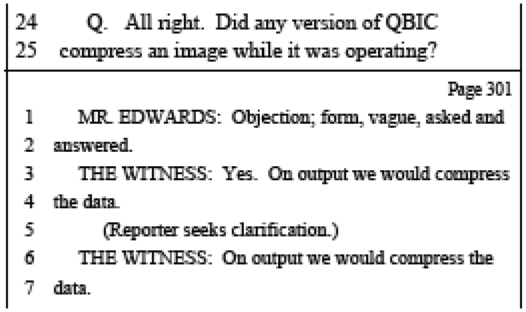 Q. All right. Did any version of QBIC compress an image while it was operating? Page 301 MR. EDWARDS: Objection; form, vague, asked and answered. THE WITNESS: Yes. On output we would compress the data. (Reporter seeks clarification.) THE WITNESS: On output we would compress the data.

Id. at 332

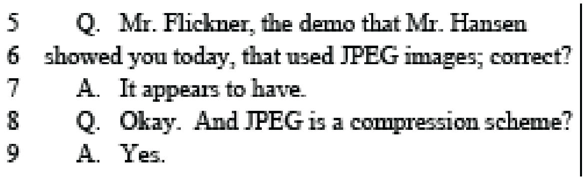

Q. Mr. Flickner, the demo that Mr. Hansen showed you today, that used JPEG images; correct?
A. It appears to have.
Q. Okay. And JPEG is a compression scheme?
A. Yes.

To the extent it is found that QBIC does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that QBIC does not anticipate this claim, QBIC renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine QBIC with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

4. Claim 29

Claim	'252 Claim Recitation	Exemplary Citations in Reference
29	The method of claim 18, wherein image processing further includes converting at least a portion of the data relating to the image into a grayscale image.	QBIC discloses the method of claim 18, wherein image processing further includes converting at least a portion of the data relating to the image into a grayscale image. The QBIC Project: Querying Images by Content Using Color, Texture, and Shape (IBM 0002390–404): "Texture features: Our texture features are based on modified versions of the coarseness, contrast, and directionality features proposed in [10]. The coarseness feature helps measure the scale of the texture (pebbles vs. boulders), and is calculated using moving windows of different sizes. The contrast feature describes the vividness of the pattern, and is a function of the variance of the gray-level histogram." *Id.* at p. 3. Query by Image and Video Content: The QBIC System (IBM 0000893–902): "QBIC has two main components: database population (the process of creating an image database) and database query. During the population, images and videos are processed to extract features describing their content-colors, textures, shapes, and camera and object motion-and the features are stored in a database." *Id.* at p. 3. Ultimedia Manager: Query by Image Content and its Applications (IBM 0002432–237): "Texture, measured as coarseness (pebbles vs. sand), contrast (presence of bright and dark image areas) and directionality (ordered vs. unordered edges in the image area), all computed from the luminance or gray level image; useful for searching for 'flat' or non-busy image regions, or image areas with many edges." *Id.* at pp. 425–26.

		Flickner Depo at 176-177 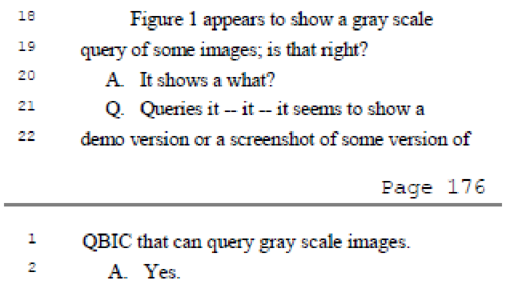 Figure 1 appears to show a gray scale query of some images; is that right? A. It shows a what? Q. Queries it -- it -- it seems to show a demo version or a screenshot of some version of Page 176 QBIC that can query gray scale images. A. Yes. Flickner Depo at 54-55 Flickner 8/29/23 Depo at 82–83

		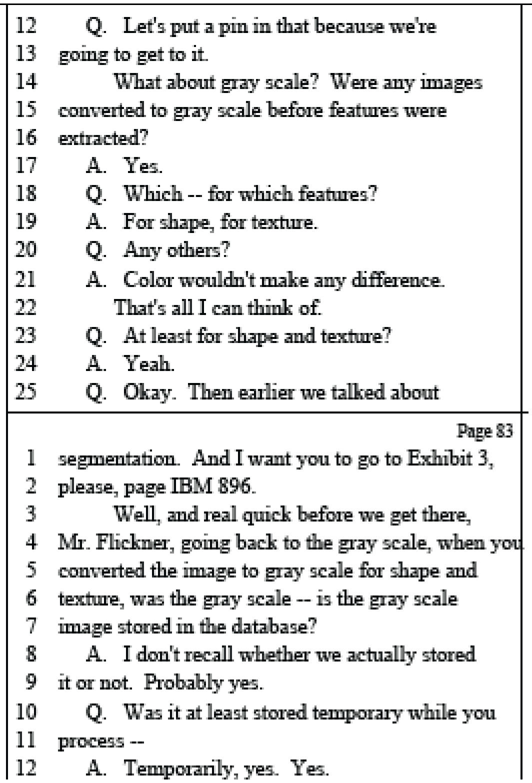 Q. Let's put a pin in that because we're going to get to it. What about gray scale? Were any images converted to gray scale before features were extracted? A. Yes. Q. Which -- for which features? A. For shape, for texture. Q. Any others? A. Color wouldn't make any difference. That's all I can think of. Q. At least for shape and texture? A. Yeah. Q. Okay. Then earlier we talked about Page 83 segmentation. And I want you to go to Exhibit 3, please, page IBM 896. Well, and real quick before we get there, Mr. Flickner, going back to the gray scale, when you converted the image to gray scale for shape and texture, was the gray scale -- is the gray scale image stored in the database? A. I don't recall whether we actually stored it or not. Probably yes. Q. Was it at least stored temporary while you process -- A. Temporarily, yes. Yes. *Id.* at 92–93

		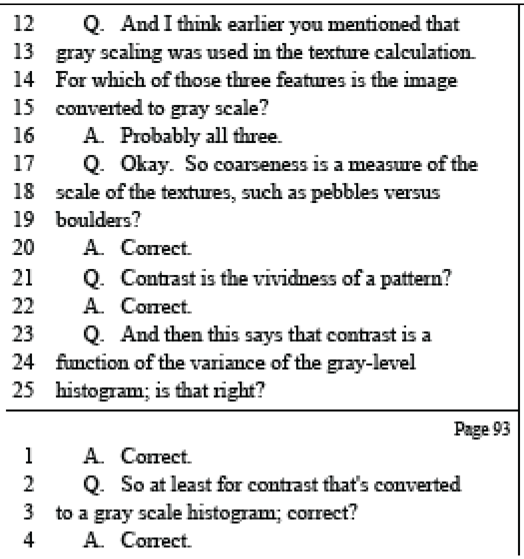 Q. And I think earlier you mentioned that gray scaling was used in the texture calculation. For which of those three features is the image converted to gray scale? A. Probably all three. Q. Okay. So coarseness is a measure of the scale of the textures, such as pebbles versus boulders? A. Correct. Q. Contrast is the vividness of a pattern? A. Correct. Q. And then this says that contrast is a function of the variance of the gray-level histogram; is that right? Page 93 A. Correct. Q. So at least for contrast that's converted to a gray scale histogram; correct? A. Correct. *Id.* at 280 – 281 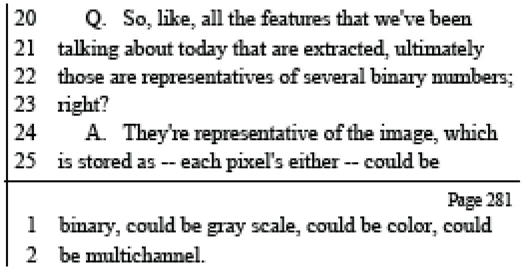 Q. So, like, all the features that we've been talking about today that are extracted, ultimately those are representatives of several binary numbers; right? A. They're representative of the image, which is stored as -- each pixel's either -- could be Page 281 binary, could be gray scale, could be color, could be multichannel. *Id.* at 301 – 302	

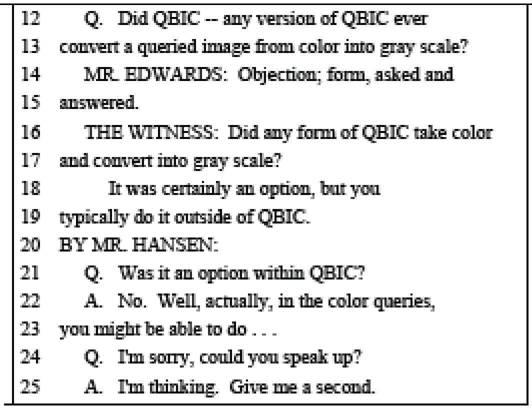

Q. Did QBIC -- any version of QBIC ever convert a queried image from color into gray scale?

MR. EDWARDS: Objection; form, asked and answered.

THE WITNESS: Did any form of QBIC take color and convert into gray scale?

It was certainly an option, but you typically do it outside of QBIC.

BY MR. HANSEN:

Q. Was it an option within QBIC?

A. No. Well, actually, in the color queries, you might be able to do . . .

Q. I'm sorry, could you speak up?

A. I'm thinking. Give me a second.

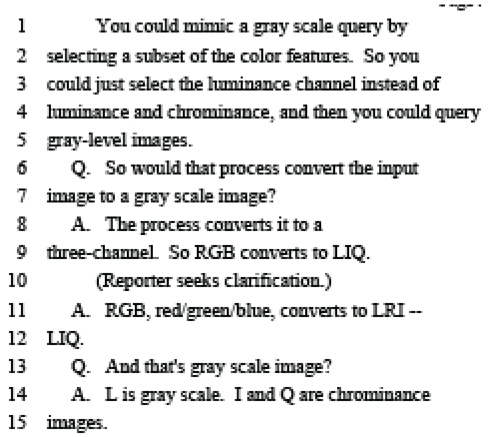

You could mimic a gray scale query by selecting a subset of the color features. So you could just select the luminance channel instead of luminance and chrominance, and then you could query gray-level images.

Q. So would that process convert the input image to a gray scale image?

A. The process converts it to a three-channel. So RGB converts to LIQ.

(Reporter seeks clarification.)

A. RGB, red/green/blue, converts to LRI -- LIQ.

Q. And that's gray scale image?

A. L is gray scale. I and Q are chrominance images.

Id. at 81-86

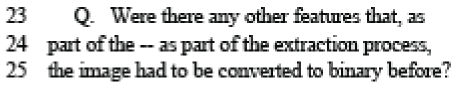

Q. Were there any other features that, as part of the -- as part of the extraction process, the image had to be converted to binary before?

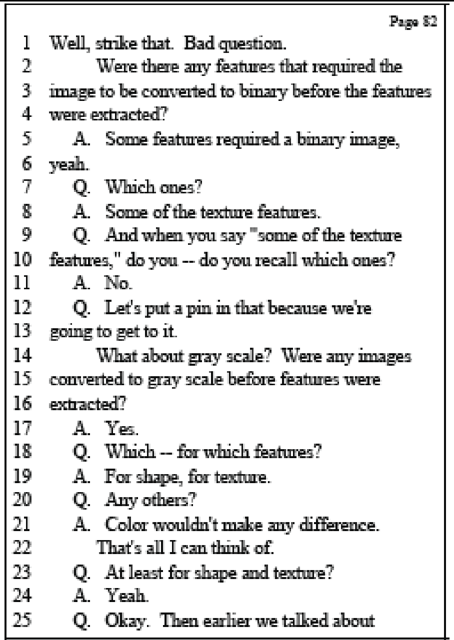
Page 82

Well, strike that. Bad question.
Were there any features that required the image to be converted to binary before the features were extracted?
A. Some features required a binary image, yeah.
Q. Which ones?
A. Some of the texture features.
Q. And when you say "some of the texture features," do you -- do you recall which ones?
A. No.
Q. Let's put a pin in that because we're going to get to it.
What about gray scale? Were any images converted to gray scale before features were extracted?
A. Yes.
Q. Which -- for which features?
A. For shape, for texture.
Q. Any others?
A. Color wouldn't make any difference.
That's all I can think of.
Q. At least for shape and texture?
A. Yeah.
Q. Okay. Then earlier we talked about

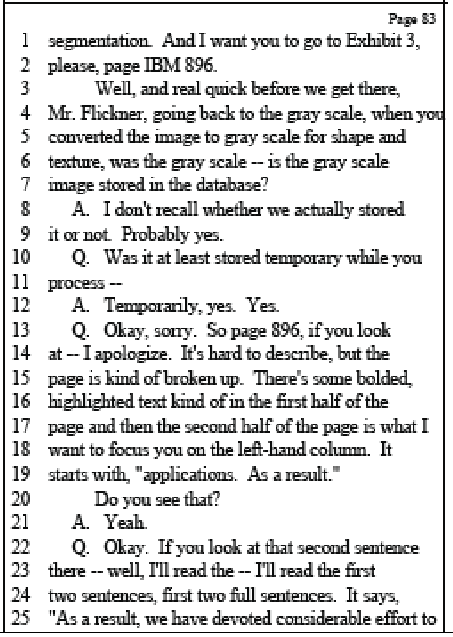
Page 83

segmentation. And I want you to go to Exhibit 3, please, page IBM 896.
Well, and real quick before we get there, Mr. Flickner, going back to the gray scale, when you converted the image to gray scale for shape and texture, was the gray scale -- is the gray scale image stored in the database?
A. I don't recall whether we actually stored it or not. Probably yes.
Q. Was it at least stored temporary while you process --
A. Temporarily, yes. Yes.
Q. Okay, sorry. So page 896, if you look at -- I apologize. It's hard to describe, but the page is kind of broken up. There's some bolded, highlighted text kind of in the first half of the page and then the second half of the page is what I want to focus you on the left-hand column. It starts with, "applications. As a result."
Do you see that?
A. Yeah.
Q. Okay. If you look at that second sentence there -- well, I'll read the -- I'll read the first two sentences, first two full sentences. It says, "As a result, we have devoted considerable effort to

		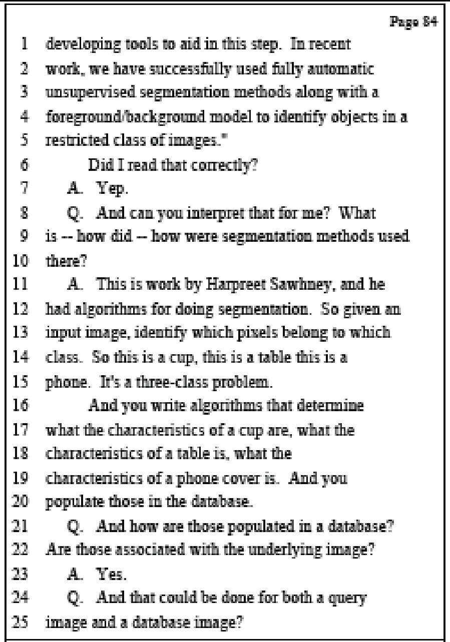 Page 84 developing tools to aid in this step. In recent work, we have successfully used fully automatic unsupervised segmentation methods along with a foreground/background model to identify objects in a restricted class of images." Did I read that correctly? A. Yep. Q. And can you interpret that for me? What is -- how did -- how were segmentation methods used there? A. This is work by Harpreet Sawhney, and he had algorithms for doing segmentation. So given an input image, identify which pixels belong to which class. So this is a cup, this is a table this is a phone. It's a three-class problem. And you write algorithms that determine what the characteristics of a cup are, what the characteristics of a table is, what the characteristics of a phone cover is. And you populate those in the database. Q. And how are those populated in a database? Are those associated with the underlying image? A. Yes. Q. And that could be done for both a query image and a database image? 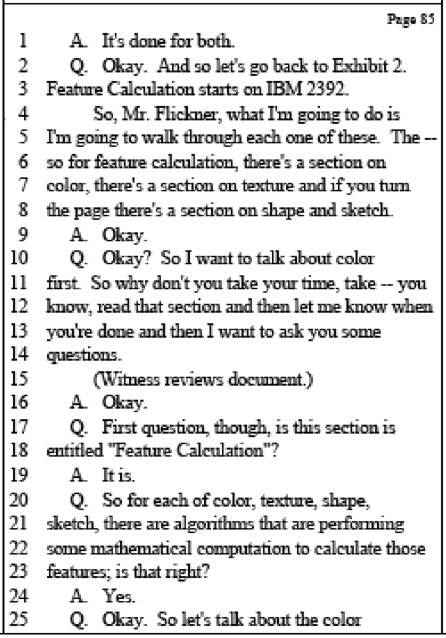 Page 85 A. It's done for both. Q. Okay. And so let's go back to Exhibit 2. Feature Calculation starts on IBM 2392. So, Mr. Flickner, what I'm going to do is I'm going to walk through each one of these. The -- so for feature calculation, there's a section on color, there's a section on texture and if you turn the page there's a section on shape and sketch. A. Okay. Q. Okay? So I want to talk about color first. So why don't you take your time, take -- you know, read that section and then let me know when you're done and then I want to ask you some questions. (Witness reviews document.) A. Okay. Q. First question, though, is this section is entitled "Feature Calculation"? A. It is. Q. So for each of color, texture, shape, sketch, there are algorithms that are performing some mathematical computation to calculate those features; is that right? A. Yes. Q. Okay. So let's talk about the color

		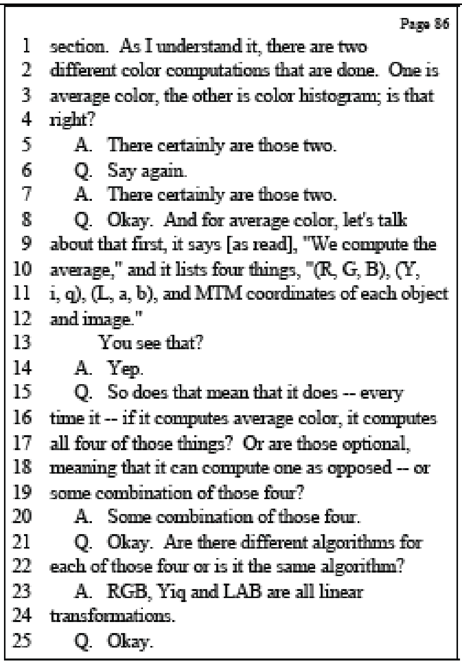 Page 86 section. As I understand it, there are two different color computations that are done. One is average color, the other is color histogram; is that right? A. There certainly are those two. Q. Say again. A. There certainly are those two. Q. Okay. And for average color, let's talk about that first, it says [as read], "We compute the average," and it lists four things, "(R, G, B), (Y, i, q), (L, a, b), and MTM coordinates of each object and image." You see that? A. Yep. Q. So does that mean that it does -- every time it -- if it computes average color, it computes all four of those things? Or are those optional, meaning that it can compute one as opposed -- or some combination of those four? A. Some combination of those four. Q. Okay. Are there different algorithms for each of those four or is it the same algorithm? A. RGB, Yiq and LAB are all linear transformations. Q. Okay. To the extent it is found that QBIC does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that QBIC does not anticipate this claim, QBIC renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine QBIC with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

5. Claim 31

Claim	’252 Claim Recitation	Exemplary Citations in Reference
[31]	The method of claim 18, further comprising the distal server using database to identify the pertinent information based on image characteristics of the image derived by the image processing platform based on the data relating to the image.	QBIC discloses the method of claim 18, further comprising the distal server using database to identify the pertinent information based on image characteristics of the image derived by the image processing platform based on the data relating to the image. The QBIC Project: Querying Images by Content Using Color, Texture, and Shape (IBM 0002390–404) at 2390: 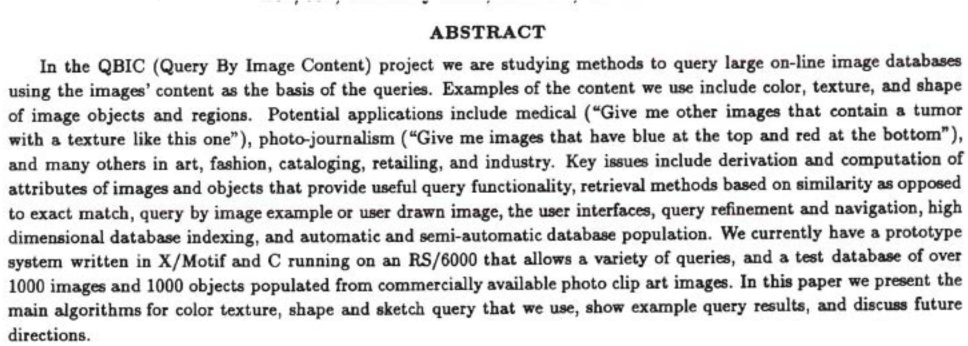 ABSTRACT In the QBIC (Query By Image Content) project we are studying methods to query large on-line image databases using the images' content as the basis of the queries. Examples of the content we use include color, texture, and shape of image objects and regions. Potential applications include medical (“Give me other images that contain a tumor with a texture like this one”), photo-journalism (“Give me images that have blue at the top and red at the bottom”), and many others in art, fashion, cataloging, retailing, and industry. Key issues include derivation and computation of attributes of images and objects that provide useful query functionality, retrieval methods based on similarity as opposed to exact match, query by image example or user drawn image, the user interfaces, query refinement and navigation, high dimensional database indexing, and automatic and semi-automatic database population. We currently have a prototype system written in X/Motif and C running on an RS/6000 that allows a variety of queries, and a test database of over 1000 images and 1000 objects populated from commercially available photo clip art images. In this paper we present the main algorithms for color texture, shape and sketch query that we use, show example query results, and discuss future directions. QBIC Demo (IBM000747):

QBIC(TM) -- IBM's Query By Image Content

Goto:

How to Use QBIC

IBM Image Search Demo

Stop QBIC Http Server

QBIC allows you to make queries of image databases large and small based on visual characteristics -- properties such as color percentages, color layout and textures occurring in the images. The queries match image colors, textures and their positions without having to describe them with in words. QBIC's content-based queries can also be combined with keywords to provide a powerful multimedia search and retrieval tool. As image collections grow larger and become more prolific, QBIC is the best tool to more efficiently search and navigate through them.

In this demo , we use QBIC to index about 470 color images based on their color and texture information. You can evaluate the QBIC Search capability by following the link located at left called IBM Image Search Demo.

For license and business related questins, please contact Ted Loewenberg by phone 408-927-1202, or by email

Flickner Depo at, e.g., 24, 41-42, 62-63 (user could take a picture or video and submit it to the QBIC system).

Flickner Depo at 18:

A. Well, we did stuff with the San Francisco Museum, the stamps database, trademark database. So we did some of the stuff with Davis --
Q. Okay.
A. -- UC Davis. I remember we talked to a couple people at Davis.
Q. So was this an active area of research in the 1990s?
A. Yeah.

Flickner Depo at 20:

		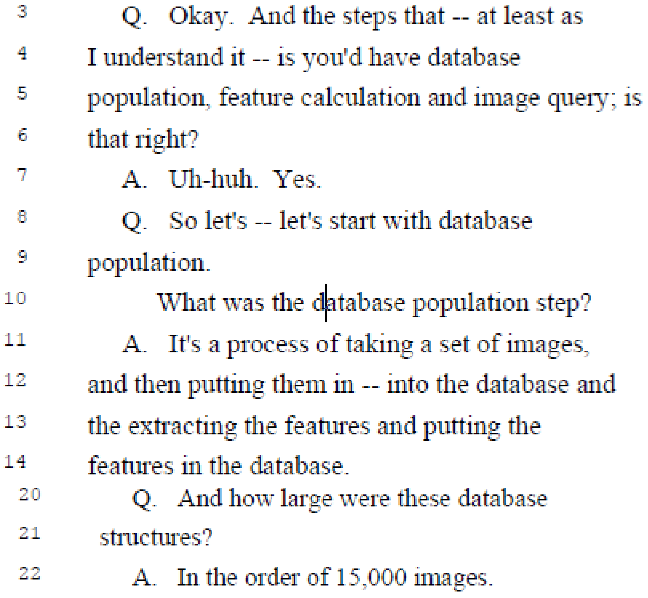 Q. Okay. And the steps that -- at least as I understand it -- is you'd have database population, feature calculation and image query; is that right? A. Uh-huh. Yes. Q. So let's -- let's start with database population. What was the database population step? A. It's a process of taking a set of images, and then putting them in -- into the database and the extracting the features and putting the features in the database. Q. And how large were these database structures? A. In the order of 15,000 images. Flickner Depo at 21:

		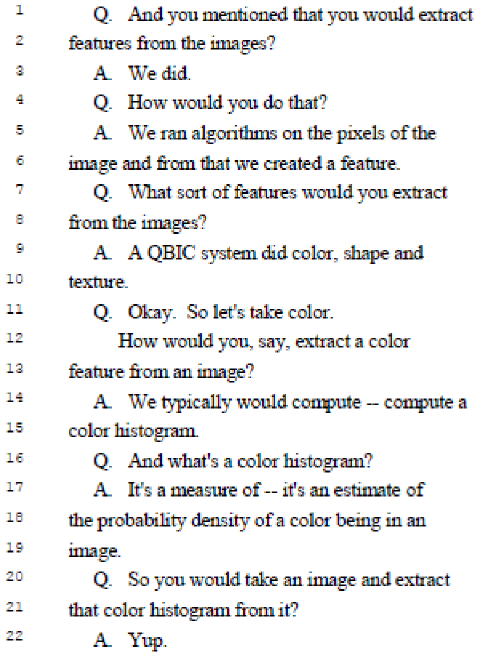 Q. And you mentioned that you would extract features from the images? A. We did. Q. How would you do that? A. We ran algorithms on the pixels of the image and from that we created a feature. Q. What sort of features would you extract from the images? A. A QBIC system did color, shape and texture. Q. Okay. So let's take color. How would you, say, extract a color feature from an image? A. We typically would compute -- compute a color histogram. Q. And what's a color histogram? A. It's a measure of -- it's an estimate of the probability density of a color being in an image. Q. So you would take an image and extract that color histogram from it? A. Yup. Flickner Depo at 22-25:

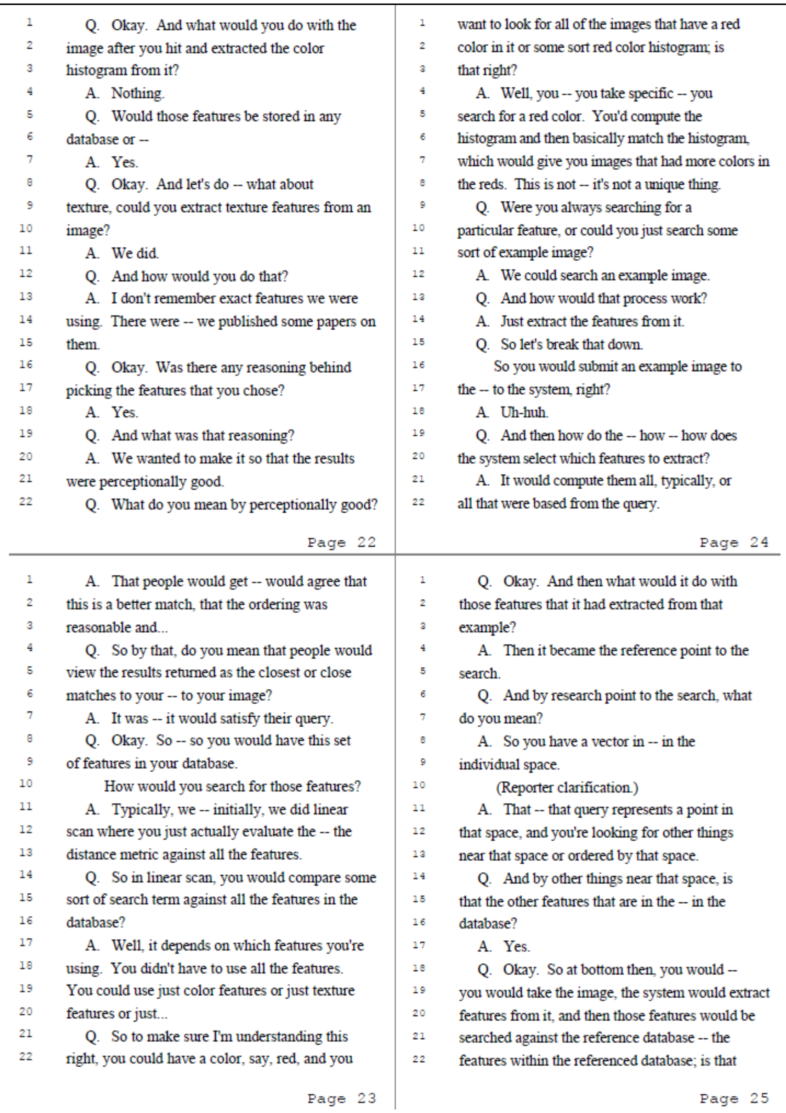

Q. Okay. And what would you do with the image after you hit and extracted the color histogram from it?
A. Nothing.
Q. Would those features be stored in any database or --
A. Yes.
Q. Okay. And let's do -- what about texture, could you extract texture features from an image?
A. We did.
Q. And how would you do that?
A. I don't remember exact features we were using. There were -- we published some papers on them.
Q. Okay. Was there any reasoning behind picking the features that you chose?
A. Yes.
Q. And what was that reasoning?
A. We wanted to make it so that the results were perceptionally good.
Q. What do you mean by perceptionally good?

Page 22

A. That people would get -- would agree that this is a better match, that the ordering was reasonable and...
Q. So by that, do you mean that people would view the results returned as the closest or close matches to your -- to your image?
A. It was -- it would satisfy their query.
Q. Okay. So -- so you would have this set of features in your database.
How would you search for those features?
A. Typically, we -- initially, we did linear scan where you just actually evaluate the -- the distance metric against all the features.
Q. So in linear scan, you would compare some sort of search term against all the features in the database?
A. Well, it depends on which features you're using. You didn't have to use all the features. You could use just color features or just texture features or just...
Q. So to make sure I'm understanding this right, you could have a color, say, red, and you

Page 23

want to look for all of the images that have a red color in it or some sort red color histogram; is that right?
A. Well, you -- you take specific -- you search for a red color. You'd compute the histogram and then basically match the histogram, which would give you images that had more colors in the reds. This is not -- it's not a unique thing.
Q. Were you always searching for a particular feature, or could you just search some sort of example image?
A. We could search an example image.
Q. And how would that process work?
A. Just extract the features from it.
Q. So let's break that down.
So you would submit an example image to the -- to the system, right?
A. Uh-huh.
Q. And then how do the -- how -- how does the system select which features to extract?
A. It would compute them all, typically, or all that were based from the query.

Page 24

Q. Okay. And then what would it do with those features that it had extracted from that example?
A. Then it became the reference point to the search.
Q. And by research point to the search, what do you mean?
A. So you have a vector in -- in the individual space.
(Reporter clarification.)
A. That -- that query represents a point in that space, and you're looking for other things near that space or ordered by that space.
Q. And by other things near that space, is that the other features that are in the -- in the database?
A. Yes.
Q. Okay. So at bottom then, you would -- you would take the image, the system would extract features from it, and then those features would be searched against the reference database -- the features within the referenced database; is that

Page 25

Flickner Depo at 25-26:

		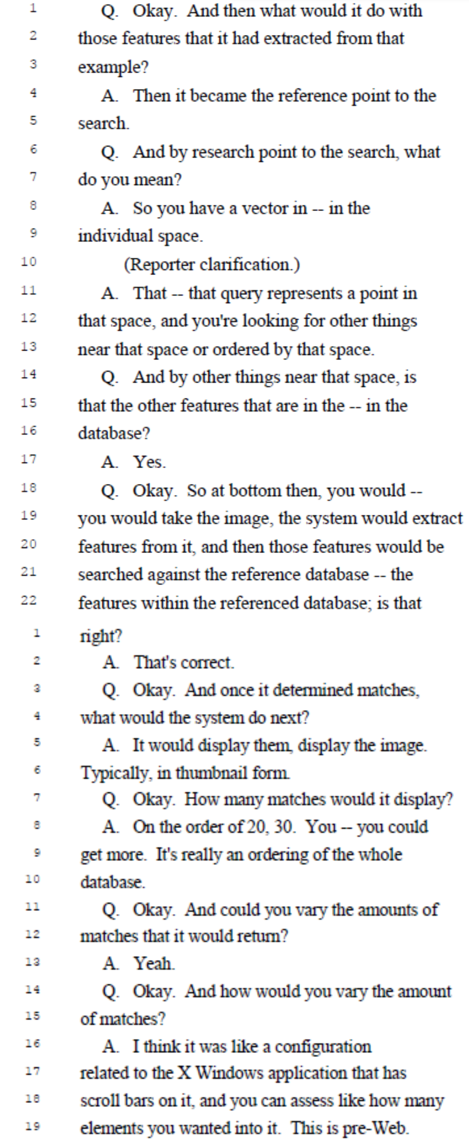 Q. Okay. And then what would it do with those features that it had extracted from that example? A. Then it became the reference point to the search. Q. And by research point to the search, what do you mean? A. So you have a vector in -- in the individual space. (Reporter clarification.) A. That -- that query represents a point in that space, and you're looking for other things near that space or ordered by that space. Q. And by other things near that space, is that the other features that are in the -- in the database? A. Yes. Q. Okay. So at bottom then, you would -- you would take the image, the system would extract features from it, and then those features would be searched against the reference database -- the features within the referenced database; is that right? A. That's correct. Q. Okay. And once it determined matches, what would the system do next? A. It would display them, display the image. Typically, in thumbnail form. Q. Okay. How many matches would it display? A. On the order of 20, 30. You -- you could get more. It's really an ordering of the whole database. Q. Okay. And could you vary the amounts of matches that it would return? A. Yeah. Q. Okay. And how would you vary the amount of matches? A. I think it was like a configuration related to the X Windows application that has scroll bars on it, and you can assess like how many elements you wanted into it. This is pre-Web.

		Flickner Depo at 36: 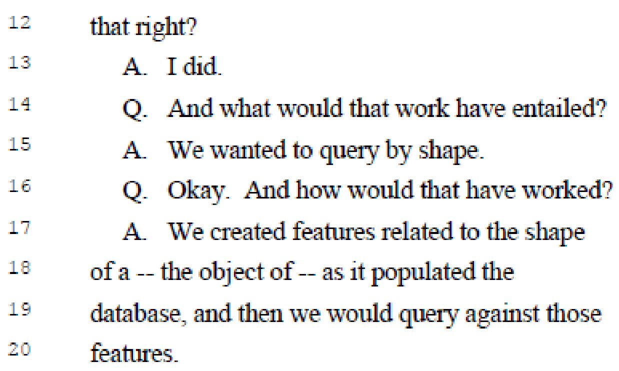 Q. So underlying this paper, you mentioned you did more work on the shape side of things; is that right? A. I did. Q. And what would that work have entailed? A. We wanted to query by shape. Q. Okay. And how would that have worked? A. We created features related to the shape of a -- the object of -- as it populated the database, and then we would query against those features. Flickner Depo 38-41:

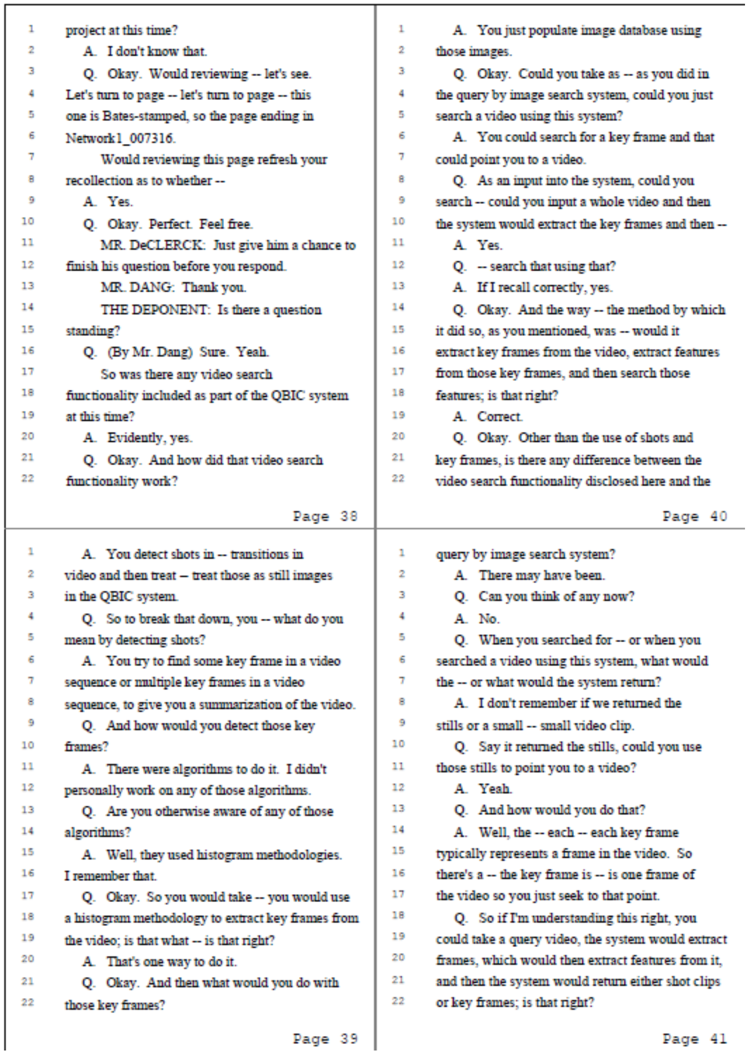

project at this time?
A. I don't know that.
Q. Okay. Would reviewing -- let's see. Let's turn to page -- let's turn to page -- this one is Bates-stamped, so the page ending in Network1_007316.
Would reviewing this page refresh your recollection as to whether --
A. Yes.
Q. Okay. Perfect. Feel free.
MR. DeCLERCK: Just give him a chance to finish his question before you respond.
MR. DANG: Thank you.
THE DEPONENT: Is there a question standing?
Q. (By Mr. Dang) Sure. Yeah.
So was there any video search functionality included as part of the QBIC system at this time?
A. Evidently, yes.
Q. Okay. And how did that video search functionality work?

Page 38

A. You detect shots in -- transitions in video and then treat -- treat those as still images in the QBIC system.
Q. So to break that down, you -- what do you mean by detecting shots?
A. You try to find some key frame in a video sequence or multiple key frames in a video sequence, to give you a summarization of the video.
Q. And how would you detect those key frames?
A. There were algorithms to do it. I didn't personally work on any of those algorithms.
Q. Are you otherwise aware of any of those algorithms?
A. Well, they used histogram methodologies. I remember that.
Q. Okay. So you would take -- you would use a histogram methodology to extract key frames from the video; is that what -- is that right?
A. That's one way to do it.
Q. Okay. And then what would you do with those key frames?

Page 39

A. You just populate image database using those images.
Q. Okay. Could you take as -- as you did in the query by image search system, could you just search a video using this system?
A. You could search for a key frame and that could point you to a video.
Q. As an input into the system, could you search -- could you input a whole video and then the system would extract the key frames and then --
A. Yes.
Q. -- search that using that?
A. If I recall correctly, yes.
Q. Okay. And the way -- the method by which it did so, as you mentioned, was -- would it extract key frames from the video, extract features from those key frames, and then search those features; is that right?
A. Correct.
Q. Okay. Other than the use of shots and key frames, is there any difference between the video search functionality disclosed here and the

Page 40

query by image search system?
A. There may have been.
Q. Can you think of any now?
A. No.
Q. When you searched for -- or when you searched a video using this system, what would the -- or what would the system return?
A. I don't remember if we returned the stills or a small -- small video clip.
Q. Say it returned the stills, could you use those stills to point you to a video?
A. Yeah.
Q. And how would you do that?
A. Well, the -- each -- each key frame typically represents a frame in the video. So there's a -- the key frame is -- is one frame of the video so you just seek to that point.
Q. So if I'm understanding this right, you could take a query video, the system would extract frames, which would then extract features from it, and then the system would return either shot clips or key frames; is that right?

Page 41

Flickner Depo at 41-42:

		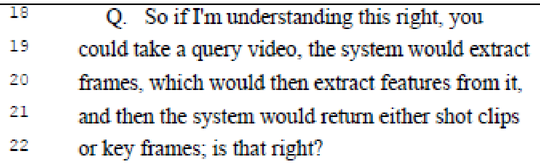 Q. So if I'm understanding this right, you could take a query video, the system would extract frames, which would then extract features from it, and then the system would return either shot clips or key frames; is that right? 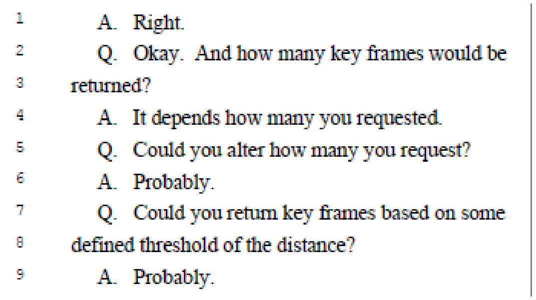 A. Right. Q. Okay. And how many key frames would be returned? A. It depends how many you requested. Q. Could you alter how many you request? A. Probably. Q. Could you return key frames based on some defined threshold of the distance? A. Probably. Flickner Depo at 48-49: 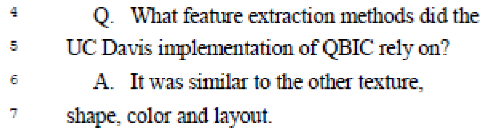 Q. What feature extraction methods did the UC Davis implementation of QBIC rely on? A. It was similar to the other texture, shape, color and layout. Flickner Depo at 50:

		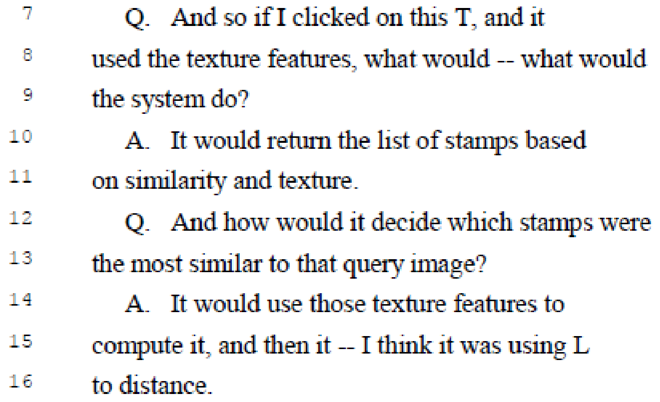 Q. And so if I clicked on this T, and it used the texture features, what would -- what would the system do? A. It would return the list of stamps based on similarity and texture. Q. And how would it decide which stamps were the most similar to that query image? A. It would use those texture features to compute it, and then it -- I think it was using L to distance. Flickner Depo at 54-55:

		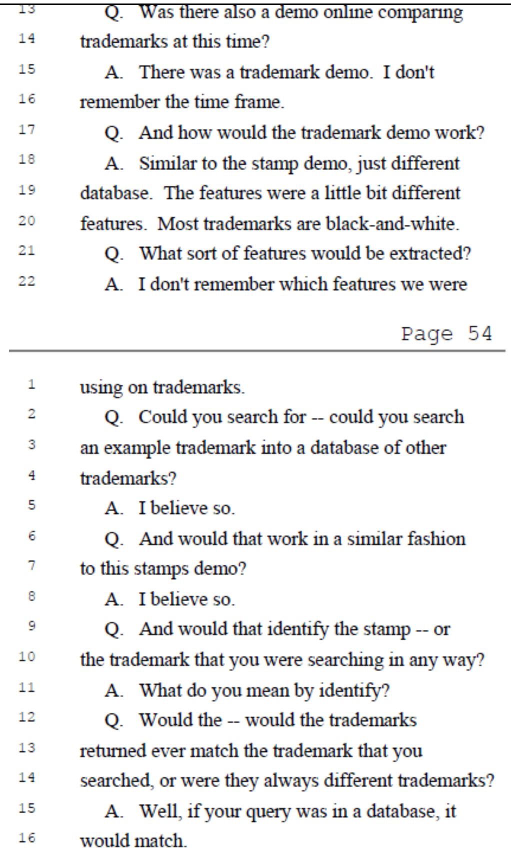 Q. Was there also a demo online comparing trademarks at this time? A. There was a trademark demo. I don't remember the time frame. Q. And how would the trademark demo work? A. Similar to the stamp demo, just different database. The features were a little bit different features. Most trademarks are black-and-white. Q. What sort of features would be extracted? A. I don't remember which features we were Page 54 using on trademarks. Q. Could you search for -- could you search an example trademark into a database of other trademarks? A. I believe so. Q. And would that work in a similar fashion to this stamps demo? A. I believe so. Q. And would that identify the stamp -- or the trademark that you were searching in any way? A. What do you mean by identify? Q. Would the -- would the trademarks returned ever match the trademark that you searched, or were they always different trademarks? A. Well, if your query was in a database, it would match. QBIC could integrate with a camera, such as a portable camera; a user could take an image using a portable camera and submit that image to the QBIC system for query. Flickner Depo at, 62-63 (user could take a picture or video and submit it to the QBIC system):

		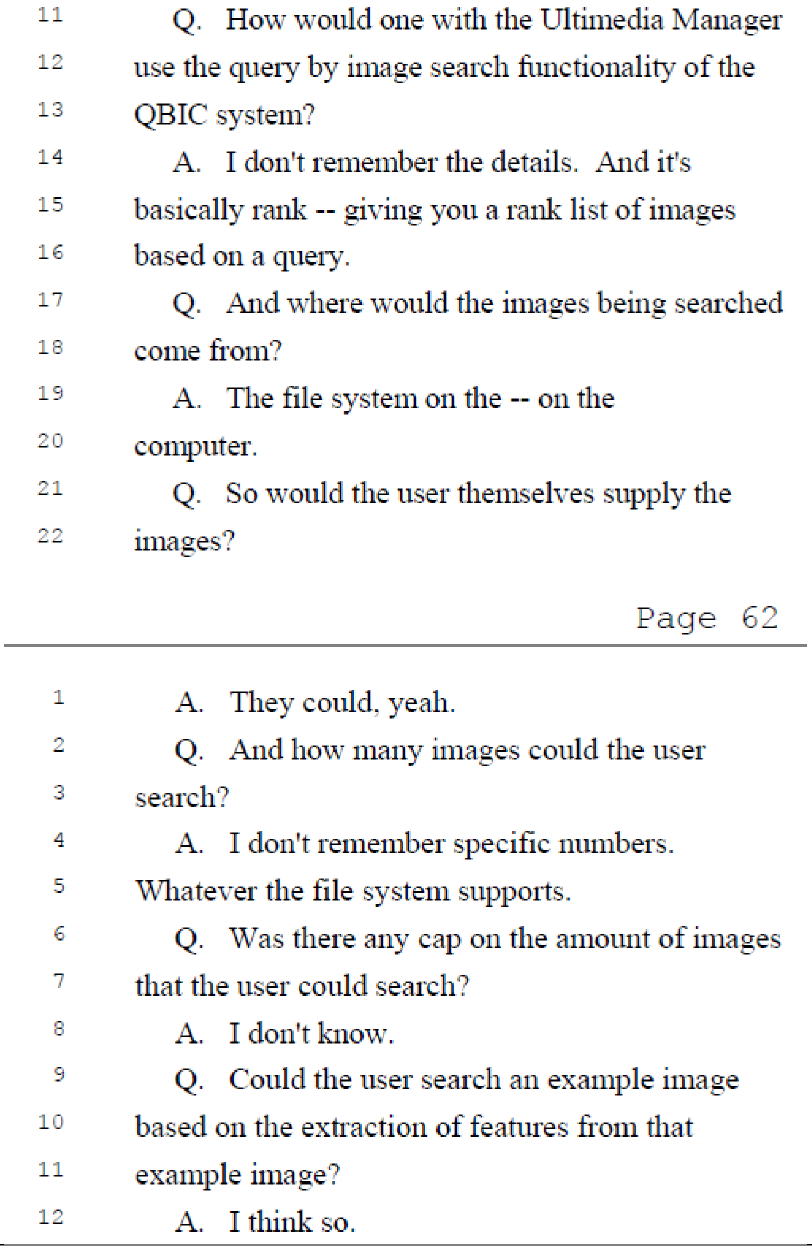 Q. How would one with the Ultimedia Manager use the query by image search functionality of the QBIC system? A. I don't remember the details. And it's basically rank -- giving you a rank list of images based on a query. Q. And where would the images being searched come from? A. The file system on the -- on the computer. Q. So would the user themselves supply the images? Page 62 A. They could, yeah. Q. And how many images could the user search? A. I don't remember specific numbers. Whatever the file system supports. Q. Was there any cap on the amount of images that the user could search? A. I don't know. Q. Could the user search an example image based on the extraction of features from that example image? A. I think so.

		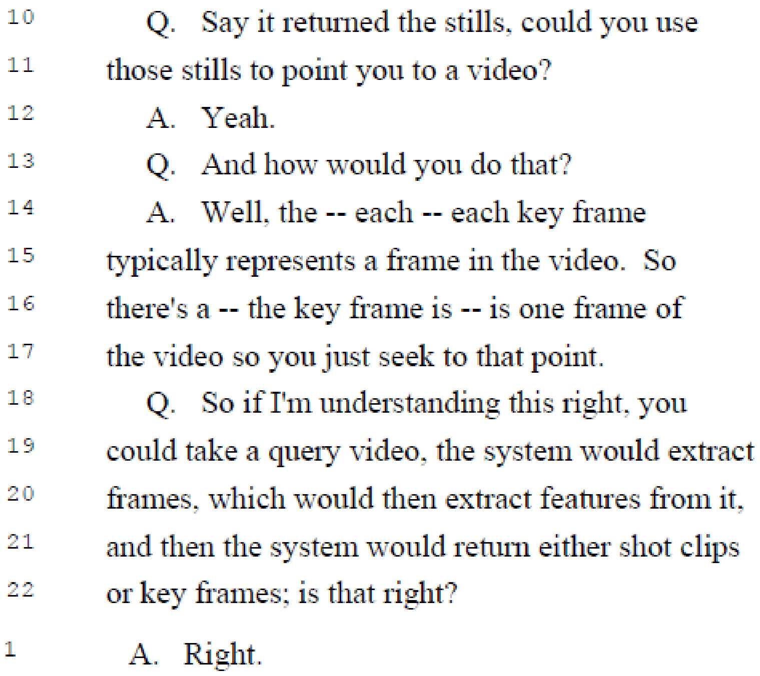 Q. Say it returned the stills, could you use those stills to point you to a video? A. Yeah. Q. And how would you do that? A. Well, the -- each -- each key frame typically represents a frame in the video. So there's a -- the key frame is -- is one frame of the video so you just seek to that point. Q. So if I'm understanding this right, you could take a query video, the system would extract frames, which would then extract features from it, and then the system would return either shot clips or key frames; is that right? A. Right. Flickner Depo at 63: 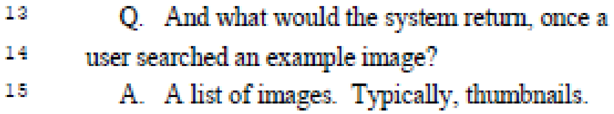 Q. And what would the system return, once a user searched an example image? A. A list of images. Typically, thumbnails. Flickner Depo at 72-73:

		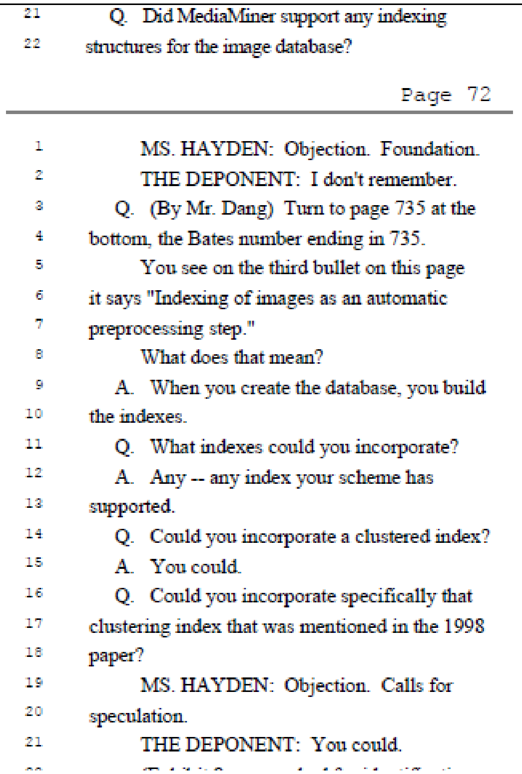 Q. Did MediaMiner support any indexing structures for the image database? Page 72 MS. HAYDEN: Objection. Foundation. THE DEPONENT: I don't remember. Q. (By Mr. Dang) Turn to page 735 at the bottom, the Bates number ending in 735. You see on the third bullet on this page it says "Indexing of images as an automatic preprocessing step." What does that mean? A. When you create the database, you build the indexes. Q. What indexes could you incorporate? A. Any -- any index your scheme has supported. Q. Could you incorporate a clustered index? A. You could. Q. Could you incorporate specifically that clustering index that was mentioned in the 1998 paper? MS. HAYDEN: Objection. Calls for speculation. THE DEPONENT: You could. Flickner Depo at 75-76:

		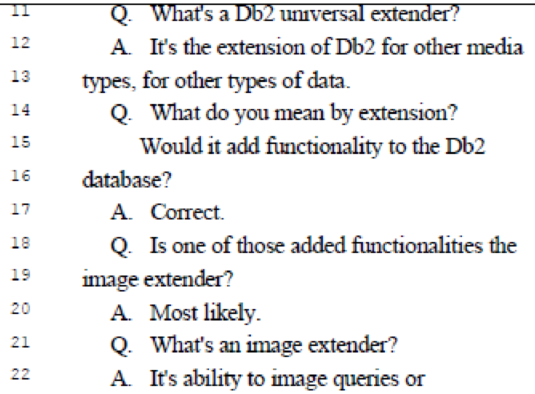 Q. What's a Db2 universal extender? A. It's the extension of Db2 for other media types, for other types of data. Q. What do you mean by extension? Would it add functionality to the Db2 database? A. Correct. Q. Is one of those added functionalities the image extender? A. Most likely. Q. What's an image extender? A. It's ability to image queries or 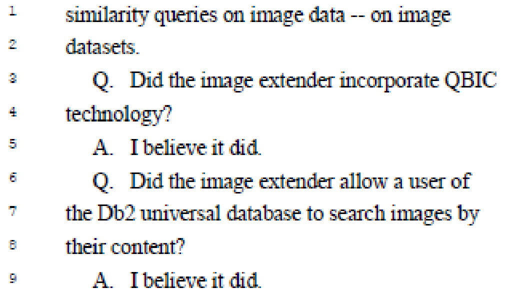 similarity queries on image data -- on image datasets. Q. Did the image extender incorporate QBIC technology? A. I believe it did. Q. Did the image extender allow a user of the Db2 universal database to search images by their content? A. I believe it did. Flickner Depo at 78-79:

		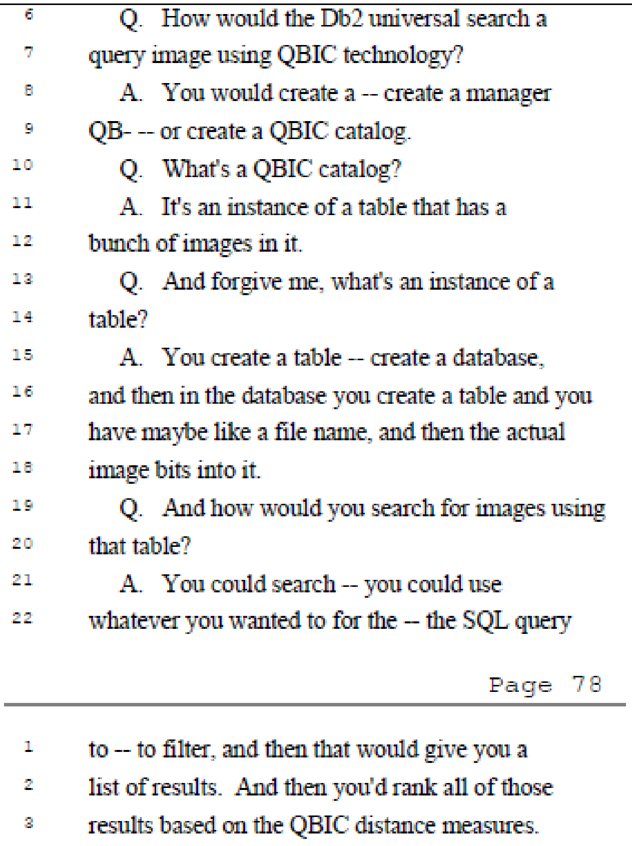 Q. How would the Db2 universal search a query image using QBIC technology? A. You would create a -- create a manager QB- -- or create a QBIC catalog. Q. What's a QBIC catalog? A. It's an instance of a table that has a bunch of images in it. Q. And forgive me, what's an instance of a table? A. You create a table -- create a database, and then in the database you create a table and you have maybe like a file name, and then the actual image bits into it. Q. And how would you search for images using that table? A. You could search -- you could use whatever you wanted to for the -- the SQL query Page 78 to -- to filter, and then that would give you a list of results. And then you'd rank all of those results based on the QBIC distance measures. Flickner Depo at 88:

		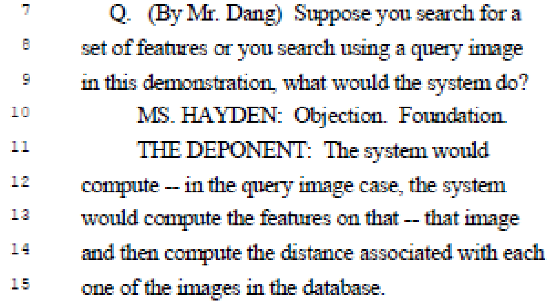 Q. (By Mr. Dang) Suppose you search for a set of features or you search using a query image in this demonstration, what would the system do? MS. HAYDEN: Objection. Foundation. THE DEPONENT: The system would compute -- in the query image case, the system would compute the features on that -- that image and then compute the distance associated with each one of the images in the database. Flickner Depo at 88-89:

		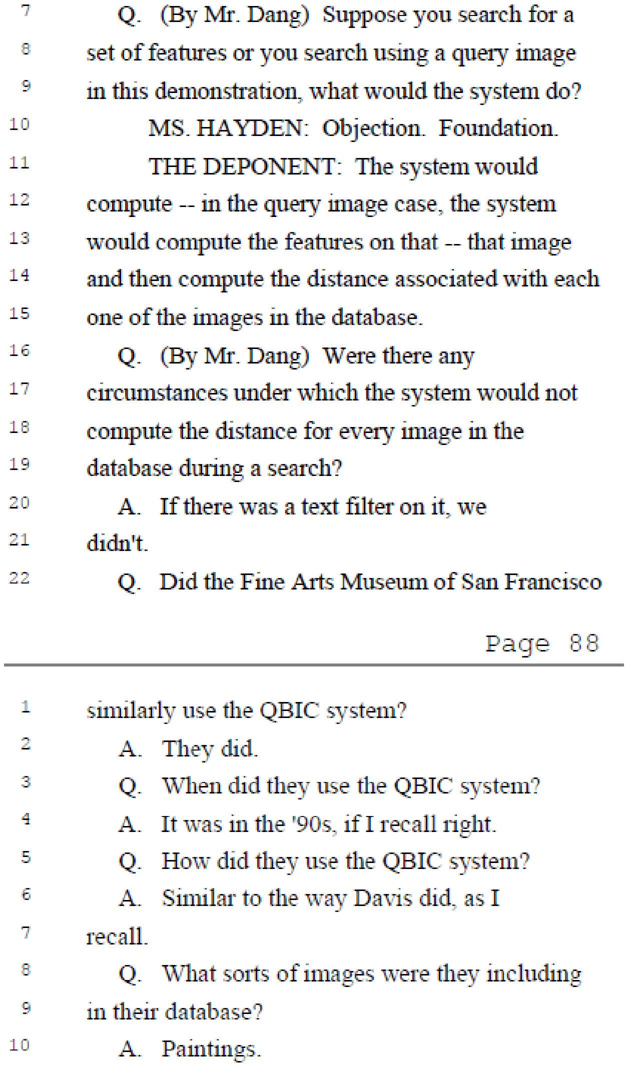 Q. (By Mr. Dang) Suppose you search for a set of features or you search using a query image in this demonstration, what would the system do? MS. HAYDEN: Objection. Foundation. THE DEPONENT: The system would compute -- in the query image case, the system would compute the features on that -- that image and then compute the distance associated with each one of the images in the database. Q. (By Mr. Dang) Were there any circumstances under which the system would not compute the distance for every image in the database during a search? A. If there was a text filter on it, we didn't. Q. Did the Fine Arts Museum of San Francisco similarly use the QBIC system? Page 88 A. They did. Q. When did they use the QBIC system? A. It was in the '90s, if I recall right. Q. How did they use the QBIC system? A. Similar to the way Davis did, as I recall. Q. What sorts of images were they including in their database? A. Paintings. Flickner Depo at 89-90:

		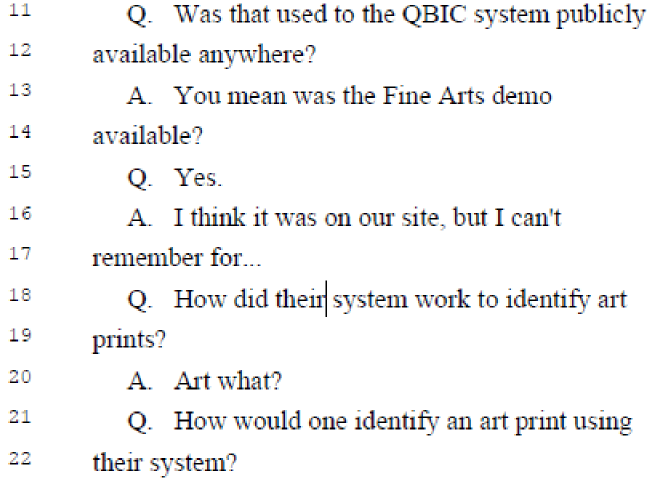Q. Was that used to the QBIC system publicly available anywhere? A. You mean was the Fine Arts demo available? Q. Yes. A. I think it was on our site, but I can't remember for... Q. How did their system work to identify art prints? A. Art what? Q. How would one identify an art print using their system? 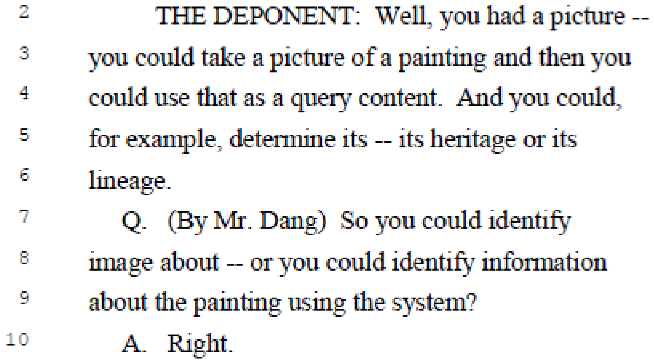THE DEPONENT: Well, you had a picture -- you could take a picture of a painting and then you could use that as a query content. And you could, for example, determine its -- its heritage or its lineage. Q. (By Mr. Dang) So you could identify image about -- or you could identify information about the painting using the system? A. Right. Flickner Depo at 91-92:

		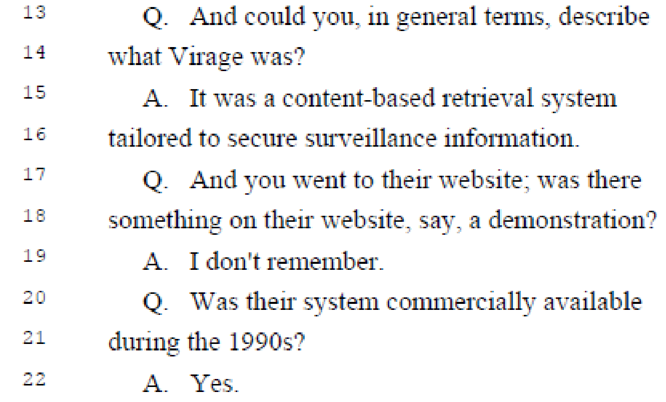 Q. And could you, in general terms, describe what Virage was? A. It was a content-based retrieval system tailored to secure surveillance information. Q. And you went to their website; was there something on their website, say, a demonstration? A. I don't remember. Q. Was their system commercially available during the 1990s? A. Yes. 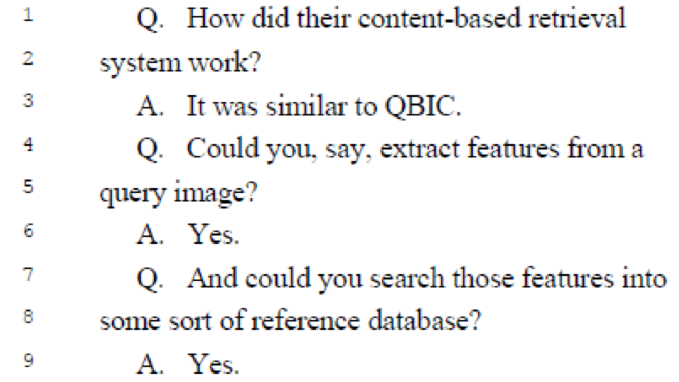 Q. How did their content-based retrieval system work? A. It was similar to QBIC. Q. Could you, say, extract features from a query image? A. Yes. Q. And could you search those features into some sort of reference database? A. Yes. Flickner Depo at 131-132: 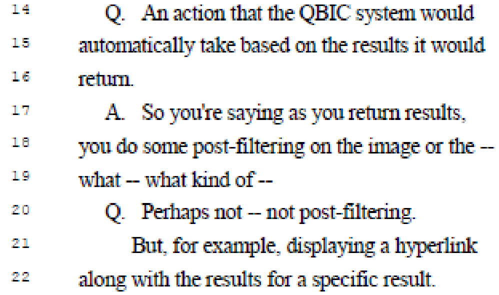 Q. An action that the QBIC system would automatically take based on the results it would return. A. So you're saying as you return results, you do some post-filtering on the image or the -- what -- what kind of -- Q. Perhaps not -- not post-filtering. But, for example, displaying a hyperlink along with the results for a specific result.  A. Yeah, you could do that.

Flickner Depo at 142:

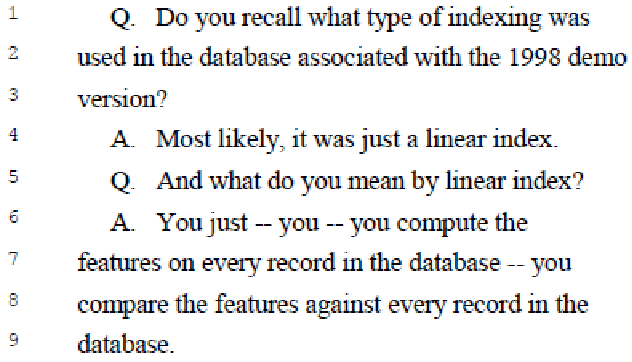

Q. Do you recall what type of indexing was
used in the database associated with the 1998 demo
version?
A. Most likely, it was just a linear index.
Q. And what do you mean by linear index?
A. You just -- you -- you compute the
features on every record in the database -- you
compare the features against every record in the
database.

Flickner Depo at 159:

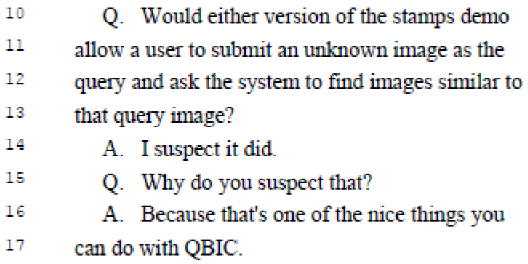

Q. Would either version of the stamps demo
allow a user to submit an unknown image as the
query and ask the system to find images similar to
that query image?
A. I suspect it did.
Q. Why do you suspect that?
A. Because that's one of the nice things you
can do with QBIC.

Flickner Depo at 162-163:

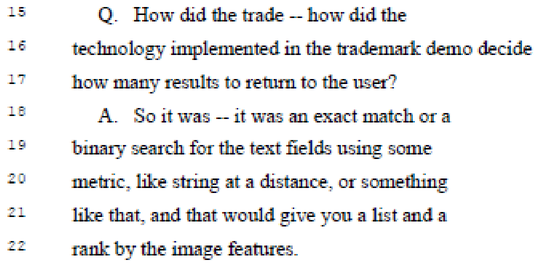

Q. How did the trade -- how did the
technology implemented in the trademark demo decide
how many results to return to the user?
A. So it was -- it was an exact match or a
binary search for the text fields using some
metric, like string at a distance, or something
like that, and that would give you a list and a
rank by the image features.

		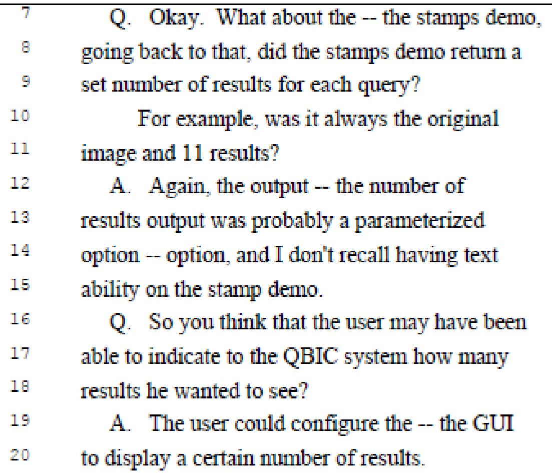 Q. Okay. What about the -- the stamps demo, going back to that, did the stamps demo return a set number of results for each query? For example, was it always the original image and 11 results? A. Again, the output -- the number of results output was probably a parameterized option -- option, and I don't recall having text ability on the stamp demo. Q. So you think that the user may have been able to indicate to the QBIC system how many results he wanted to see? A. The user could configure the -- the GUI to display a certain number of results. Flickner Depo at 168-169:

		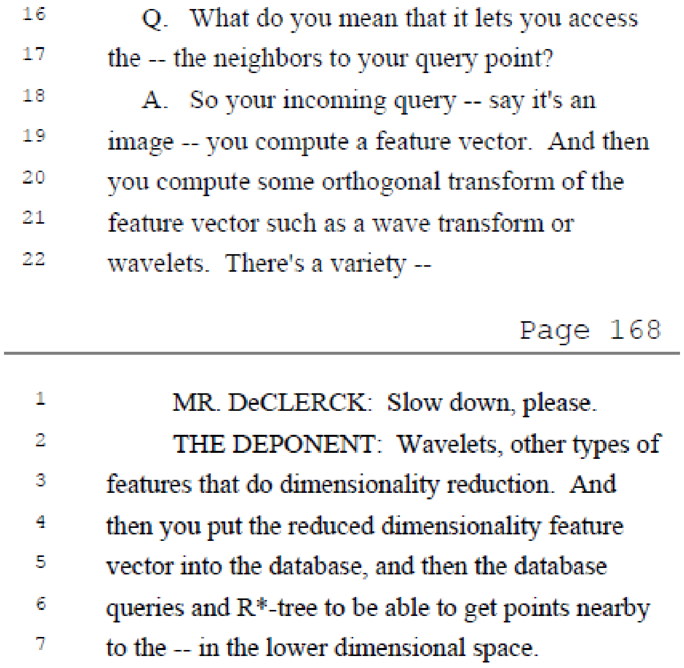 Q. What do you mean that it lets you access the -- the neighbors to your query point? A. So your incoming query -- say it's an image -- you compute a feature vector. And then you compute some orthogonal transform of the feature vector such as a wave transform or wavelets. There's a variety -- Page 168 MR. DeCLERCK: Slow down, please. THE DEPONENT: Wavelets, other types of features that do dimensionality reduction. And then you put the reduced dimensionality feature vector into the database, and then the database queries and R*-tree to be able to get points nearby to the -- in the lower dimensional space. Flickner Depo at 181:

		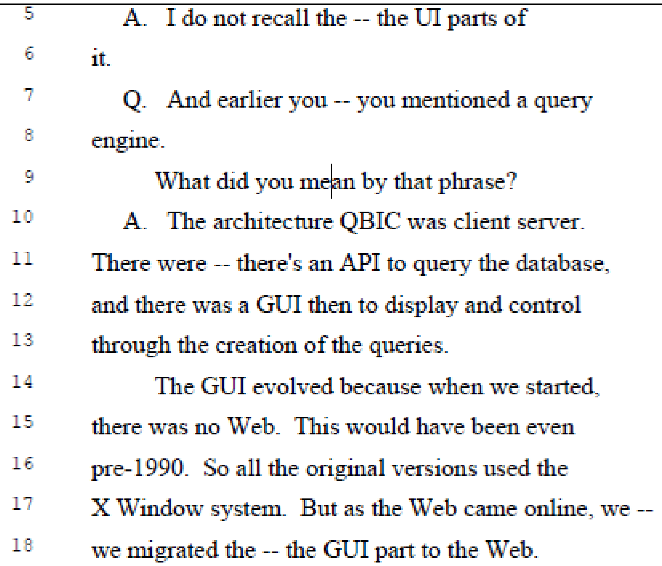 A. I do not recall the -- the UI parts of it. Q. And earlier you -- you mentioned a query engine. What did you mean by that phrase? A. The architecture QBIC was client server. There were -- there's an API to query the database, and there was a GUI then to display and control through the creation of the queries. The GUI evolved because when we started, there was no Web. This would have been even pre-1990. So all the original versions used the X Window system. But as the Web came online, we -- we migrated the -- the GUI part to the Web.

		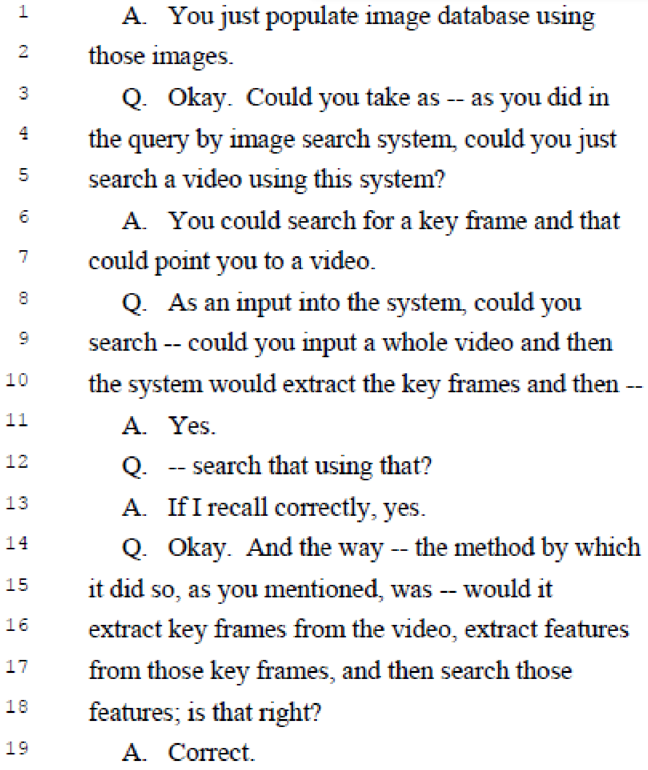 A. You just populate image database using those images. Q. Okay. Could you take as -- as you did in the query by image search system, could you just search a video using this system? A. You could search for a key frame and that could point you to a video. Q. As an input into the system, could you search -- could you input a whole video and then the system would extract the key frames and then -- A. Yes. Q. -- search that using that? A. If I recall correctly, yes. Q. Okay. And the way -- the method by which it did so, as you mentioned, was -- would it extract key frames from the video, extract features from those key frames, and then search those features; is that right? A. Correct. Flickner Depo at 201:

		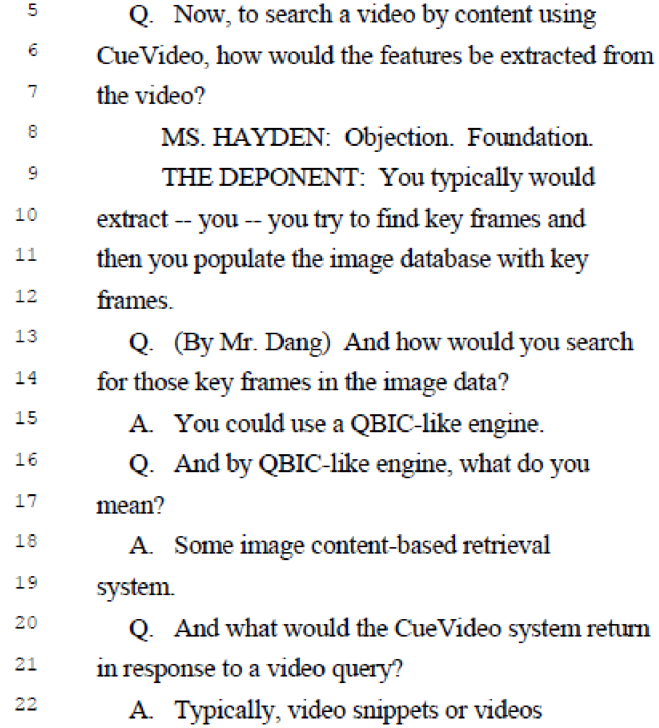 Q. Now, to search a video by content using CueVideo, how would the features be extracted from the video? MS. HAYDEN: Objection. Foundation. THE DEPONENT: You typically would extract -- you -- you try to find key frames and then you populate the image database with key frames. Q. (By Mr. Dang) And how would you search for those key frames in the image data? A. You could use a QBIC-like engine. Q. And by QBIC-like engine, what do you mean? A. Some image content-based retrieval system. Q. And what would the CueVideo system return in response to a video query? A. Typically, video snippets or videos Flickner Depo at 201-202:

		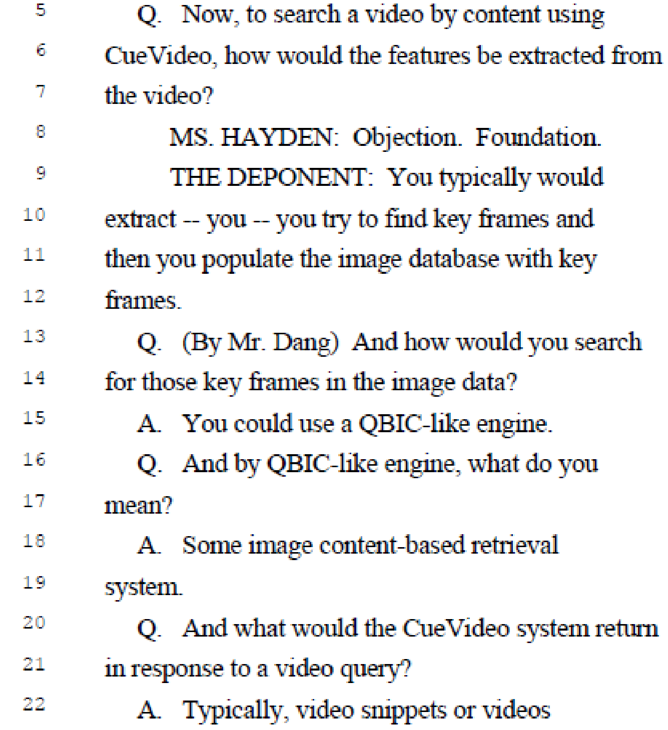 Q. Now, to search a video by content using CueVideo, how would the features be extracted from the video? MS. HAYDEN: Objection. Foundation. THE DEPONENT: You typically would extract -- you -- you try to find key frames and then you populate the image database with key frames. Q. (By Mr. Dang) And how would you search for those key frames in the image data? A. You could use a QBIC-like engine. Q. And by QBIC-like engine, what do you mean? A. Some image content-based retrieval system. Q. And what would the CueVideo system return in response to a video query? A. Typically, video snippets or videos

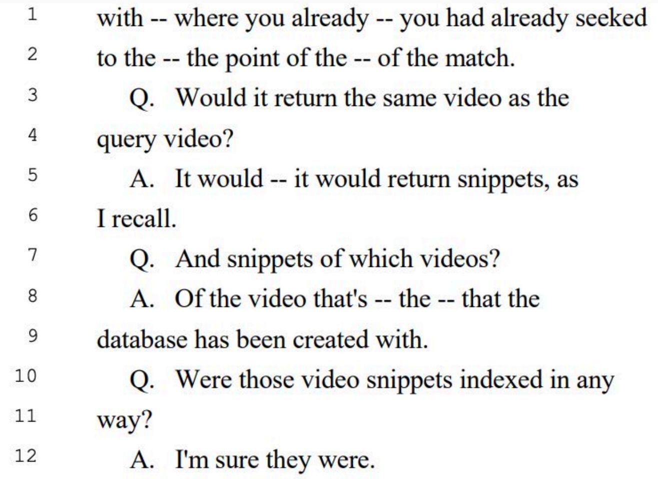

with -- where you already -- you had already seeked to the -- the point of the -- of the match.

Q. Would it return the same video as the query video?

A. It would -- it would return snippets, as I recall.

Q. And snippets of which videos?

A. Of the video that's -- the -- that the database has been created with.

Q. Were those video snippets indexed in any way?

A. I'm sure they were.

Flickner Depo at 209-210:

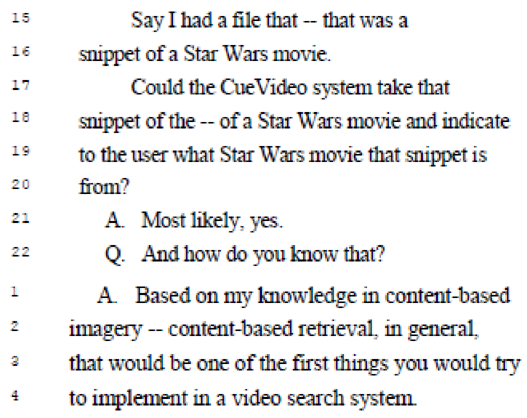

Q. Okay. Let me ask -- let me ask that question using -- using an example.

Say I had a file that -- that was a snippet of a Star Wars movie.

Could the CueVideo system take that snippet of the -- of a Star Wars movie and indicate to the user what Star Wars movie that snippet is from?

A. Most likely, yes.

Q. And how do you know that?

A. Based on my knowledge in content-based imagery -- content-based retrieval, in general, that would be one of the first things you would try to implement in a video search system.

Flickner Depo at 212-213:

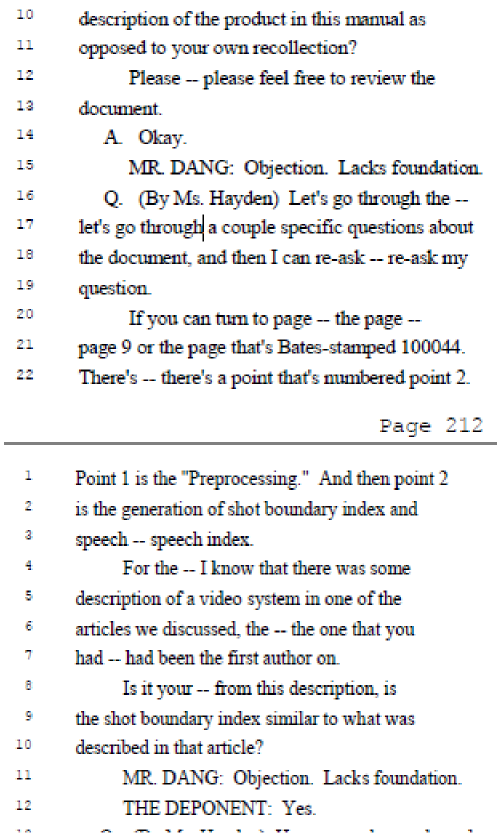

Q. So comparing your -- the understanding of
CueVideo that you've described so far and this
manual, would you -- would you defer to the
description of the product in this manual as
opposed to your own recollection?
Please -- please feel free to review the
document.
A. Okay.
MR. DANG: Objection. Lacks foundation.
Q. (By Ms. Hayden) Let's go through the --
let's go through a couple specific questions about
the document, and then I can re-ask -- re-ask my
question.
If you can turn to page -- the page --
page 9 or the page that's Bates-stamped 100044.
There's -- there's a point that's numbered point 2.

Page 212

Point 1 is the "Preprocessing." And then point 2
is the generation of shot boundary index and
speech -- speech index.
For the -- I know that there was some
description of a video system in one of the
articles we discussed, the -- the one that you
had -- had been the first author on.
Is it your -- from this description, is
the shot boundary index similar to what was
described in that article?
MR. DANG: Objection. Lacks foundation.
THE DEPONENT: Yes.

Flickner Depo at 214-215:

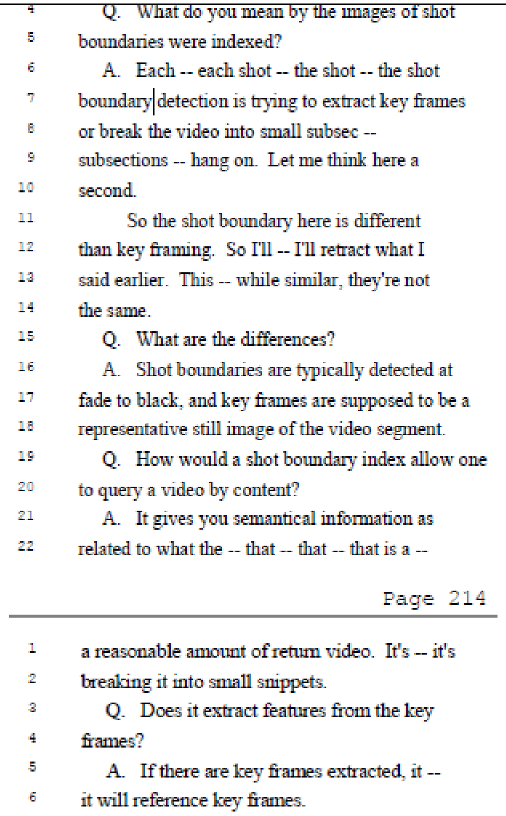

Q. What do you mean by the images of shot boundaries were indexed?

A. Each -- each shot -- the shot -- the shot boundary detection is trying to extract key frames or break the video into small subsec -- subsections -- hang on. Let me think here a second.

So the shot boundary here is different than key framing. So I'll -- I'll retract what I said earlier. This -- while similar, they're not the same.

Q. What are the differences?

A. Shot boundaries are typically detected at fade to black, and key frames are supposed to be a representative still image of the video segment.

Q. How would a shot boundary index allow one to query a video by content?

A. It gives you semantical information as related to what the -- that -- that -- that is a --

Page 214

a reasonable amount of return video. It's -- it's breaking it into small snippets.

Q. Does it extract features from the key frames?

A. If there are key frames extracted, it -- it will reference key frames.

Query by Image and Video Content: The QBIC System (IBM 0000893–902):
"We have developed the QBIC (Query by Image Content) system to explore content-based retrieval methods. QBIC allows queries on large image and video databases based on • example images, • user-constructed sketches and drawings, • selected color and texture patterns, • camera and object motion, and • other graphical information." *Id*. at pp. 1–3.

"To achieve this functionality, QBIC has two main components: database population (the process of creating an image database) and database query. During the population, images and videos are processed to extract features describing their content-colors,

		textures, shapes, and camera and object motion-and the features are stored in a database. During the query, the user composes a query graphically. Features are generated from the graphical query and then input to a matching engine that finds images or videos from the database with similar features. Figure 2 shows the system architecture." *Id*. at p. 3. "For both population and query, the QBIC data model has • still images or scenes (full images) that contain objects (subsets of an image), and • video shots that consist of sets of contiguous frames and contain motion objects. For still images, the QBIC data model distinguishes between 'scenes' (or images) and 'objects.' A scene is an image or single representative frame of video. An object is a part of a scene-for example, the fox in Figure 3-or a moving entity in a video. For still image database population, features are extracted from images and objects and stored in a database as shown in the top left part of Figure 2. Videos are broken into clips called shots. Representative frames, or r-frames, are generated for each extracted shot. R-frames are treated as still images, and features are extracted and stored in the database. Further processing of shots generates motion objects-for example, a car moving across the screen. Queries are allowed on objects ('Find images with a red, round object'), scenes ('Find images that have approximately 30-percent red and IS-percent blue colors'), shots ('Find all shots panning from left to right'), or any combination ('Find images that have 30 percent red and contain a blue textured object')." *Id*. at p. 3. "**SAMPLE QUERIES** For each full-scene image, identified image object, r-frame, and identified video object resulting from the above processing, a set of features is computed to allow content-based queries. The features are computed and stored during database population. We present a brief description of the features and the associated queries. Mathematical details on the features and matching methods can be found in Ashley et al. and Niblack et al. Average color queries let users find images or objects that are similar to a selected color, say from a color wheel, or to the color of an object. The feature used in the query is a 3D vector of Munsell color coordinates. Histogram color queries return items with

		matching color distributions—say, a fabric pattern with approximately 40 percent red and 20 percent blue. For this case, the underlying feature is a 256-element histogram computed over a quantized version of the color space. Figure 7 shows a histogram query on still images and a color query on video r-frames. Note that in the query specification for the histogram query of Figure 7, the user has selected percentages of two colors (blue and white) by adjusting sliders. Using such a query, an advertising agent could, for example, search for a picture of a beach scene, one predominantly blue (for sky and water) and white (for sand and clouds); or find images with similar color spreads for a uniform ad campaign. The average color query demonstrates a query against a video shot database where the user is searching for red r-frames. Again, the query specification is on the left and the best hits are on the right. Figure 8 shows an example texture query. In this case, the query is specified by selecting from a sampler—a set of prestored example images. The underlying texture features are mathematical representations of coarseness, contrast, and directionality features. Coarseness measures the scale of a texture (pebbles versus boulders), contrast describes its vividness, and directionality describes whether it has a favored direction (like grass) or not (like a smooth object). An object shape query is shown in Figure 8. In this case, the query specification is the drawn shape on the left. Area, circularity, eccentricity, major-axis direction, features derived from the object moments, and a set of tangent angles around the object perimeter are the features used to characterize and match shapes. Figure 9 illustrates query by sketch. In this case, the query specification is a freehand drawing of the dominant lines and edges in the image. The sketch feature is an automatically extracted reduced-resolution 'edge map.' Matching is done by using a template-matching technique. A multiobject query asking for images that contain both a red round object and a green textured object is shown in the bottom of Figure 9. The features are standard color and texture. The matching is done by combining the color and texture distances. Combining distances is applied to arbitrary sets of objects and features to implement logical And semantics. . . . . **ANNOTATION AND DATABASE POPULATION TOOLS**. Automatic methods (such as our Positional Color query) that don't rely on object identification, methods

		that identify objects automatically as in the museum image example, fast and easy-to-use semiautomatic outlining tools, and motion-based segmentation algorithms will enable additional application areas. **FEATURE EXTRACTION AND MATCHING METHODS**. New mathematical representations of video, image, and object attributes that capture 'interesting' features for retrieval are needed. Features that describe new image properties such as alternate texture measures or that are based on fractals or wavelet representations, for example, may offer advantages of representation, indexability, and ease of similarity matching. **INTEGRATION WITH TEXT AND PARAMETRIC ANNOTATION**. Query by visual content complements and extends existing query methods. Systems must be able to integrate queries combining date, subject matter, price, and availability with content properties such as color, texture, and shape." *Id.* at pp. 7–8. Image, Audio, and Video Extenders Administration and Programming (IBM 000001–617): "Your applications can even search for images by content. Imagine an application that uses visual examples to search for images. With such an application, users could select an example image and have the application find other images that have colors or textures similar to those in the example. With DB2 extenders' Query by Image Content (QBIC) capability, you can create applications that search for images in this visual way." *Id.* at p. 4. "To make an image available for searching by content, you catalog the image. When you catalog an image, the Image Extender analyzes the image, by computing the feature values for the image, and stores the values in a QBIC catalog. When you search for an image by content, your query identifies one or more features for the search (such as average color), a source for each feature (such as an example image), and a target set of cataloged images. The Image Extender computes the feature value of the source and compares it to the cataloged feature values for the target images. It then computes a score that indicates how similar the feature values of the target images are to the source. You can have the Image Extender return the images whose features are most similar to the source. The Image Extender will return the handle of each image and the image

		score. You can also have the Image Extender return only the score of a single image." *Id.* at pp. 20–21. "In still image database population, the images are reduced to a standard-sized icon called a thumbnail and annotated with any available text information. Object identification is an optional but key part of this step. It lets users manually, semiautomatically, or fully automatically identify interesting regions-which we call objects-in the images. . . . For video data, database population has three major components: • shot detection, • representative frame creation for each shot, and • derivation of a layered representation of coherently moving structures/ objects. . . . Once the shot boundaries have been detected, each shot is represented using an r-frame. R-frames are used for several purposes. First, during database population, r-frames are treated as still images in which objects can be identified by using the previously described methods." *Id.* at pp. 3–6. Image, Audio, and Video Extenders Administration and Programming (IBM 000001–617): "Your applications can even search for images by content. Imagine an application that uses visual examples to search for images. With such an application, users could select an example image and have the application find other images that have colors or textures similar to those in the example. With DB2 extenders' Query by Image Content (QBIC) capability, you can create applications that search for images in this visual way." *Id.* at p. 4. "To make an image available for searching by content, you catalog the image. When you catalog an image, the Image Extender analyzes the image, by computing the feature values for the image, and stores the values in a QBIC catalog. When you search for an image by content, your query identifies one or more features for the search (such as average color), a source for each feature (such as an example image), and a target set of cataloged images. The Image Extender computes the feature value of the source and compares it to the cataloged feature values for the target

		images. It then computes a score that indicates how similar the feature values of the target images are to the source. You can have the Image Extender return the images whose features are most similar to the source. The Image Extender will return the handle of each image and the image score. You can also have the Image Extender return only the score of a single image." *Id*. at pp. 20–21. 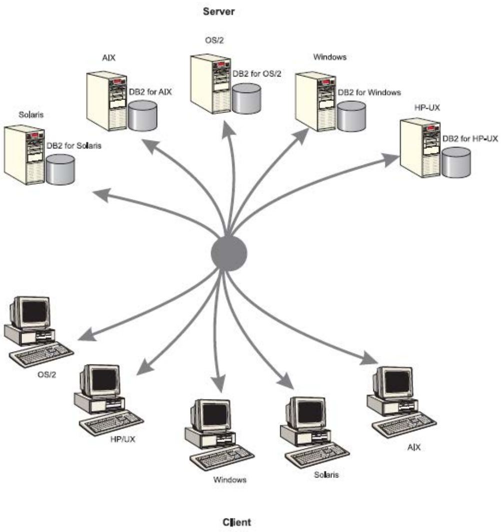 *Figure 6. DB2 extender platforms* *Id*. at p. 12 (Fig. 6). "The DB2 Extenders run in the DB2 client/server environment. This environment consists of a database server and one or more remote database clients. The DB2 extender services run on the server. Before you can access them, you have to start them." *Id*. at p. 47.

		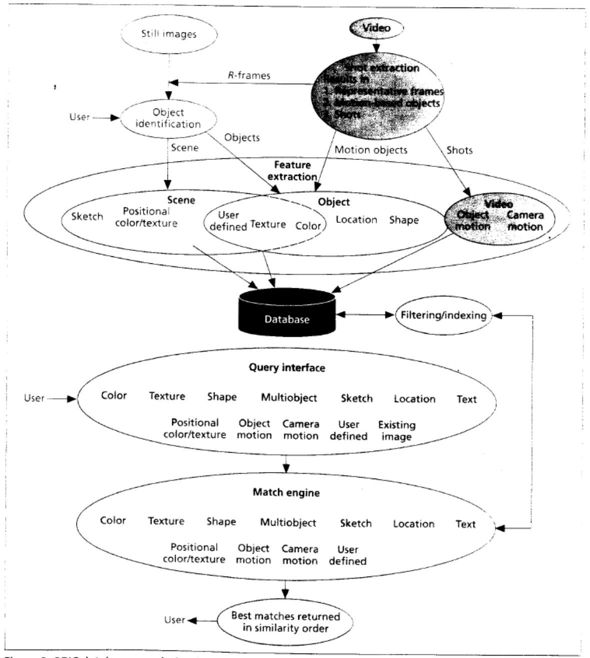 Figure 2. QBIC database population (top) and query (bottom) architecture. *Id*. at Fig. 2.

		QBIC Demo (IBM000747) – showing compatibility with images captured by a mobile phone camera 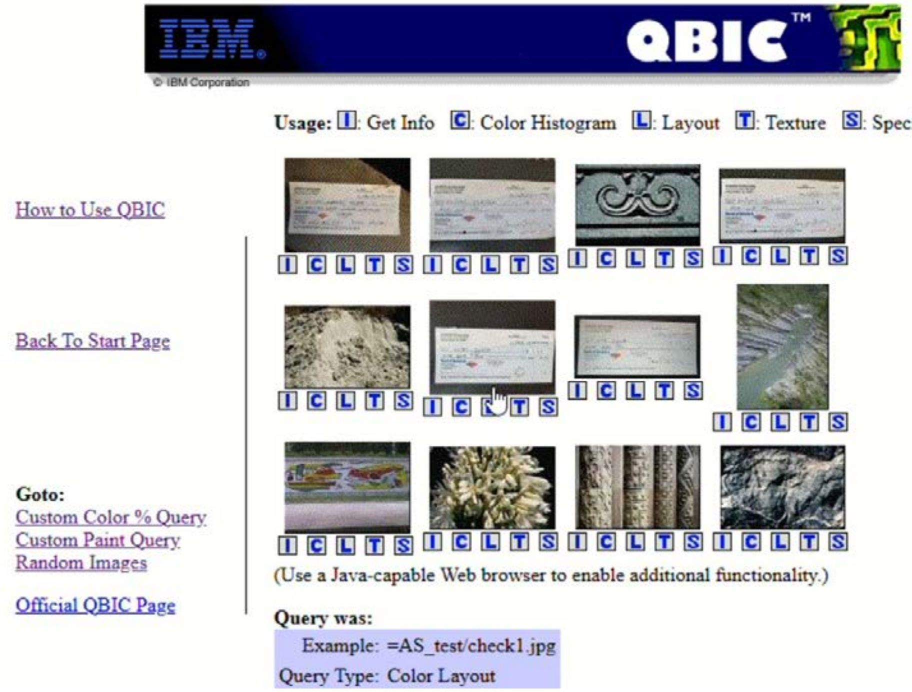 (Use a Java-capable Web browser to enable additional functionality.) Query was: Example: =AS_test/check1.jpg Query Type: Color Layout

		 check1.jpg.thumb0.jpg check2_orig.jpg.thumb0.jpg check3.jpg.thumb0.jpg check4.jpg.thumb0.jpg check5.jpg.thumb0.jpg check5_rotate_crop_fix.jpg.thumb0.jpg

		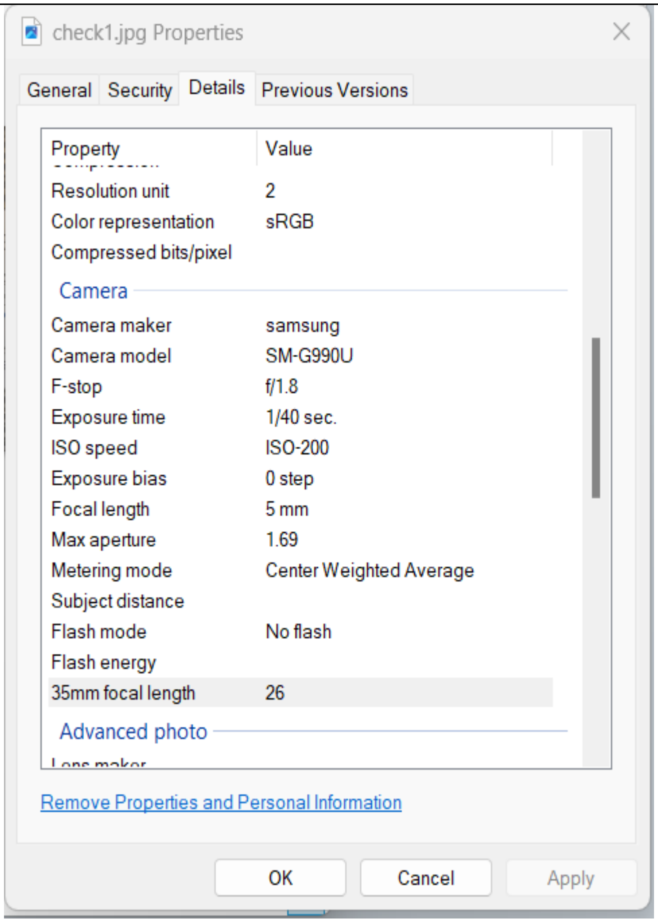

Property	Value
Resolution unit	2
Color representation	sRGB
Compressed bits/pixel	
Camera	
Camera maker	samsung
Camera model	SM-G990U
F-stop	f/1.8
Exposure time	1/40 sec.
ISO speed	ISO-200
Exposure bias	0 step
Focal length	5 mm
Max aperture	1.69
Metering mode	Center Weighted Average
Subject distance	
Flash mode	No flash
Flash energy	
35mm focal length	26
Advanced photo	
Lens maker	

		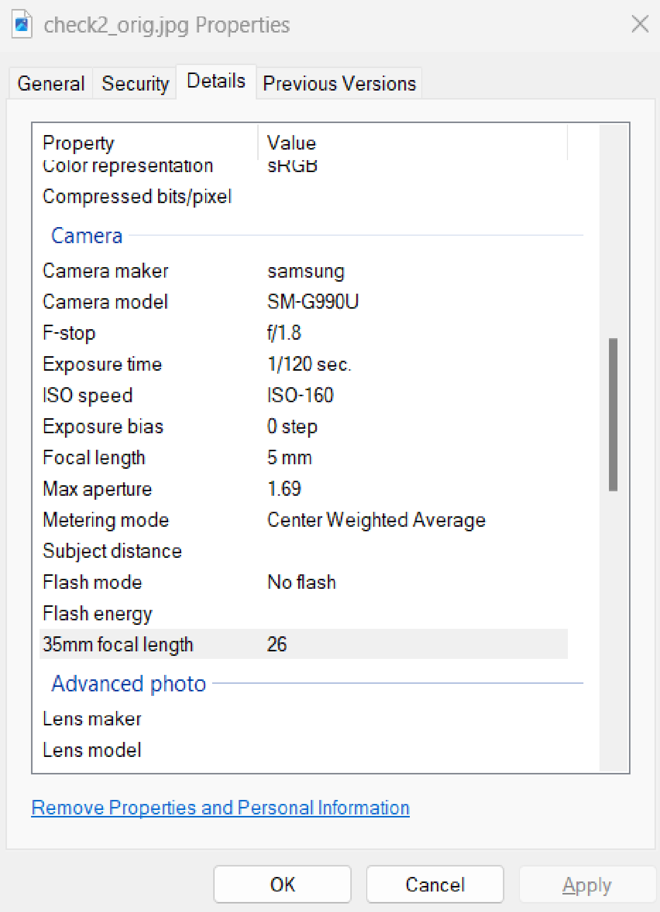

check2_orig.jpg Properties

Property	Value
Color representation	sRGB
Compressed bits/pixel	
Camera	
Camera maker	samsung
Camera model	SM-G990U
F-stop	f/1.8
Exposure time	1/120 sec.
ISO speed	ISO-160
Exposure bias	0 step
Focal length	5 mm
Max aperture	1.69
Metering mode	Center Weighted Average
Subject distance	
Flash mode	No flash
Flash energy	
35mm focal length	26
Advanced photo	
Lens maker	
Lens model	

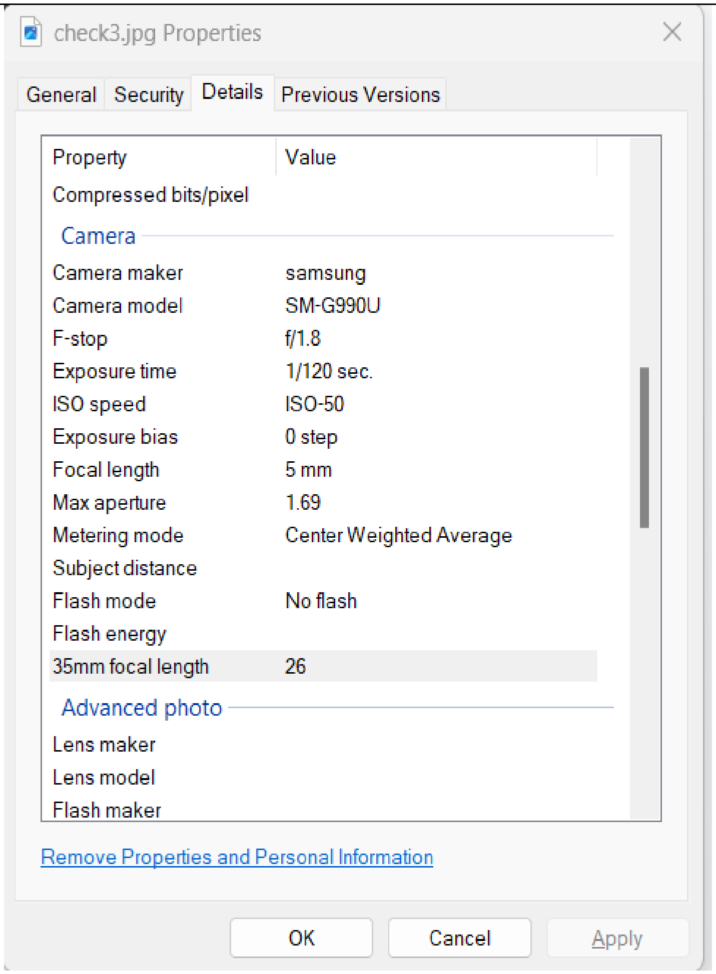

Property	Value
Compressed bits/pixel	
Camera	
Camera maker	samsung
Camera model	SM-G990U
F-stop	f/1.8
Exposure time	1/120 sec.
ISO speed	ISO-50
Exposure bias	0 step
Focal length	5 mm
Max aperture	1.69
Metering mode	Center Weighted Average
Subject distance	
Flash mode	No flash
Flash energy	
35mm focal length	26
Advanced photo	
Lens maker	
Lens model	
Flash maker	

		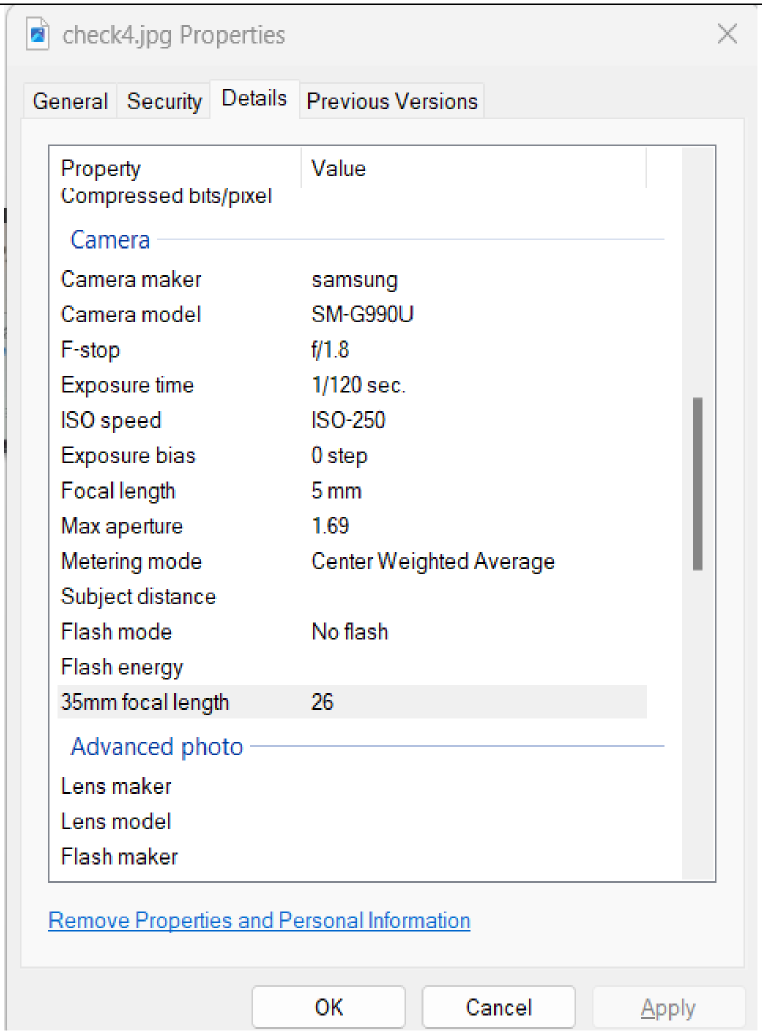

check4.jpg Properties

Property	Value
Compressed bits/pixel	
Camera	
Camera maker	samsung
Camera model	SM-G990U
F-stop	f/1.8
Exposure time	1/120 sec.
ISO speed	ISO-250
Exposure bias	0 step
Focal length	5 mm
Max aperture	1.69
Metering mode	Center Weighted Average
Subject distance	
Flash mode	No flash
Flash energy	
35mm focal length	26
Advanced photo	
Lens maker	
Lens model	
Flash maker	

		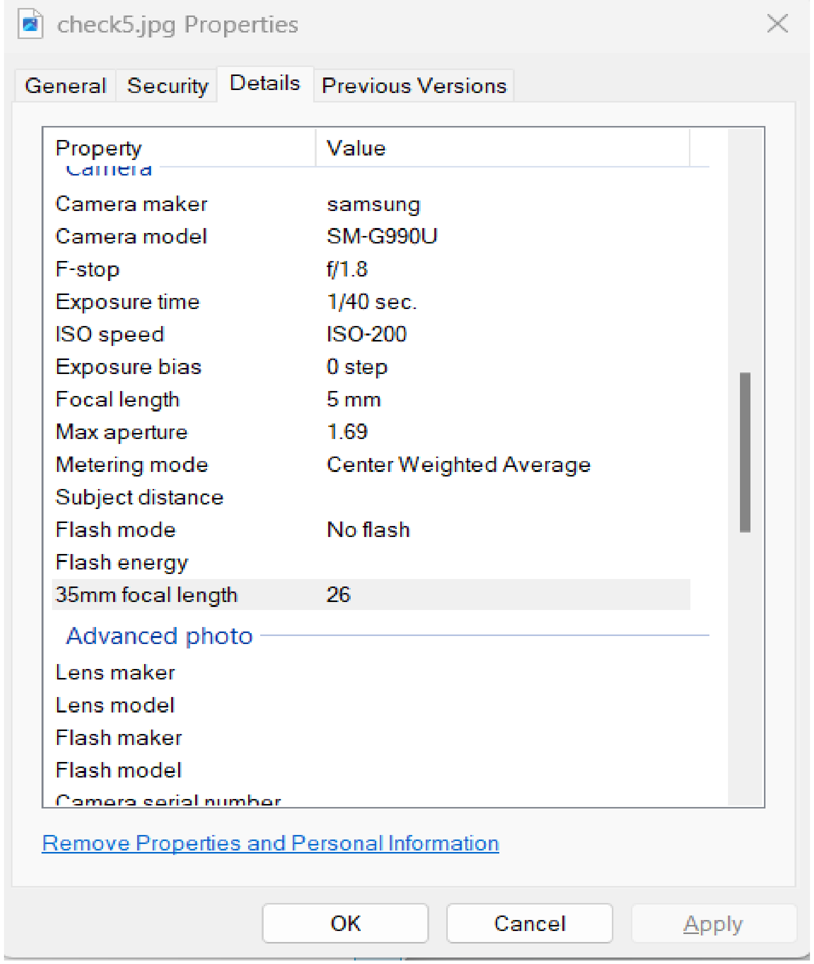 Property: Value Camera Camera maker: samsung Camera model: SM-G990U F-stop: f/1.8 Exposure time: 1/40 sec. ISO speed: ISO-200 Exposure bias: 0 step Focal length: 5 mm Max aperture: 1.69 Metering mode: Center Weighted Average Subject distance Flash mode: No flash Flash energy 35mm focal length: 26 Advanced photo Lens maker Lens model Flash maker Flash model Camera serial number The QBIC Project: Querying Images by Content Using Color, Texture, and Shape (IBM 0002390–404) IBM0002390 at 2390:

		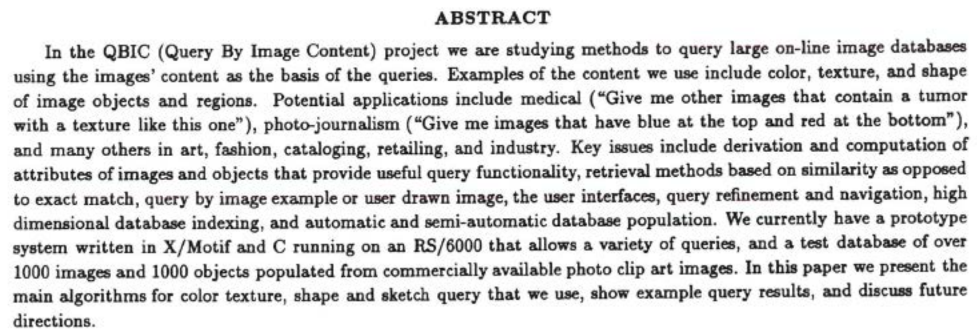 **ABSTRACT** In the QBIC (Query By Image Content) project we are studying methods to query large on-line image databases using the images' content as the basis of the queries. Examples of the content we use include color, texture, and shape of image objects and regions. Potential applications include medical ("Give me other images that contain a tumor with a texture like this one"), photo-journalism ("Give me images that have blue at the top and red at the bottom"), and many others in art, fashion, cataloging, retailing, and industry. Key issues include derivation and computation of attributes of images and objects that provide useful query functionality, retrieval methods based on similarity as opposed to exact match, query by image example or user drawn image, the user interfaces, query refinement and navigation, high dimensional database indexing, and automatic and semi-automatic database population. We currently have a prototype system written in X/Motif and C running on an RS/6000 that allows a variety of queries, and a test database of over 1000 images and 1000 objects populated from commercially available photo clip art images. In this paper we present the main algorithms for color texture, shape and sketch query that we use, show example query results, and discuss future directions. IBM0002390 at 2391: 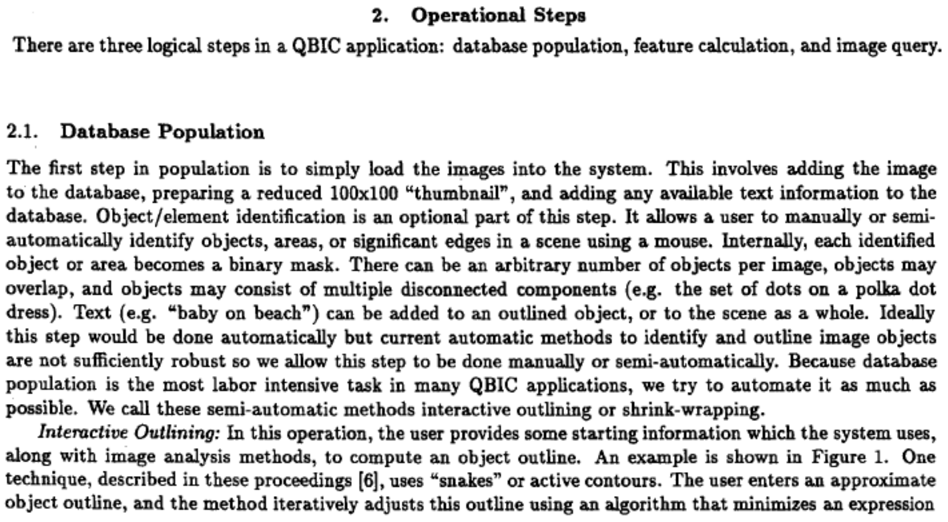 **2. Operational Steps** There are three logical steps in a QBIC application: database population, feature calculation, and image query. **2.1. Database Population** The first step in population is to simply load the images into the system. This involves adding the image to the database, preparing a reduced 100x100 "thumbnail", and adding any available text information to the database. Object/element identification is an optional part of this step. It allows a user to manually or semi-automatically identify objects, areas, or significant edges in a scene using a mouse. Internally, each identified object or area becomes a binary mask. There can be an arbitrary number of objects per image, objects may overlap, and objects may consist of multiple disconnected components (e.g. the set of dots on a polka dot dress). Text (e.g. "baby on beach") can be added to an outlined object, or to the scene as a whole. Ideally this step would be done automatically but current automatic methods to identify and outline image objects are not sufficiently robust so we allow this step to be done manually or semi-automatically. Because database population is the most labor intensive task in many QBIC applications, we try to automate it as much as possible. We call these semi-automatic methods interactive outlining or shrink-wrapping. *Interactive Outlining:* In this operation, the user provides some starting information which the system uses, along with image analysis methods, to compute an object outline. An example is shown in Figure 1. One technique, described in these proceedings [6], uses "snakes" or active contours. The user enters an approximate object outline, and the method iteratively adjusts this outline using an algorithm that minimizes an expression

		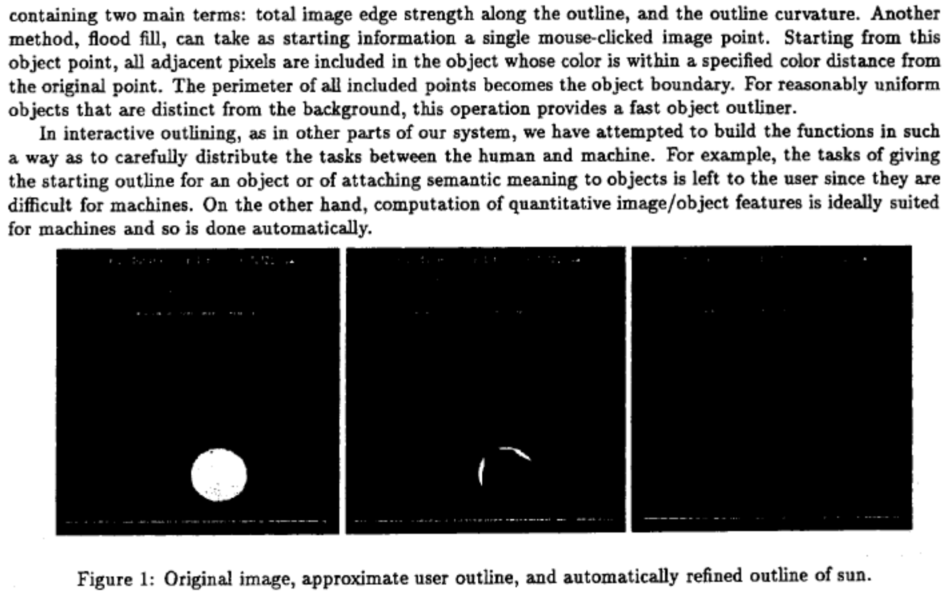 containing two main terms: total image edge strength along the outline, and the outline curvature. Another method, flood fill, can take as starting information a single mouse-clicked image point. Starting from this object point, all adjacent pixels are included in the object whose color is within a specified color distance from the original point. The perimeter of all included points becomes the object boundary. For reasonably uniform objects that are distinct from the background, this operation provides a fast object outliner. In interactive outlining, as in other parts of our system, we have attempted to build the functions in such a way as to carefully distribute the tasks between the human and machine. For example, the tasks of giving the starting outline for an object or of attaching semantic meaning to objects is left to the user since they are difficult for machines. On the other hand, computation of quantitative image/object features is ideally suited for machines and so is done automatically. Figure 1: Original image, approximate user outline, and automatically refined outline of sun. IBM0002390 at 2391-92: 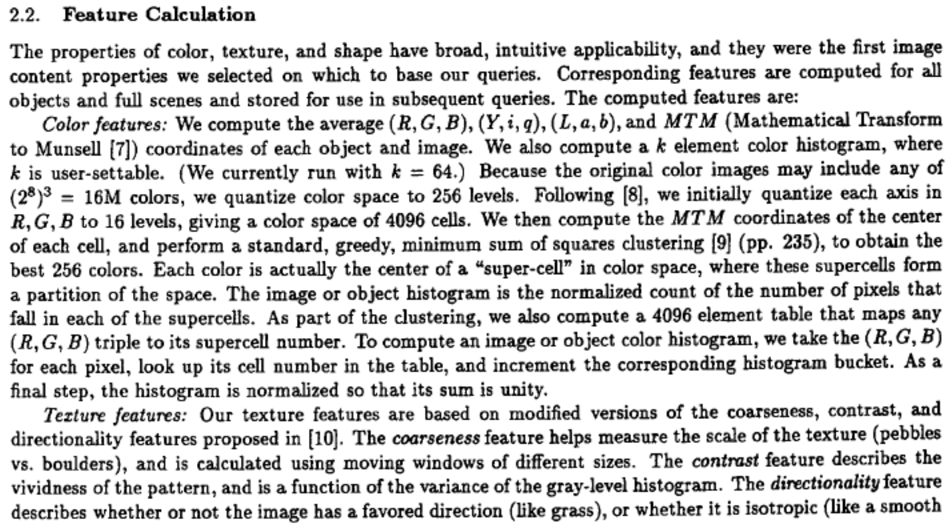 **2.2. Feature Calculation** The properties of color, texture, and shape have broad, intuitive applicability, and they were the first image content properties we selected on which to base our queries. Corresponding features are computed for all objects and full scenes and stored for use in subsequent queries. The computed features are: *Color features:* We compute the average (R, G, B), (Y, i, q), (L, a, b), and MTM (Mathematical Transform to Munsell [7]) coordinates of each object and image. We also compute a k element color histogram, where k is user-settable. (We currently run with $k = 64$.) Because the original color images may include any of $(2^8)^3 = 16$M colors, we quantize color space to 256 levels. Following [8], we initially quantize each axis in R, G, B to 16 levels, giving a color space of 4096 cells. We then compute the MTM coordinates of the center of each cell, and perform a standard, greedy, minimum sum of squares clustering [9] (pp. 235), to obtain the best 256 colors. Each color is actually the center of a "super-cell" in color space, where these supercells form a partition of the space. The image or object histogram is the normalized count of the number of pixels that fall in each of the supercells. As part of the clustering, we also compute a 4096 element table that maps any (R, G, B) triple to its supercell number. To compute an image or object color histogram, we take the (R, G, B) for each pixel, look up its cell number in the table, and increment the corresponding histogram bucket. As a final step, the histogram is normalized so that its sum is unity. *Texture features:* Our texture features are based on modified versions of the coarseness, contrast, and directionality features proposed in [10]. The *coarseness* feature helps measure the scale of the texture (pebbles vs. boulders), and is calculated using moving windows of different sizes. The *contrast* feature describes the vividness of the pattern, and is a function of the variance of the gray-level histogram. The *directionality* feature describes whether or not the image has a favored direction (like grass), or whether it is isotropic (like a smooth

		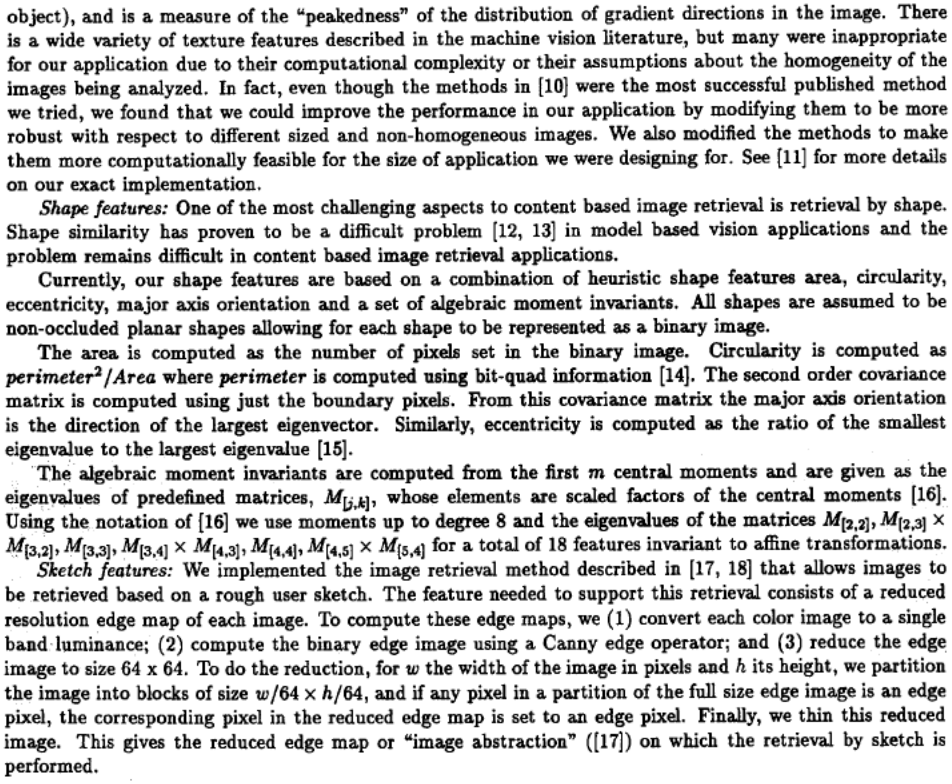 object), and is a measure of the "peakedness" of the distribution of gradient directions in the image. There is a wide variety of texture features described in the machine vision literature, but many were inappropriate for our application due to their computational complexity or their assumptions about the homogeneity of the images being analyzed. In fact, even though the methods in [10] were the most successful published method we tried, we found that we could improve the performance in our application by modifying them to be more robust with respect to different sized and non-homogeneous images. We also modified the methods to make them more computationally feasible for the size of application we were designing for. See [11] for more details on our exact implementation. *Shape features:* One of the most challenging aspects to content based image retrieval is retrieval by shape. Shape similarity has proven to be a difficult problem [12, 13] in model based vision applications and the problem remains difficult in content based image retrieval applications. Currently, our shape features are based on a combination of heuristic shape features area, circularity, eccentricity, major axis orientation and a set of algebraic moment invariants. All shapes are assumed to be non-occluded planar shapes allowing for each shape to be represented as a binary image. The area is computed as the number of pixels set in the binary image. Circularity is computed as $perimeter^2/Area$ where $perimeter$ is computed using bit-quad information [14]. The second order covariance matrix is computed using just the boundary pixels. From this covariance matrix the major axis orientation is the direction of the largest eigenvector. Similarly, eccentricity is computed as the ratio of the smallest eigenvalue to the largest eigenvalue [15]. The algebraic moment invariants are computed from the first m central moments and are given as the eigenvalues of predefined matrices, $M_{[j,k]}$, whose elements are scaled factors of the central moments [16]. Using the notation of [16] we use moments up to degree 8 and the eigenvalues of the matrices $M_{[2,2]}, M_{[2,3]} \times M_{[3,2]}, M_{[3,3]}, M_{[3,4]} \times M_{[4,3]}, M_{[4,4]}, M_{[4,5]} \times M_{[5,4]}$ for a total of 18 features invariant to affine transformations. *Sketch features:* We implemented the image retrieval method described in [17, 18] that allows images to be retrieved based on a rough user sketch. The feature needed to support this retrieval consists of a reduced resolution edge map of each image. To compute these edge maps, we (1) convert each color image to a single band luminance; (2) compute the binary edge image using a Canny edge operator; and (3) reduce the edge image to size 64 x 64. To do the reduction, for w the width of the image in pixels and h its height, we partition the image into blocks of size $w/64 \times h/64$, and if any pixel in a partition of the full size edge image is an edge pixel, the corresponding pixel in the reduced edge map is set to an edge pixel. Finally, we thin this reduced image. This gives the reduced edge map or "image abstraction" ([17]) on which the retrieval by sketch is performed. IBM0002390 at 2393-95:

		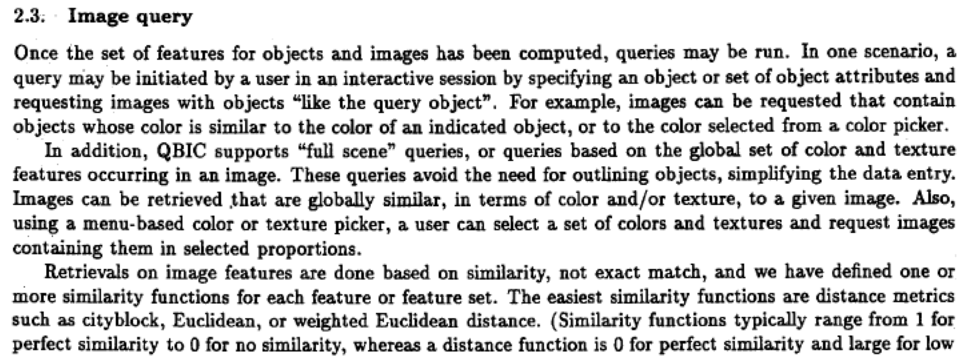 **2.3. Image query** Once the set of features for objects and images has been computed, queries may be run. In one scenario, a query may be initiated by a user in an interactive session by specifying an object or set of object attributes and requesting images with objects "like the query object". For example, images can be requested that contain objects whose color is similar to the color of an indicated object, or to the color selected from a color picker. In addition, QBIC supports "full scene" queries, or queries based on the global set of color and texture features occurring in an image. These queries avoid the need for outlining objects, simplifying the data entry. Images can be retrieved that are globally similar, in terms of color and/or texture, to a given image. Also, using a menu-based color or texture picker, a user can select a set of colors and textures and request images containing them in selected proportions. Retrievals on image features are done based on similarity, not exact match, and we have defined one or more similarity functions for each feature or feature set. The easiest similarity functions are distance metrics such as cityblock, Euclidean, or weighted Euclidean distance. (Similarity functions typically range from 1 for perfect similarity to 0 for no similarity, whereas a distance function is 0 for perfect similarity and large for low

similarity, so we must convert as necessary to obtain consistent measures.) Different similarity measures may be used, and they may vary from application to application.

The matching methods and similarity functions for the different feature types are described below:

Color: For average color, the distance between a query object and database object is weighted Euclidean distance. The weights are user adjustable so that we may, for example, downweight the L component of an (L, a, b) query. Best results, as judged by our subjective evaluation in on-going experiments, are obtained by representing the average color in the MTM color space, and by inversely weighting each component by its standard deviation over the samples in the database.

Improved results are obtained when color matching is done not on average color, but on the distribution of colors that occur in an object or image. In this way, a red and blue object matches other red and blue objects better than to a purple object. To match color histograms, we use a method similar to the one described in [8]. Let X be the query histogram and Y the histogram of an item in the database, both normalized. We compute the element by element difference histogram Z. Then the similarity between X and Y is given by $||Z|| = Z^T AZ$, where A is a symmetric color similarity matrix with

$$a(i,j) \equiv 1 - d(c_i, c_j)/d_{max}$$

where c_i and c_j are the ith and jth colors in the color histograms, and $d(c_i, c_j)$ is the MTM color distance, and d_{max} is the maximum distance between any two colors ([8] used Luv distance.) This metric gives the (weighted) length of the difference vector between X and Y, weighted by A, which accounts for the perceptual distance between different pairs of colors. The result is a measure that accounts for both the difference between the amounts in any given color (e.g. a particular shade of red), as well as differences between similar colors (e.g. red and orange).

Based on our experience, we also found it useful for the user to be able to request images with x% of color 1, y% of color 2, etc. where both the percentages and the colors are explicitly specified by the user. This allows searches such as for a beach scene (25% white (for the sand) and 50% blue (for the sky and water)), and for many scenes for which the user has seen before, and for which he/she can remember approximate color percentages. The color picker that allows this type of query is shown in Figure 2.

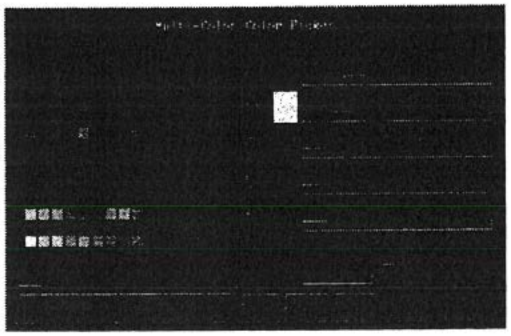

Figure 2: Multi-Color Color Picker with two selected colors. The palette on the left displays the set of selectable colors. Once a color is selected, it appears in a box on the right, and the amount of the color can be specified by the sliders. In this case, 25% light blue and 15% white have been selected.

Texture: Texture distance is computed as weighted Euclidean distance in the three dimensional texture space. The most common weighting factors are the inverse variances for each component, computed over the

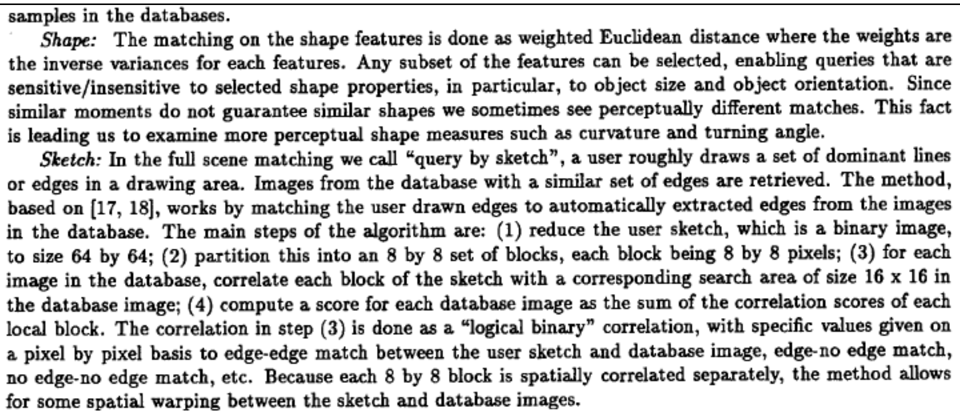

samples in the databases.

Shape: The matching on the shape features is done as weighted Euclidean distance where the weights are the inverse variances for each features. Any subset of the features can be selected, enabling queries that are sensitive/insensitive to selected shape properties, in particular, to object size and object orientation. Since similar moments do not guarantee similar shapes we sometimes see perceptually different matches. This fact is leading us to examine more perceptual shape measures such as curvature and turning angle.

Sketch: In the full scene matching we call "query by sketch", a user roughly draws a set of dominant lines or edges in a drawing area. Images from the database with a similar set of edges are retrieved. The method, based on [17, 18], works by matching the user drawn edges to automatically extracted edges from the images in the database. The main steps of the algorithm are: (1) reduce the user sketch, which is a binary image, to size 64 by 64; (2) partition this into an 8 by 8 set of blocks, each block being 8 by 8 pixels; (3) for each image in the database, correlate each block of the sketch with a corresponding search area of size 16 x 16 in the database image; (4) compute a score for each database image as the sum of the correlation scores of each local block. The correlation in step (3) is done as a "logical binary" correlation, with specific values given on a pixel by pixel basis to edge-edge match between the user sketch and database image, edge-no edge match, no edge-no edge match, etc. Because each 8 by 8 block is spatially correlated separately, the method allows for some spatial warping between the sketch and database images.

IBM0002390 at 2395

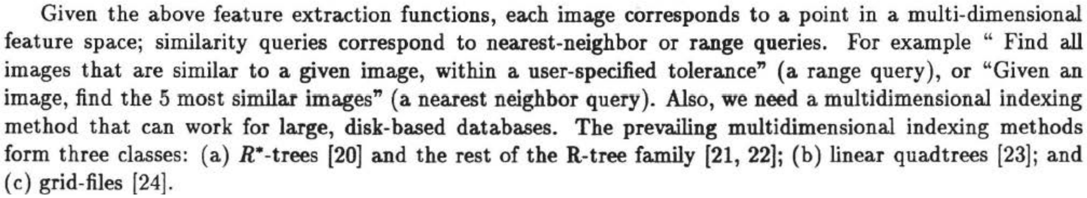

Given the above feature extraction functions, each image corresponds to a point in a multi-dimensional feature space; similarity queries correspond to nearest-neighbor or range queries. For example " Find all images that are similar to a given image, within a user-specified tolerance" (a range query), or "Given an image, find the 5 most similar images" (a nearest neighbor query). Also, we need a multidimensional indexing method that can work for large, disk-based databases. The prevailing multidimensional indexing methods form three classes: (a) R^*-trees [20] and the rest of the R-tree family [21, 22]; (b) linear quadtrees [23]; and (c) grid-files [24].

IBM 0002390 at 2396-97

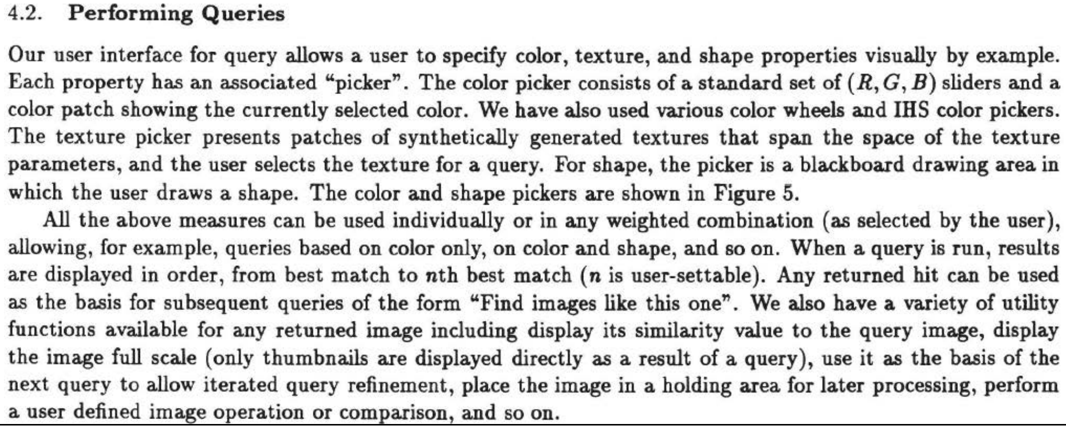

4.2. Performing Queries

Our user interface for query allows a user to specify color, texture, and shape properties visually by example. Each property has an associated "picker". The color picker consists of a standard set of (R, G, B) sliders and a color patch showing the currently selected color. We have also used various color wheels and IHS color pickers. The texture picker presents patches of synthetically generated textures that span the space of the texture parameters, and the user selects the texture for a query. For shape, the picker is a blackboard drawing area in which the user draws a shape. The color and shape pickers are shown in Figure 5.

All the above measures can be used individually or in any weighted combination (as selected by the user), allowing, for example, queries based on color only, on color and shape, and so on. When a query is run, results are displayed in order, from best match to nth best match (n is user-settable). Any returned hit can be used as the basis for subsequent queries of the form "Find images like this one". We also have a variety of utility functions available for any returned image including display its similarity value to the query image, display the image full scale (only thumbnails are displayed directly as a result of a query), use it as the basis of the next query to allow iterated query refinement, place the image in a holding area for later processing, perform a user defined image operation or comparison, and so on.

5. Sample Query Results

Sample query results are shown in the following figures. The reproductions are in black and white and do not show the color component which, for these images, is an important component. Figures 6 through 11 and 12 show a color query, a texture query, a combined color and texture query, a shape query, and a query by sketch. All queries are done on a database of about 1000 objects with a wide diversity of photo clip art subject matter (buildings, people, landscapes, animals, etc.). Formal experiments to access the recall and precision rate of QBIC queries are in progress.

(buildings, people, landscapes, animals, etc.). Formal experiments to access the recall and precision rate of QBIC queries are in progress.

The QBIC Project: Querying Images by Content Using Color, Texture, and Shape (IBM 0002390–404) at 2400, Fig. 4

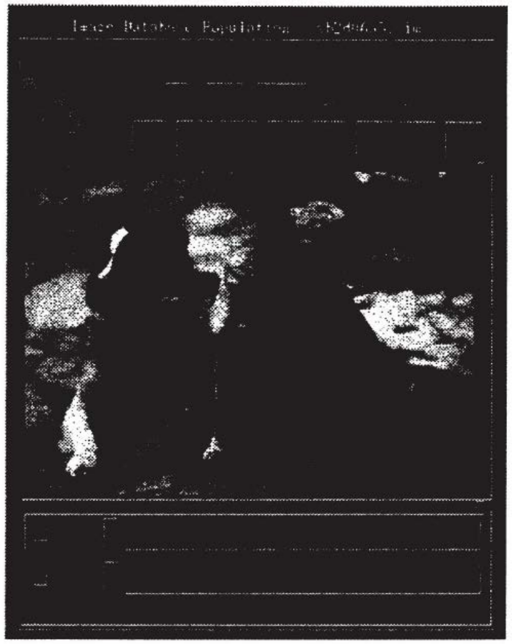

Figure 4: Object Identification Tool. Note for example, the right dog's ear, which has been outlined.

The QBIC Project: Querying Images by Content Using Color, Texture, and Shape (IBM 0002390–404) at Figs. 7–9

		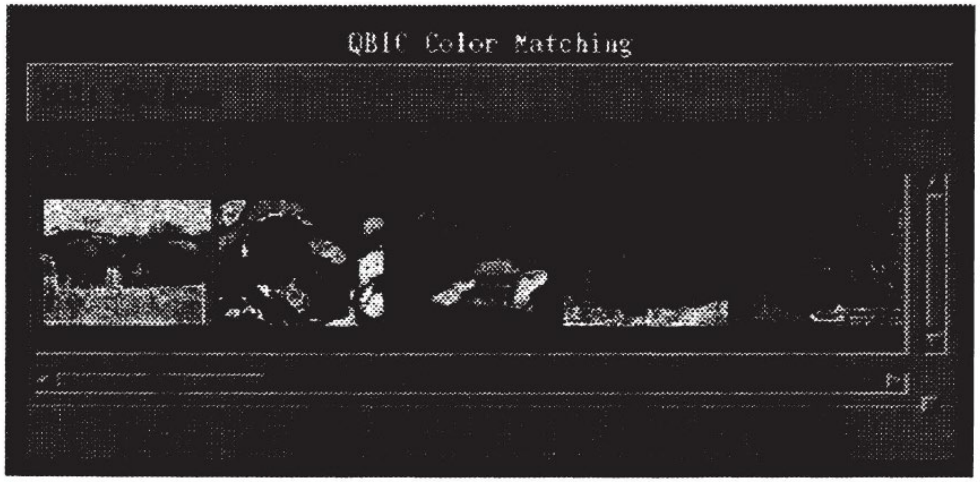 Figure 7: Result of query by example using color only.

		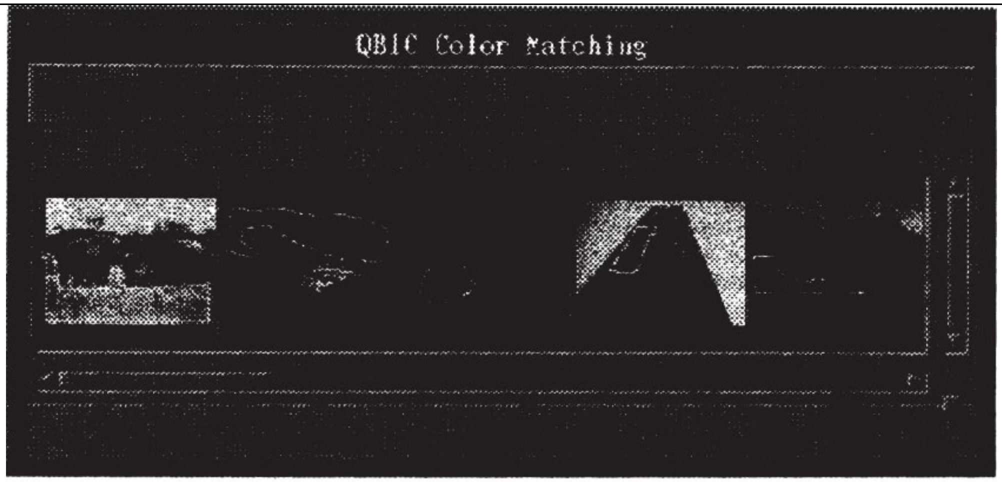 Figure 8: Result of query by example using texture only. 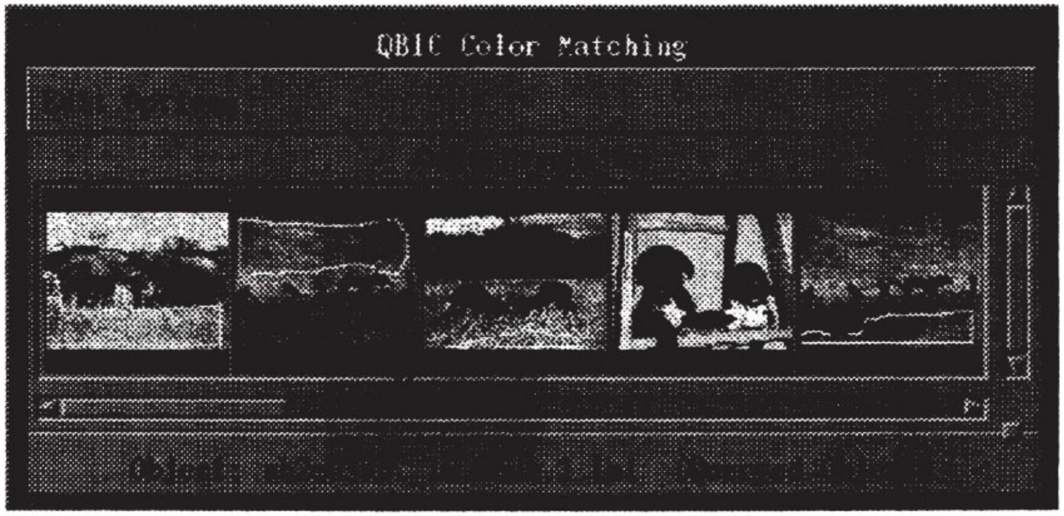 Figure 9: Result of query by example using color and texture.

		Image, Audio, and Video Extenders Administration and Programming (IBM 000001–617): "Your applications can even search for images by content. Imagine an application that uses visual examples to search for images. With such an application, users could select an example image and have the application find other images that have colors or textures similar to those in the example. With DB2 extenders' Query by Image Content (QBIC) capability, you can create applications that search for images in this visual way." *Id*. at p. 4.

		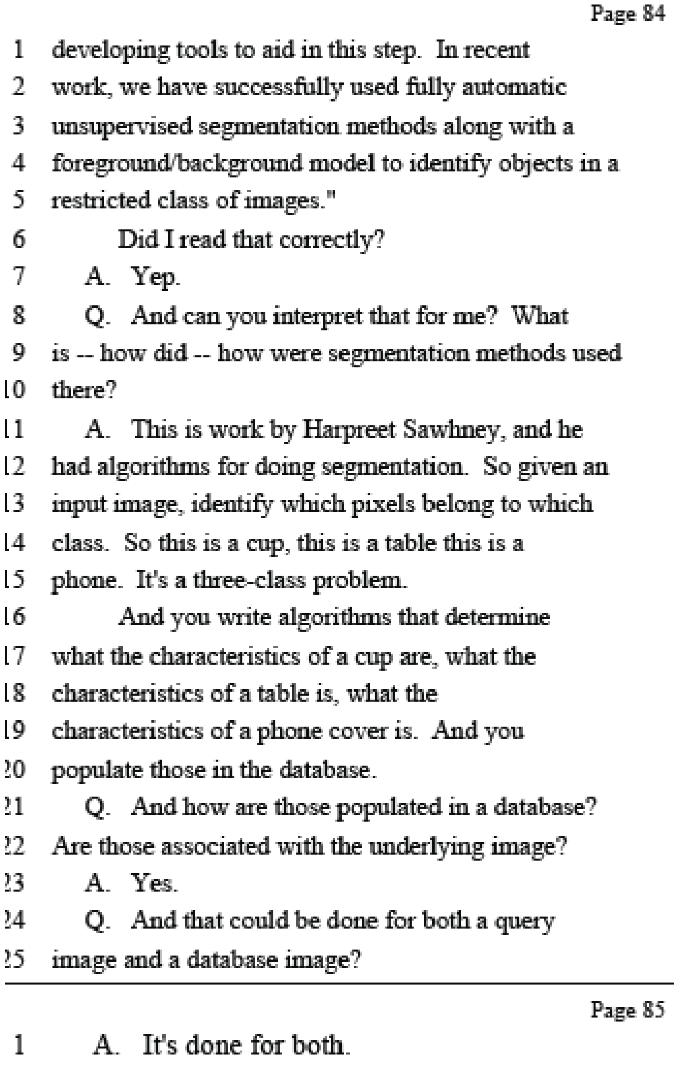 Page 84 developing tools to aid in this step. In recent work, we have successfully used fully automatic unsupervised segmentation methods along with a foreground/background model to identify objects in a restricted class of images." Did I read that correctly? A. Yep. Q. And can you interpret that for me? What is -- how did -- how were segmentation methods used there? A. This is work by Harpreet Sawhney, and he had algorithms for doing segmentation. So given an input image, identify which pixels belong to which class. So this is a cup, this is a table this is a phone. It's a three-class problem. And you write algorithms that determine what the characteristics of a cup are, what the characteristics of a table is, what the characteristics of a phone cover is. And you populate those in the database. Q. And how are those populated in a database? Are those associated with the underlying image? A. Yes. Q. And that could be done for both a query image and a database image? Page 85 A. It's done for both.

		Id. at 305: 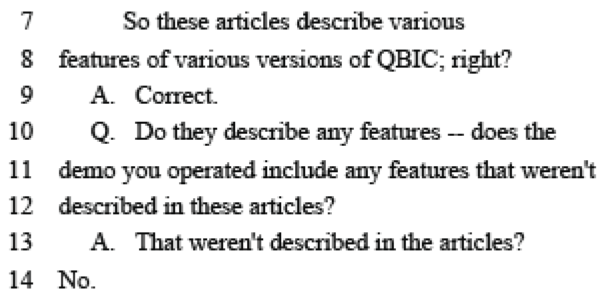 So these articles describe various features of various versions of QBIC; right? A. Correct. Q. Do they describe any features -- does the demo you operated include any features that weren't described in these articles? A. That weren't described in the articles? No. *Id.* at 35: 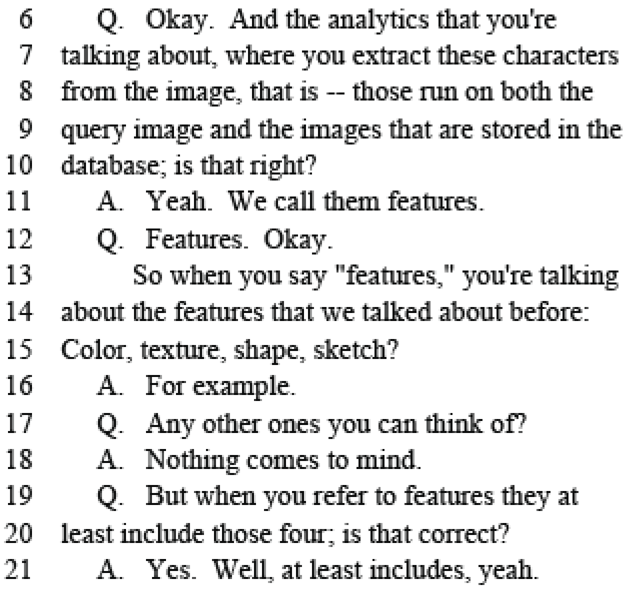 Q. Okay. And the analytics that you're talking about, where you extract these characters from the image, that is -- those run on both the query image and the images that are stored in the database; is that right? A. Yeah. We call them features. Q. Features. Okay. So when you say "features," you're talking about the features that we talked about before: Color, texture, shape, sketch? A. For example. Q. Any other ones you can think of? A. Nothing comes to mind. Q. But when you refer to features they at least include those four; is that correct? A. Yes. Well, at least includes, yeah.

		Id. at 109-110: 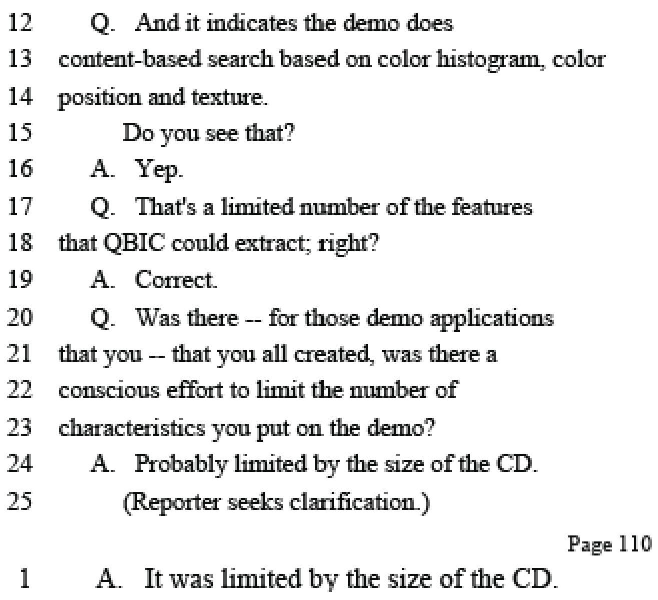 Q. And it indicates the demo does content-based search based on color histogram, color position and texture. Do you see that? A. Yep. Q. That's a limited number of the features that QBIC could extract; right? A. Correct. Q. Was there -- for those demo applications that you -- that you all created, was there a conscious effort to limit the number of characteristics you put on the demo? A. Probably limited by the size of the CD. (Reporter seeks clarification.) Page 110 A. It was limited by the size of the CD. *Id.* at 311: 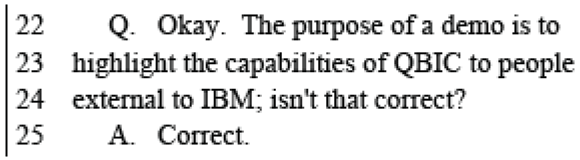 Q. Okay. The purpose of a demo is to highlight the capabilities of QBIC to people external to IBM; isn't that correct? A. Correct. *Id.* at 313-317:

		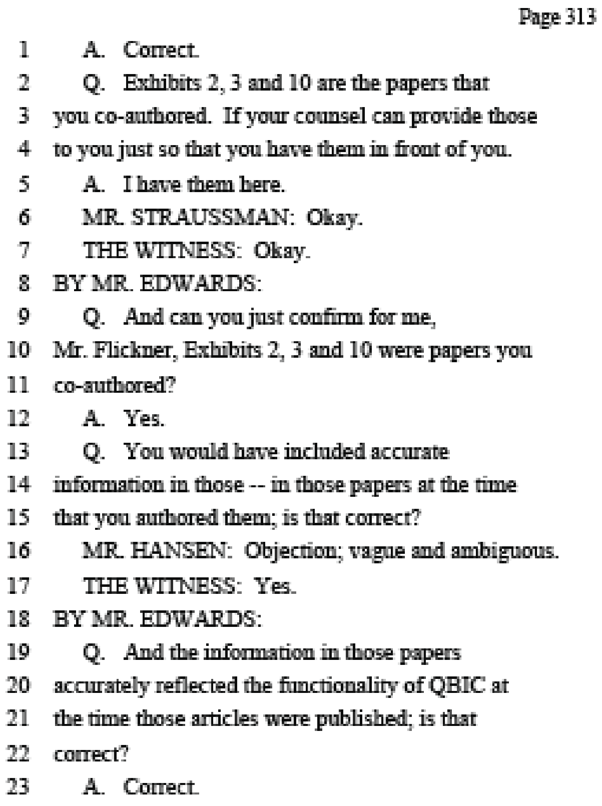Page 313 A. Correct. Q. Exhibits 2, 3 and 10 are the papers that you co-authored. If your counsel can provide those to you just so that you have them in front of you. A. I have them here. MR. STRAUSSMAN: Okay. THE WITNESS: Okay. BY MR. EDWARDS: Q. And can you just confirm for me, Mr. Flickner, Exhibits 2, 3 and 10 were papers you co-authored? A. Yes. Q. You would have included accurate information in those -- in those papers at the time that you authored them; is that correct? MR. HANSEN: Objection; vague and ambiguous. THE WITNESS: Yes. BY MR. EDWARDS: Q. And the information in those papers accurately reflected the functionality of QBIC at the time those articles were published; is that correct? A. Correct.

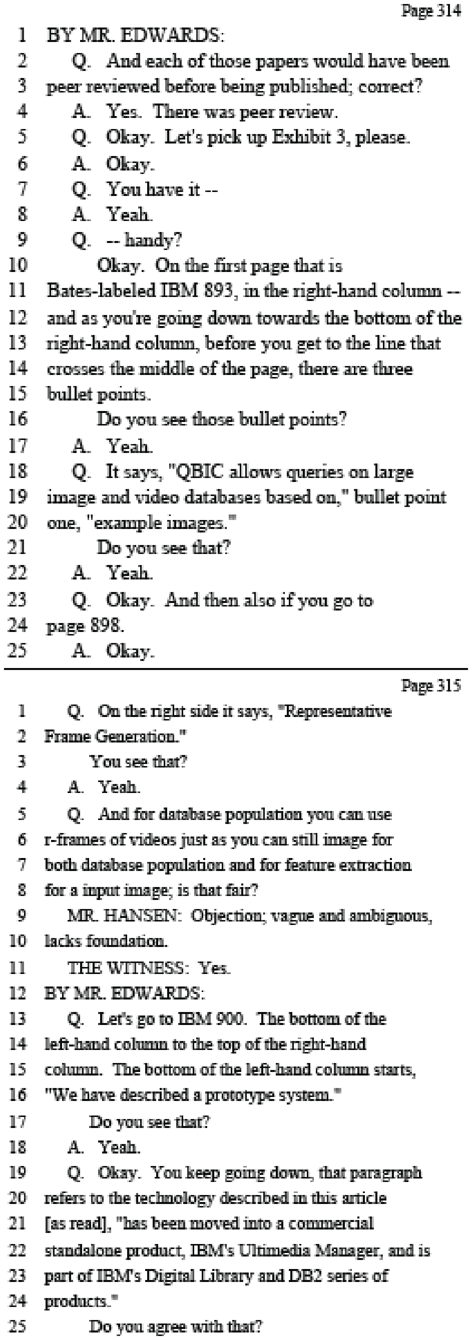

Page 314

BY MR. EDWARDS:

Q. And each of those papers would have been peer reviewed before being published; correct?

A. Yes. There was peer review.

Q. Okay. Let's pick up Exhibit 3, please.

A. Okay.

Q. You have it --

A. Yeah.

Q. -- handy?

Okay. On the first page that is Bates-labeled IBM 893, in the right-hand column -- and as you're going down towards the bottom of the right-hand column, before you get to the line that crosses the middle of the page, there are three bullet points.

Do you see those bullet points?

A. Yeah.

Q. It says, "QBIC allows queries on large image and video databases based on," bullet point one, "example images."

Do you see that?

A. Yeah.

Q. Okay. And then also if you go to page 898.

A. Okay.

Page 315

Q. On the right side it says, "Representative Frame Generation."

You see that?

A. Yeah.

Q. And for database population you can use r-frames of videos just as you can still image for both database population and for feature extraction for a input image; is that fair?

MR. HANSEN: Objection; vague and ambiguous, lacks foundation.

THE WITNESS: Yes.

BY MR. EDWARDS:

Q. Let's go to IBM 900. The bottom of the left-hand column to the top of the right-hand column. The bottom of the left-hand column starts, "We have described a prototype system."

Do you see that?

A. Yeah.

Q. Okay. You keep going down, that paragraph refers to the technology described in this article [as read], "has been moved into a commercial standalone product, IBM's Ultimedia Manager, and is part of IBM's Digital Library and DB2 series of products."

Do you agree with that?

		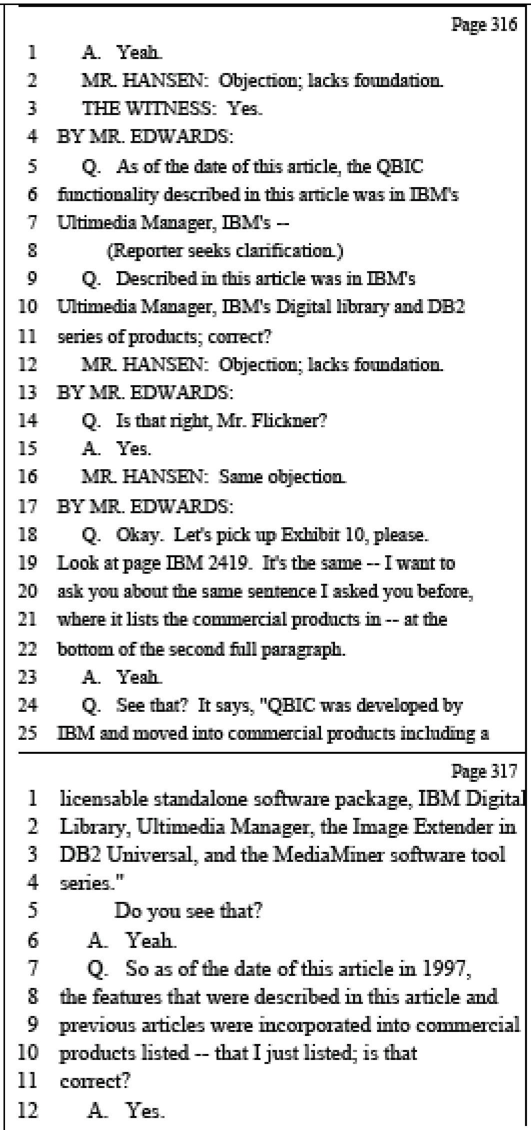 Page 316 A. Yeah. MR. HANSEN: Objection; lacks foundation. THE WITNESS: Yes. BY MR. EDWARDS: Q. As of the date of this article, the QBIC functionality described in this article was in IBM's Ultimedia Manager, IBM's -- (Reporter seeks clarification.) Q. Described in this article was in IBM's Ultimedia Manager, IBM's Digital library and DB2 series of products; correct? MR. HANSEN: Objection; lacks foundation. BY MR. EDWARDS: Q. Is that right, Mr. Flickner? A. Yes. MR. HANSEN: Same objection. BY MR. EDWARDS: Q. Okay. Let's pick up Exhibit 10, please. Look at page IBM 2419. It's the same -- I want to ask you about the same sentence I asked you before, where it lists the commercial products in -- at the bottom of the second full paragraph. A. Yeah. Q. See that? It says, "QBIC was developed by IBM and moved into commercial products including a Page 317 licensable standalone software package, IBM Digital Library, Ultimedia Manager, the Image Extender in DB2 Universal, and the MediaMiner software tool series." Do you see that? A. Yeah. Q. So as of the date of this article in 1997, the features that were described in this article and previous articles were incorporated into commercial products listed -- that I just listed; is that correct? A. Yes.

		Image, Audio, and Video Extenders Administration and Programming (IBM 000001–617): "Your applications can even search for images by content. Imagine an application that uses visual examples to search for images. With such an application, users could select an example image and have the application find other images that have colors or textures similar to those in the example. With DB2 extenders' Query by Image Content (QBIC) capability, you can create applications that search for images in this visual way." *Id*. at p. 4. " A QBIC catalog is a set of files that hold data about the visual features of images. The Image Extender uses this data to search for images by content. You create a QBIC catalog for each column of images in a user table that you want to make available for searching by content. When you create a QBIC catalog you identify the features for which you want the Image Extender to analyze, store, and later query data. You can also add or drop features from a QBIC catalog after the catalog is created. A QBIC catalog can hold data for the following image features: **Average color** The sum of the color values for all pixels in an image divided by the number of pixels in the image. (A pixel is the smallest element of an image that can be assigned color and intensity.) For example, if 50% of an image consists of blue pixels and the other 50% red pixels, the image has an average color value of purple. Average color is used to search for images that have a predominant color. If an image has a predominant color, the average color will be similar to the predominant color. **Histogram color** Measures the distribution of colors in an image against a spectrum of 64 colors. For each of the 64 colors, histogram color identifies the percentage of pixels in an image that have that color. For example, the histogram color of an image might be 40% white pixels, 50% blue, and 10% red; none of the pixels in the image have any of the remaining colors in the histogram spectrum. Histogram color is used to search for images that have a variety of colors. **Positional color**

		The average color value for the pixels in a specified area in an image. For example, the upper right-hand corner of an image might show a bright yellow sun; the positional color of this area of the image is bright yellow. Positional color is used to search for images that have a predominant color in a particular area. **Texture** Measures the coarseness, contrast, and directionality of an image. Coarseness indicates the size of repeating items in an image (for example, pebbles versus boulders). Contrast identifies the brightness variations in an image (light versus dark). Directionality indicates whether a direction predominates in an image (as in the vertical direction of a picket fence) or does not predominate (as in an image of sand). Texture is used to search for images that have a particular pattern. To make an image available for searching by content, you catalog the image. When you catalog an image, the Image Extender analyzes the image, by computing the feature values for the image, and stores the values in a QBIC catalog. When you search for an image by content, your query identifies one or more features for the search (such as average color), a source for each feature (such as an example image), and a target set of cataloged images. The Image Extender computes the feature value of the source and compares it to the cataloged feature values for the target images. It then computes a score that indicates how similar the feature values of the target images are to the source. You can have the Image Extender return the images whose features are most similar to the source. The Image Extender will return the handle of each image and the image score. You can also have the Image Extender return only the score of a single image." *Id.* at pp. 20–21. QBIC Demo (IBM000747):

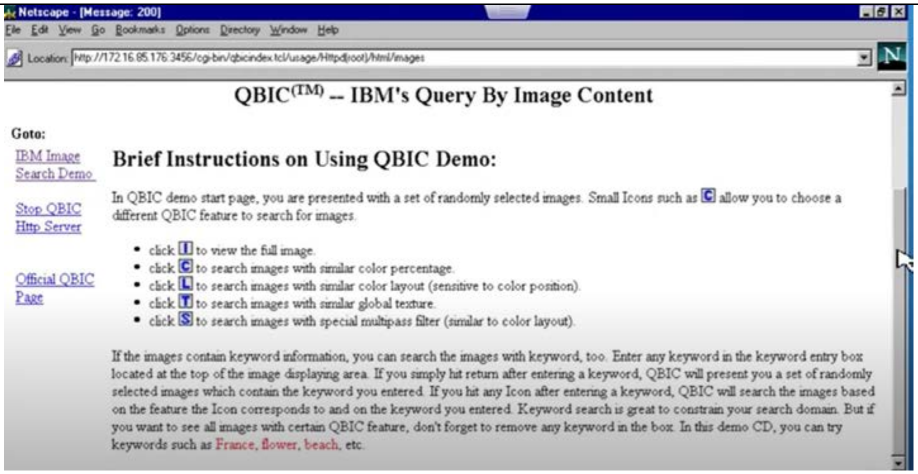

QBIC(TM) -- IBM's Query By Image Content

Brief Instructions on Using QBIC Demo:

In QBIC demo start page, you are presented with a set of randomly selected images. Small Icons such as C allow you to choose a different QBIC feature to search for images.

- click I to view the full image.
- click C to search images with similar color percentage.
- click L to search images with similar color layout (sensitive to color position).
- click T to search images with similar global texture.
- click S to search images with special multipass filter (similar to color layout).

If the images contain keyword information, you can search the images with keyword, too. Enter any keyword in the keyword entry box located at the top of the image displaying area. If you simply hit return after entering a keyword, QBIC will present you a set of randomly selected images which contain the keyword you entered. If you hit any Icon after entering a keyword, QBIC will search the images based on the feature the Icon corresponds to and on the keyword you entered. Keyword search is great to constrain your search domain. But if you want to see all images with certain QBIC feature, don't forget to remove any keyword in the box. In this demo CD, you can try keywords such as France, flower, beach, etc.

QBIC Demo (IBM000747):

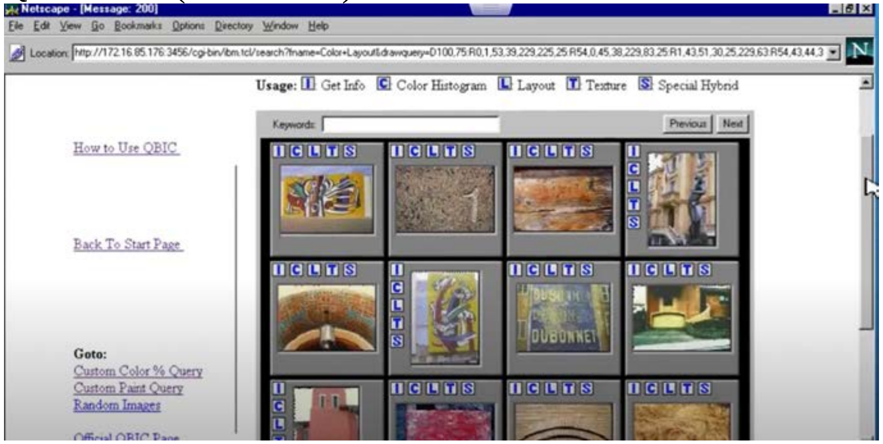

QBIC Demo (IBM000747):

		 Location: http://172.16.85.176:3456/cgi-bin/ibm.tcl/histpicker Keywords: Red: 59 Green: 131 Blue: 59 9% (255,255,255) 8% (59,131,59) 4% (106,182,106) To specify your color histogram query, "paint" a graphical representation of your pattern of colors and approximate percentages that you want. First select a *percentage* by clicking in the painting area to block off a rectangle of the desired size (percentage of the full area). All new regions are painted white, initially. Then select a color from the QBIC Demo (IBM000747): 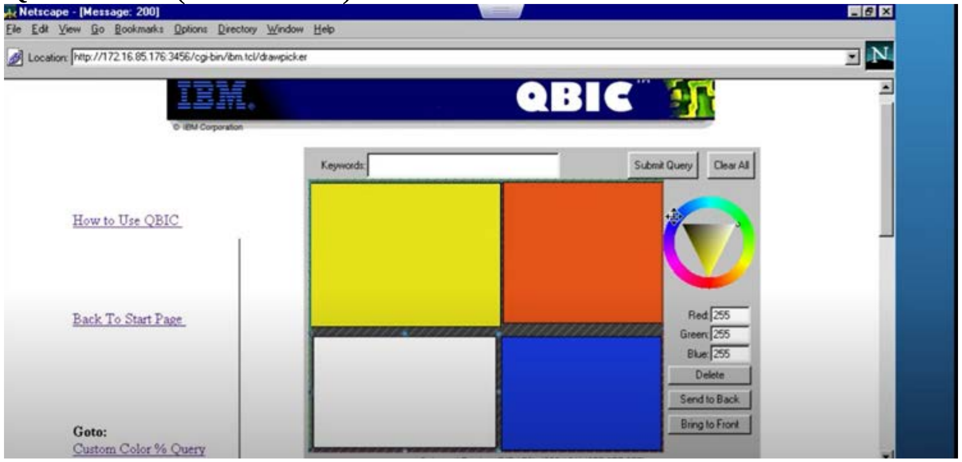 Location: http://172.16.85.176:3456/cgi-bin/ibm.tcl/drawpicker Keywords: Red: 255 Green: 255 Blue: 255 QBIC Demo (IBM000747):

Location: http://172.16.85.176:3456/cgi-bin/ibm.tcl/search?fname=HybridColor&france10%2Ejpg.x=0

Query was:

Example:	=france10.jpg
Query Type:	QbHybridColorFeatureClass

QBIC Demo (IBM000747):

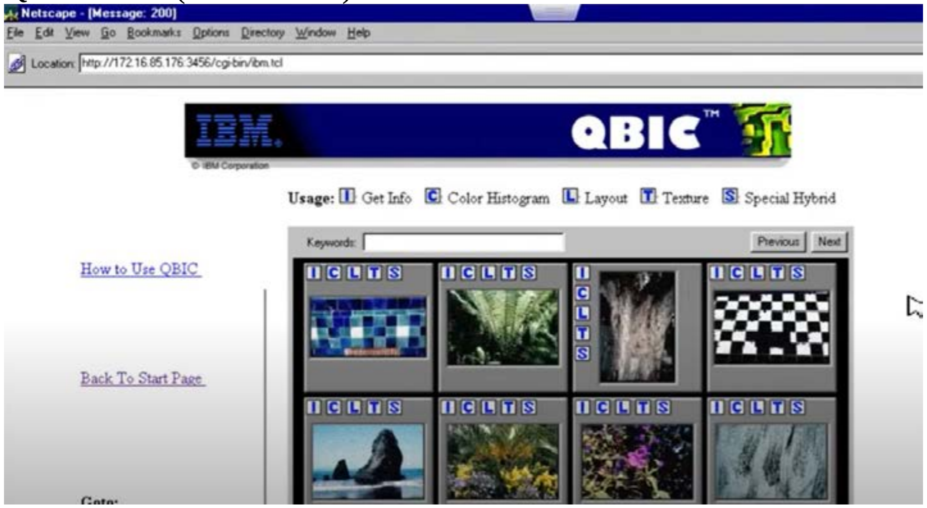

Location: http://172.16.85.176:3456/cgi-bin/ibm.tcl

QBIC Demo (IBM000747):

		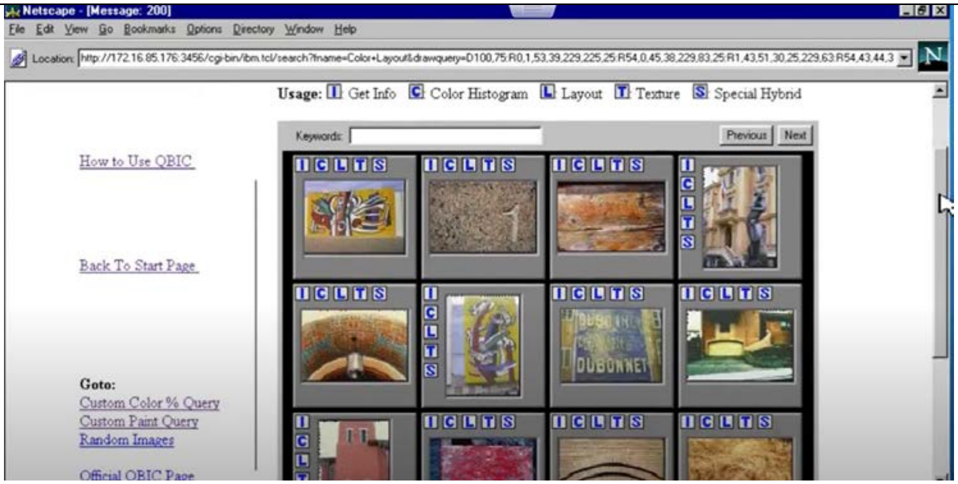 Location: http://172.16.85.176:3456/cgi-bin/ibm.tcl/search?fname=Color+Layout&drawquery=D100,75:R0,1,53,39,229,225,25:R54,0,45,38,229,83,25:R1,43,51,30,25,229,63:R54,43,44,3 QBIC Demo (IBM000747):  Location: http://172.16.85.176:3456/cgi-bin/ibm.tcl/info?france50%2Ejpg

QBIC (Query by Image Content) at 1

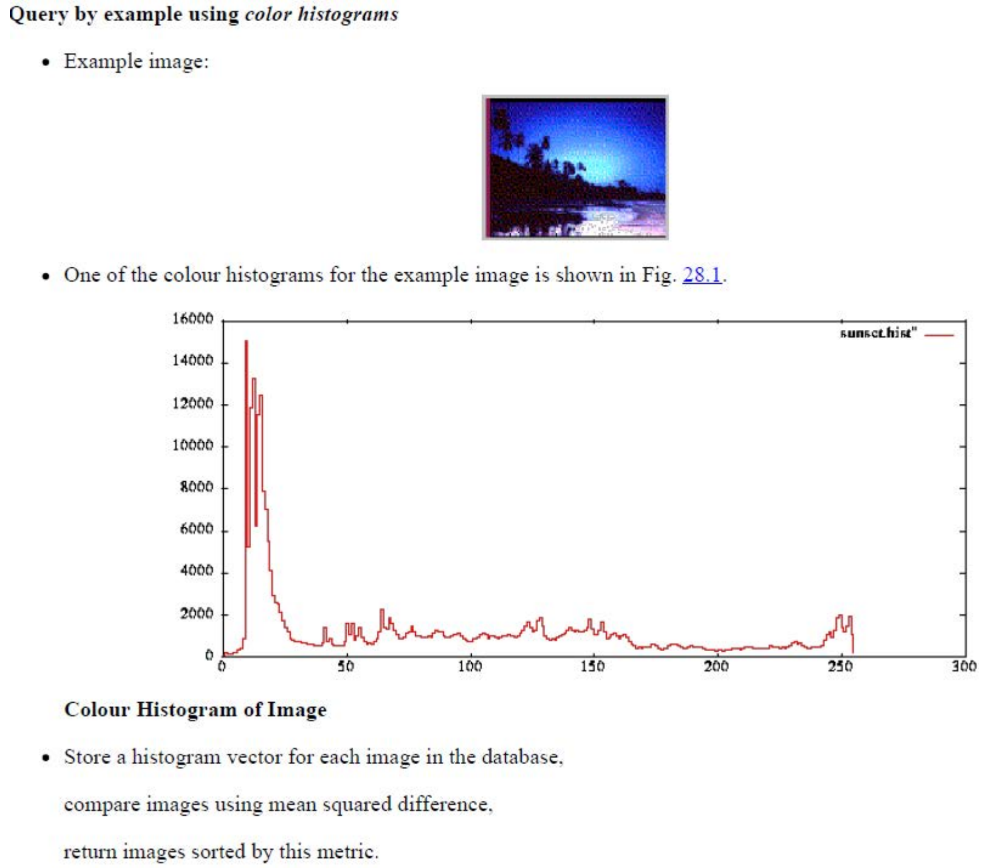

Query by example using *color histograms*

- Example image:

- One of the colour histograms for the example image is shown in Fig. 28.1.

Colour Histogram of Image

- Store a histogram vector for each image in the database,

  compare images using mean squared difference,

  return images sorted by this metric.

IBM000001 at 39 – 41

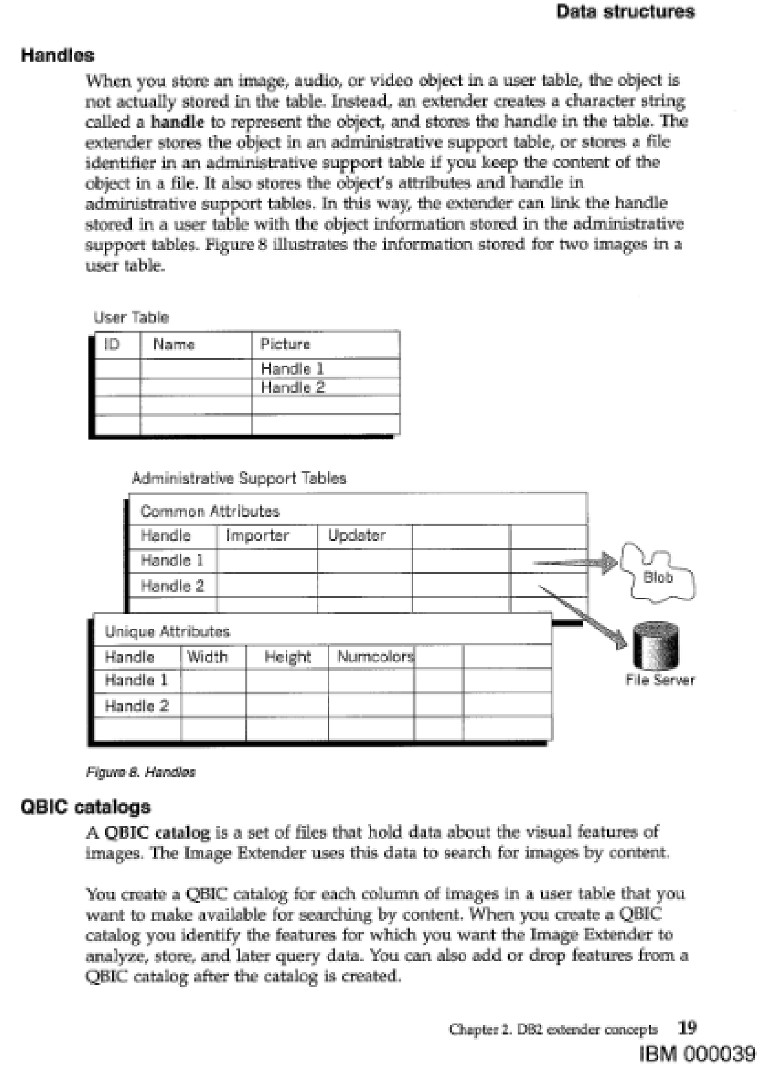

Handles

When you store an image, audio, or video object in a user table, the object is not actually stored in the table. Instead, an extender creates a character string called a **handle** to represent the object, and stores the handle in the table. The extender stores the object in an administrative support table, or stores a file identifier in an administrative support table if you keep the content of the object in a file. It also stores the object's attributes and handle in administrative support tables. In this way, the extender can link the handle stored in a user table with the object information stored in the administrative support tables. Figure 8 illustrates the information stored for two images in a user table.

Figure 8. Handles

QBIC catalogs

A **QBIC catalog** is a set of files that hold data about the visual features of images. The Image Extender uses this data to search for images by content.

You create a QBIC catalog for each column of images in a user table that you want to make available for searching by content. When you create a QBIC catalog you identify the features for which you want the Image Extender to analyze, store, and later query data. You can also add or drop features from a QBIC catalog after the catalog is created.

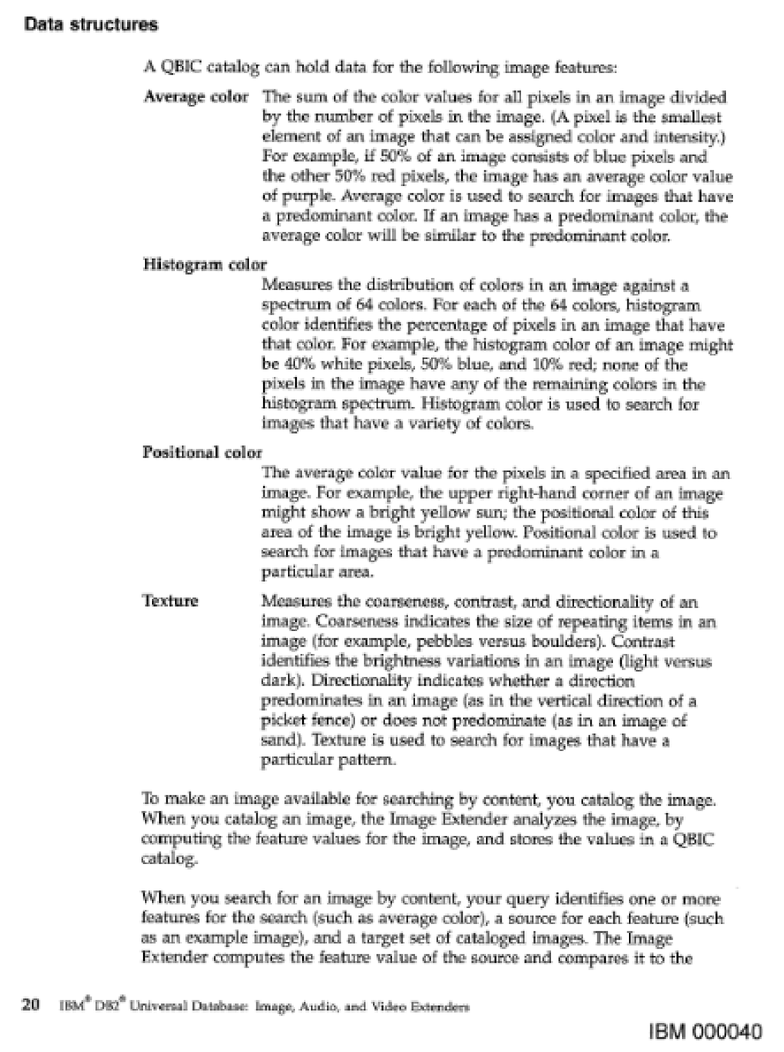

Data structures

A QBIC catalog can hold data for the following image features:

Average color	The sum of the color values for all pixels in an image divided by the number of pixels in the image. (A pixel is the smallest element of an image that can be assigned color and intensity.) For example, if 50% of an image consists of blue pixels and the other 50% red pixels, the image has an average color value of purple. Average color is used to search for images that have a predominant color. If an image has a predominant color, the average color will be similar to the predominant color.
Histogram color	Measures the distribution of colors in an image against a spectrum of 64 colors. For each of the 64 colors, histogram color identifies the percentage of pixels in an image that have that color. For example, the histogram color of an image might be 40% white pixels, 50% blue, and 10% red; none of the pixels in the image have any of the remaining colors in the histogram spectrum. Histogram color is used to search for images that have a variety of colors.
Positional color	The average color value for the pixels in a specified area in an image. For example, the upper right-hand corner of an image might show a bright yellow sun; the positional color of this area of the image is bright yellow. Positional color is used to search for images that have a predominant color in a particular area.
Texture	Measures the coarseness, contrast, and directionality of an image. Coarseness indicates the size of repeating items in an image (for example, pebbles versus boulders). Contrast identifies the brightness variations in an image (light versus dark). Directionality indicates whether a direction predominates in an image (as in the vertical direction of a picket fence) or does not predominate (as in an image of sand). Texture is used to search for images that have a particular pattern.

To make an image available for searching by content, you catalog the image. When you catalog an image, the Image Extender analyzes the image, by computing the feature values for the image, and stores the values in a QBIC catalog.

When you search for an image by content, your query identifies one or more features for the search (such as average color), a source for each feature (such as an example image), and a target set of cataloged images. The Image Extender computes the feature value of the source and compares it to the

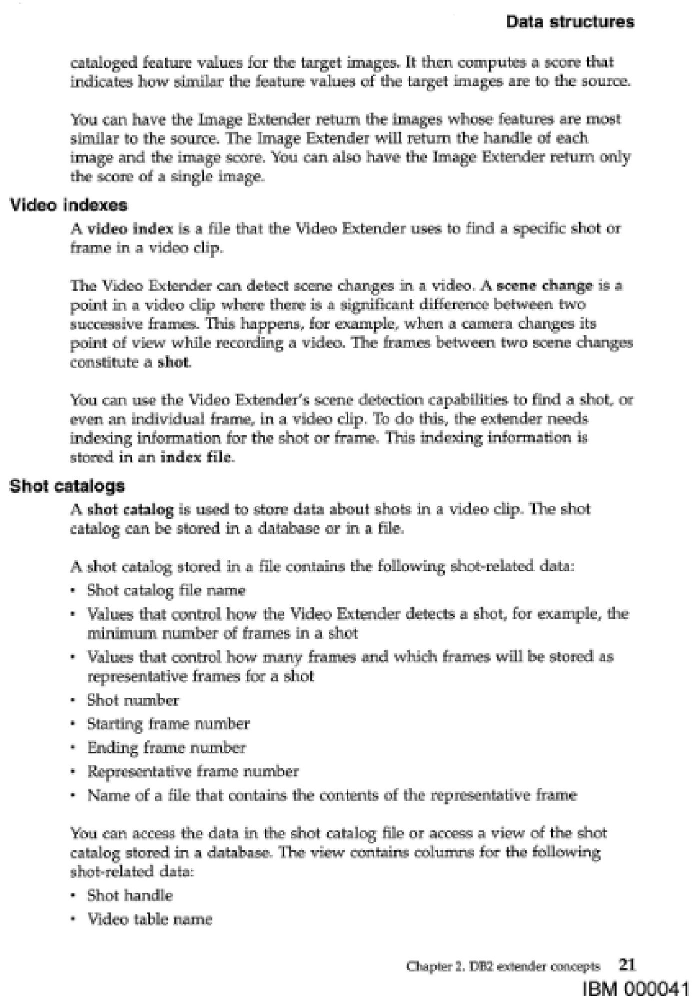

cataloged feature values for the target images. It then computes a score that indicates how similar the feature values of the target images are to the source.

You can have the Image Extender return the images whose features are most similar to the source. The Image Extender will return the handle of each image and the image score. You can also have the Image Extender return only the score of a single image.

Video indexes

A **video index** is a file that the Video Extender uses to find a specific shot or frame in a video clip.

The Video Extender can detect scene changes in a video. A **scene change** is a point in a video clip where there is a significant difference between two successive frames. This happens, for example, when a camera changes its point of view while recording a video. The frames between two scene changes constitute a **shot**.

You can use the Video Extender's scene detection capabilities to find a shot, or even an individual frame, in a video clip. To do this, the extender needs indexing information for the shot or frame. This indexing information is stored in an **index file.**

Shot catalogs

A **shot catalog** is used to store data about shots in a video clip. The shot catalog can be stored in a database or in a file.

A shot catalog stored in a file contains the following shot-related data:

- Shot catalog file name
- Values that control how the Video Extender detects a shot, for example, the minimum number of frames in a shot
- Values that control how many frames and which frames will be stored as representative frames for a shot
- Shot number
- Starting frame number
- Ending frame number
- Representative frame number
- Name of a file that contains the contents of the representative frame

You can access the data in the shot catalog file or access a view of the shot catalog stored in a database. The view contains columns for the following shot-related data:

- Shot handle
- Video table name

IBM0002263 at 2264

Description

Ultimedia Manager 1.1 opens new dimensions in visual queries, helping you manage your image database more efficiently. With its easy-to-use interface, even entry-level users can perform sophisticated searches. You can search by image content, data description, or a combination of both. Once the images are found, you can view the images in full resolution, use the images in another application, or read the associated textual comments.

You can use Ultimedia Manager 1.1 either as a stand-alone product or in a LAN environment. As a stand-alone product, Ultimedia Manager 1.1 employs full image classification and image query capabilities. Within the LAN environment, the Ultimedia Manager image catalogs can be sorted on a DB2 for OS/2 or DB2 for AIX database server. The image catalogs can be accessed by Windows or OS/2 clients using the Client Search Feature or the Visualizer Ultimedia Query.

Ultimedia Manager runs as an OS/2 application to a DB2 for OS/2 or DB2 for AIX database server in a LAN, or as a stand-alone application with a local DB2 for OS/2 database. The Client Search Feature runs on DOS or Windows, and accesses files in dBASE flat file format. A database manager is not required to access catalogs in dBASE file format.

Image Classification

With the Image Classifier, you can create visual indexes based on color, shape, category, and layout elements, which are areas of the image containing the patterns you want to use for searches. You can store these elements in catalogs in DB2 for OS/2 databases. The catalogs can be used locally or, through the use of utility programs, be transferred to DB2 for OS/2 or AIX servers. In addition, you can convert the catalog data into dBASE flat file format for access by Windows client applications.

The Image Classifier and Analyzer allows automatic or interactive classification of images for later image queries, flexibility in storage of catalog information, and access to multiple users. You can organize large image catalogs so that later searches for images can be performed easily. You can classify automatically using a batch program, or interactively by specifying areas of interest in an image and indicating the visual cue (color, shape, or texture) to be saved for later queries.

Image Query

Once the image catalogs are created, you can use the Image Query component to find image and business information. Business information, such as character and numeric data, are stored in the database tables using normal table population methods. Image Query is provided in Ultimedia Manager and Client Search Feature.

Image Query allows searches by three methods:

- Image content, such as color, shape, texture, or layout
- Data description, such as business data stored in the database (character or numeric) or descriptive text (category)
- Combinations of image content and data description.

Image Content: To search by image content, specify the attributes (color, shape, texture, or layout). You can drag and drop visual samples into an image query, select a color from the color wheel, or sketch a shape.

Data Description: To search by data description, enter character and numeric data values in the database tables using the simple fill-in-the-blanks query panel. You can enter data such as name, number, date, or a range of dates. You can integrate character or numeric data as part of the image query, enabling a more effective search for better results. The simple fill-in-the-blanks method for specifying desired conditions is so easy even non-programmers can use database table information as part of their searches.

Combinations: Image queries using either image content or data description, or combinations of both, allow you to search through thousands of images more quickly and effectively than the traditional text and keyword search methods.

Image Formats

Ultimedia Manager 1.1 understands these image formats:

- Audio Visual Connection (AVC) images
- OS/2 bitmap
- Windows 3.x bitmap
- PCX
- Intel and Motorola TIFF
- TARGA TGA
- TIFF (FAX)
- GIF
- Kodak Photo-CD

Product Positioning

The advanced image search technology in Ultimedia Manager, query by image content, was developed in conjunction with the IBM Almaden Research Center. It provides function not available with conventional, text-only image management systems.

Ultimedia Manager fulfills image search and retrieval needs of business users and creative professionals through use of online image libraries or collections. It is intended for users who have large image libraries and need fast, flexible access to their image data, users who may also want to add character or numeric data to their image queries, and users who feel limited by text-based systems.

Ultimedia Manager 1.1 works with the DB2 Family of products and the Visualizer Ultimedia Query product. Catalogs created by Ultimedia Manager can be stored in DB2 for OS/2 database tables. Image catalogs also can be stored in DB2 for OS/2 or DB2 for AIX database servers in client/server environments.

OS/2 and Windows clients in a LAN environment using the Client Search Feature and the Visualizer Ultimedia Query can access catalogs created with Ultimedia Manager 1.1. Visualizer Ultimedia Query provides users with the capability to:

- Add multimedia columns (such as image, video, audio, and formatted text) to DB2 for OS/2 database tables
- Query and display images
- View or listen to video and audio information
- Display formatted text

Users of Ultimedia Manager 1.1 and Client Search Feature are focused on image management and query applications. Visualizer Ultimedia Query users are

295-042 -2-

IBM 0002264

IBM000777 at 794

		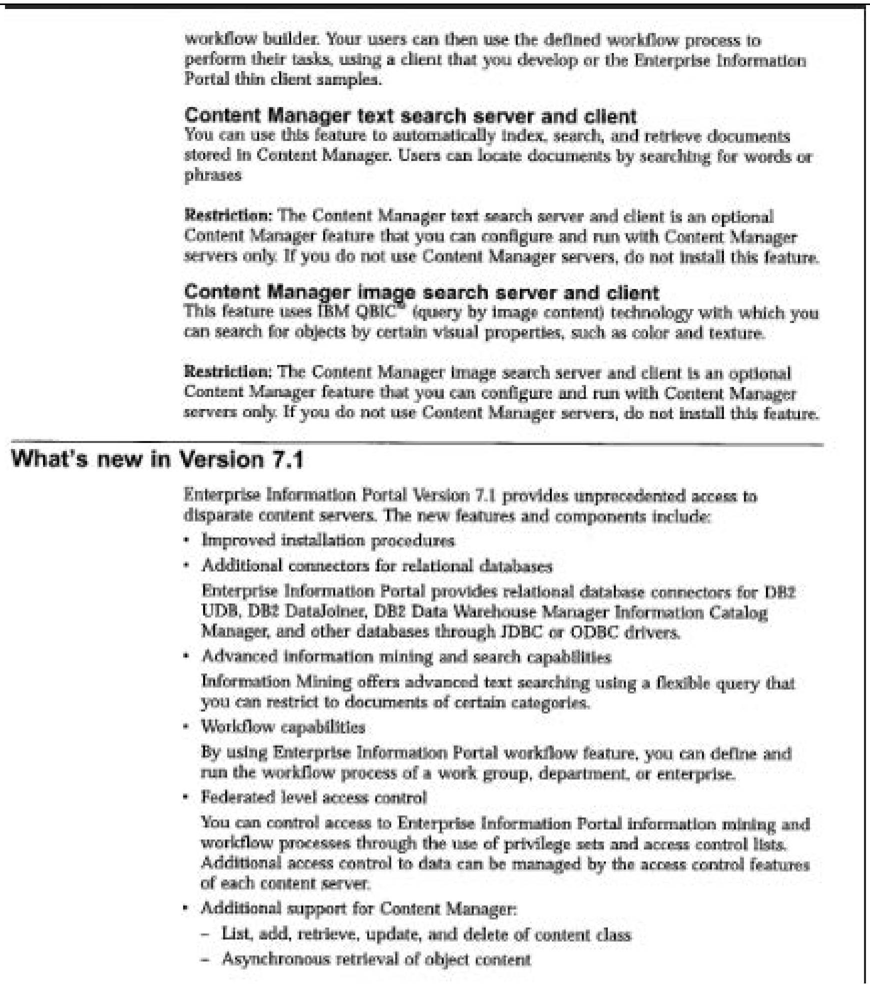 workflow builder. Your users can then use the defined workflow process to perform their tasks, using a client that you develop or the Enterprise Information Portal thin client samples. **Content Manager text search server and client** You can use this feature to automatically index, search, and retrieve documents stored in Content Manager. Users can locate documents by searching for words or phrases **Restriction:** The Content Manager text search server and client is an optional Content Manager feature that you can configure and run with Content Manager servers only. If you do not use Content Manager servers, do not install this feature. **Content Manager image search server and client** This feature uses IBM QBIC™ (query by image content) technology with which you can search for objects by certain visual properties, such as color and texture. **Restriction:** The Content Manager image search server and client is an optional Content Manager feature that you can configure and run with Content Manager servers only. If you do not use Content Manager servers, do not install this feature. **What's new in Version 7.1** Enterprise Information Portal Version 7.1 provides unprecedented access to disparate content servers. The new features and components include: • Improved installation procedures • Additional connectors for relational databases Enterprise Information Portal provides relational database connectors for DB2 UDB, DB2 DataJoiner, DB2 Data Warehouse Manager Information Catalog Manager, and other databases through JDBC or ODBC drivers. • Advanced information mining and search capabilities Information Mining offers advanced text searching using a flexible query that you can restrict to documents of certain categories. • Workflow capabilities By using Enterprise Information Portal workflow feature, you can define and run the workflow process of a work group, department, or enterprise. • Federated level access control You can control access to Enterprise Information Portal information mining and workflow processes through the use of privilege sets and access control lists. Additional access control to data can be managed by the access control features of each content server. • Additional support for Content Manager: – List, add, retrieve, update, and delete of content class – Asynchronous retrieval of object content Flickner 8/29/23 Depo at 34:

		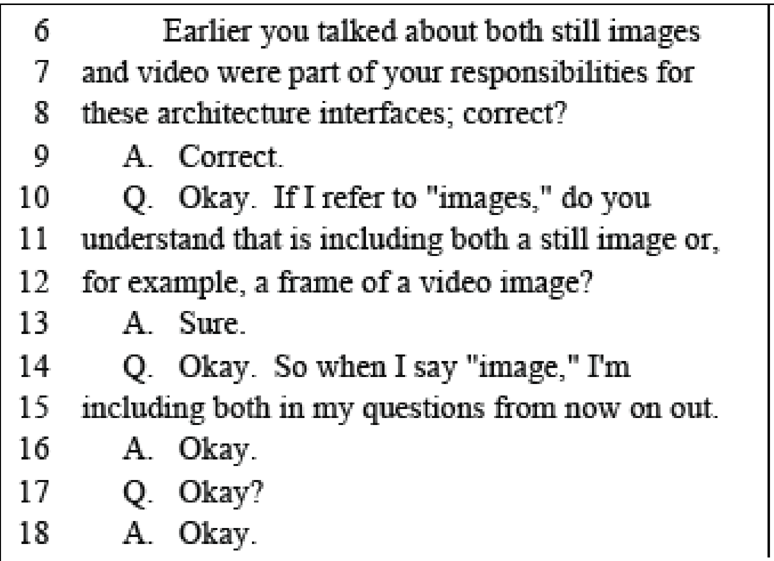 Earlier you talked about both still images and video were part of your responsibilities for these architecture interfaces; correct? A. Correct. Q. Okay. If I refer to "images," do you understand that is including both a still image or, for example, a frame of a video image? A. Sure. Q. Okay. So when I say "image," I'm including both in my questions from now on out. A. Okay. Q. Okay? A. Okay. *Id.* at 60–61:	

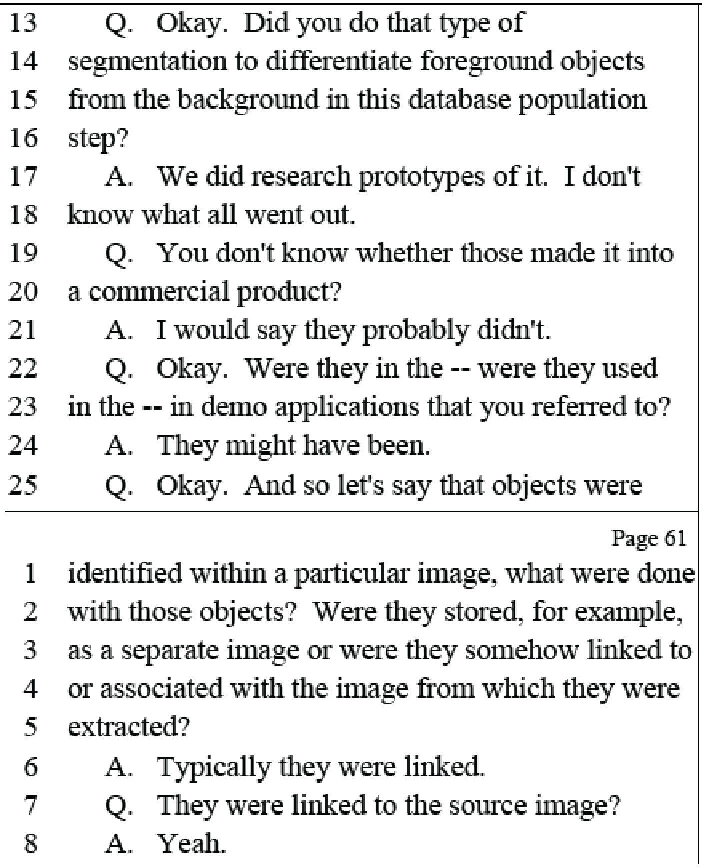

Q. Okay. Did you do that type of segmentation to differentiate foreground objects from the background in this database population step?
A. We did research prototypes of it. I don't know what all went out.
Q. You don't know whether those made it into a commercial product?
A. I would say they probably didn't.
Q. Okay. Were they in the -- were they used in the -- in demo applications that you referred to?
A. They might have been.
Q. Okay. And so let's say that objects were

Page 61

identified within a particular image, what were done with those objects? Were they stored, for example, as a separate image or were they somehow linked to or associated with the image from which they were extracted?
A. Typically they were linked.
Q. They were linked to the source image?
A. Yeah.

Id. at 32–33:

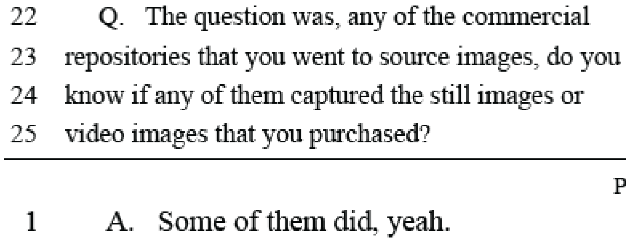

Q. The question was, any of the commercial repositories that you went to source images, do you know if any of them captured the still images or video images that you purchased?

P

A. Some of them did, yeah.

		Id. at 74–75:

			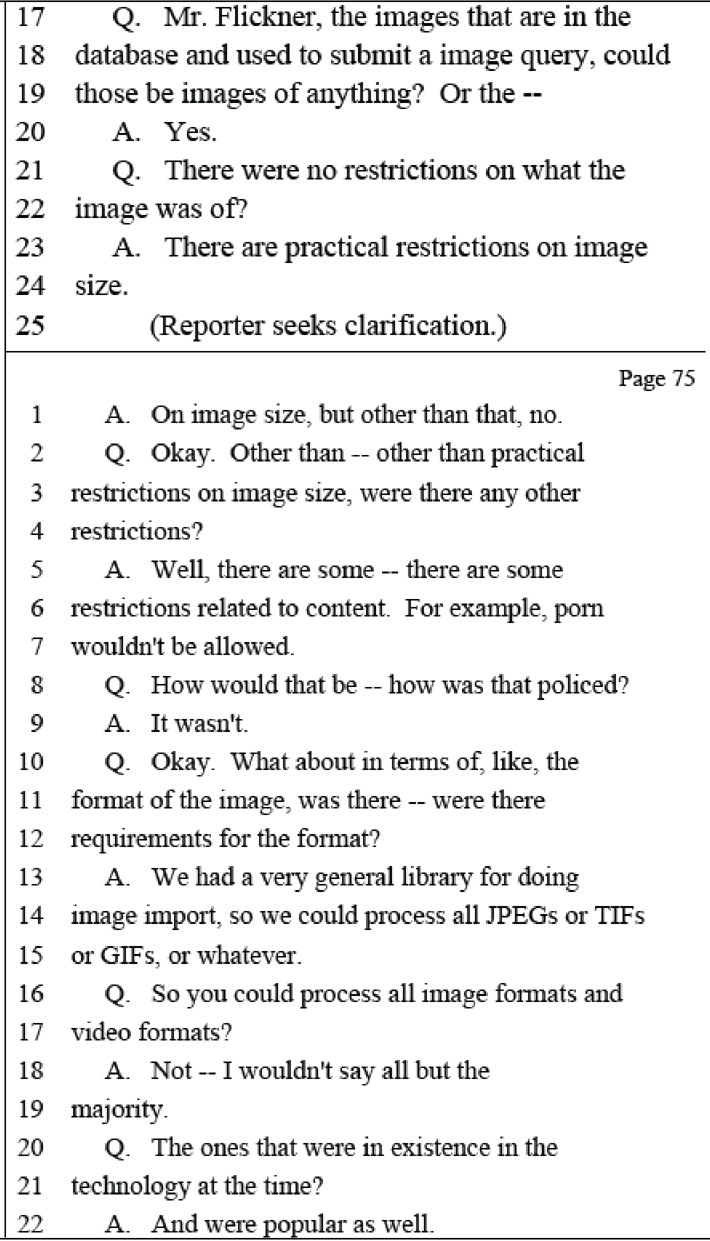 Q. Mr. Flickner, the images that are in the database and used to submit a image query, could those be images of anything? Or the -- A. Yes. Q. There were no restrictions on what the image was of? A. There are practical restrictions on image size. (Reporter seeks clarification.) Page 75 A. On image size, but other than that, no. Q. Okay. Other than -- other than practical restrictions on image size, were there any other restrictions? A. Well, there are some -- there are some restrictions related to content. For example, porn wouldn't be allowed. Q. How would that be -- how was that policed? A. It wasn't. Q. Okay. What about in terms of, like, the format of the image, was there -- were there requirements for the format? A. We had a very general library for doing image import, so we could process all JPEGs or TIFs or GIFs, or whatever. Q. So you could process all image formats and video formats? A. Not -- I wouldn't say all but the majority. Q. The ones that were in existence in the technology at the time? A. And were popular as well.

		Id. at 114: 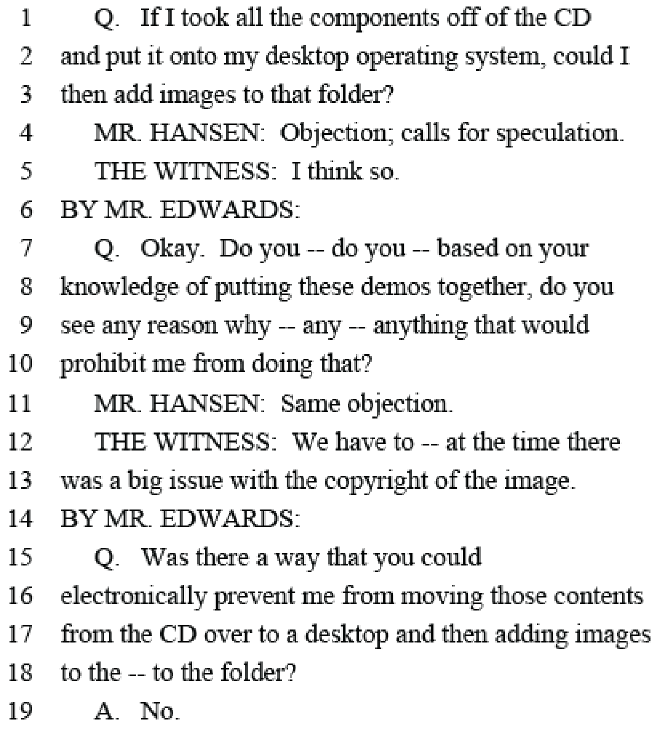 Q. If I took all the components off of the CD and put it onto my desktop operating system, could I then add images to that folder? MR. HANSEN: Objection; calls for speculation. THE WITNESS: I think so. BY MR. EDWARDS: Q. Okay. Do you -- do you -- based on your knowledge of putting these demos together, do you see any reason why -- any -- anything that would prohibit me from doing that? MR. HANSEN: Same objection. THE WITNESS: We have to -- at the time there was a big issue with the copyright of the image. BY MR. EDWARDS: Q. Was there a way that you could electronically prevent me from moving those contents from the CD over to a desktop and then adding images to the -- to the folder? A. No. *Id.* at 202:

		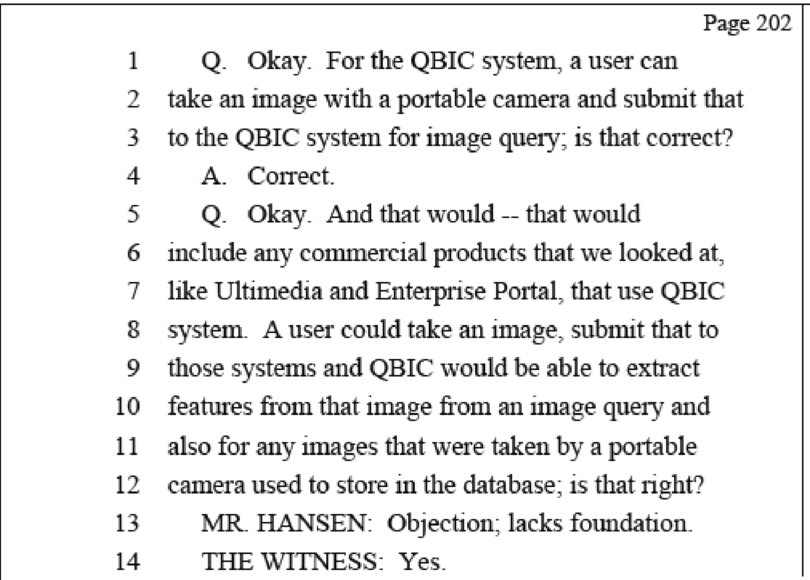 Page 202 Q. Okay. For the QBIC system, a user can take an image with a portable camera and submit that to the QBIC system for image query; is that correct? A. Correct. Q. Okay. And that would -- that would include any commercial products that we looked at, like Ultimedia and Enterprise Portal, that use QBIC system. A user could take an image, submit that to those systems and QBIC would be able to extract features from that image from an image query and also for any images that were taken by a portable camera used to store in the database; is that right? MR. HANSEN: Objection; lacks foundation. THE WITNESS: Yes. *Id.* at 300: 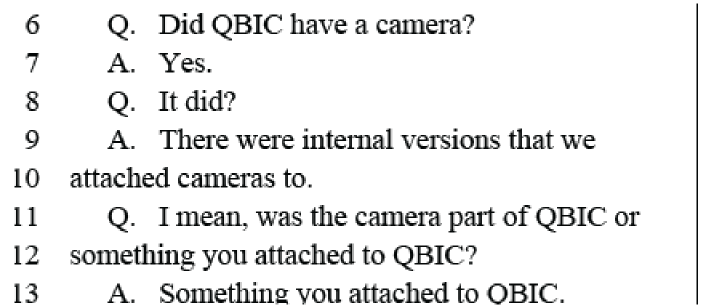 Q. Did QBIC have a camera? A. Yes. Q. It did? A. There were internal versions that we attached cameras to. Q. I mean, was the camera part of QBIC or something you attached to QBIC? A. Something you attached to QBIC. *Id.* at 312:

		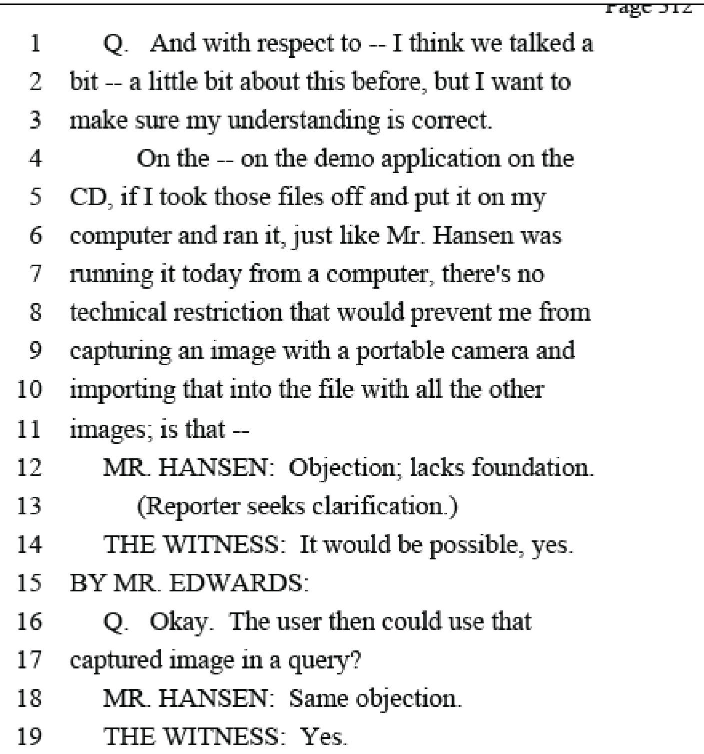 Page 312 Q. And with respect to -- I think we talked a bit -- a little bit about this before, but I want to make sure my understanding is correct. On the -- on the demo application on the CD, if I took those files off and put it on my computer and ran it, just like Mr. Hansen was running it today from a computer, there's no technical restriction that would prevent me from capturing an image with a portable camera and importing that into the file with all the other images; is that -- MR. HANSEN: Objection; lacks foundation. (Reporter seeks clarification.) THE WITNESS: It would be possible, yes. BY MR. EDWARDS: Q. Okay. The user then could use that captured image in a query? MR. HANSEN: Same objection. THE WITNESS: Yes. Flickner 8/29/23 Depo at 42:

		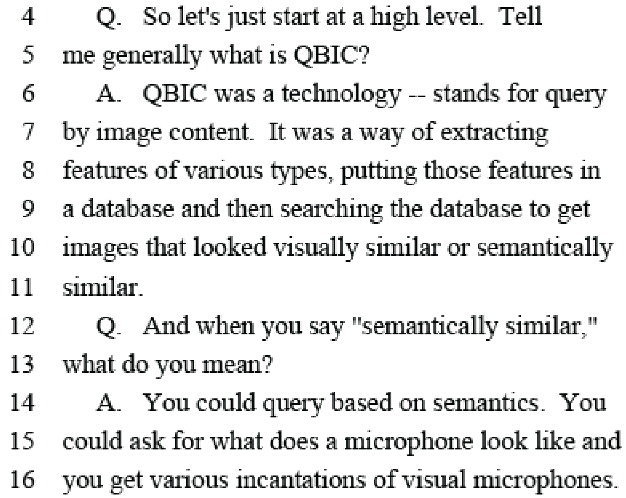 Q. So let's just start at a high level. Tell me generally what is QBIC? A. QBIC was a technology -- stands for query by image content. It was a way of extracting features of various types, putting those features in a database and then searching the database to get images that looked visually similar or semantically similar. Q. And when you say "semantically similar," what do you mean? A. You could query based on semantics. You could ask for what does a microphone look like and you get various incantations of visual microphones. *Id.* at 205–206:

		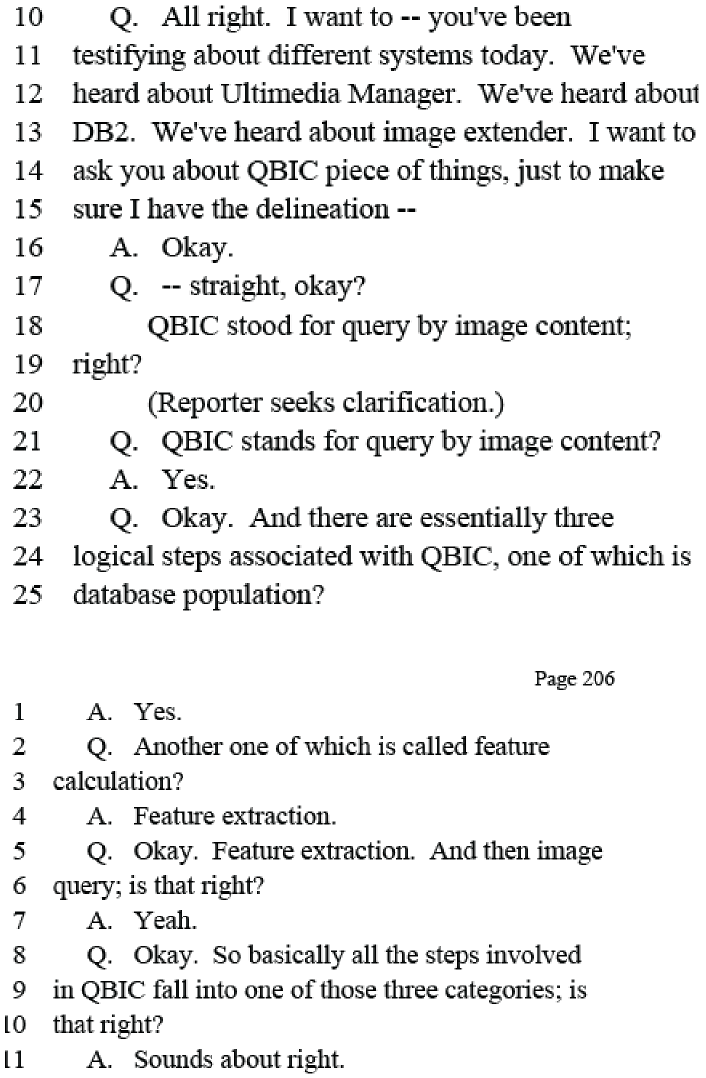 Q. All right. I want to -- you've been testifying about different systems today. We've heard about Ultimedia Manager. We've heard about DB2. We've heard about image extender. I want to ask you about QBIC piece of things, just to make sure I have the delineation -- A. Okay. Q. -- straight, okay? QBIC stood for query by image content; right? (Reporter seeks clarification.) Q. QBIC stands for query by image content? A. Yes. Q. Okay. And there are essentially three logical steps associated with QBIC, one of which is database population? Page 206 A. Yes. Q. Another one of which is called feature calculation? A. Feature extraction. Q. Okay. Feature extraction. And then image query; is that right? A. Yeah. Q. Okay. So basically all the steps involved in QBIC fall into one of those three categories; is that right? A. Sounds about right. Flickner 8/29/23 Depo at Ex. 5:

		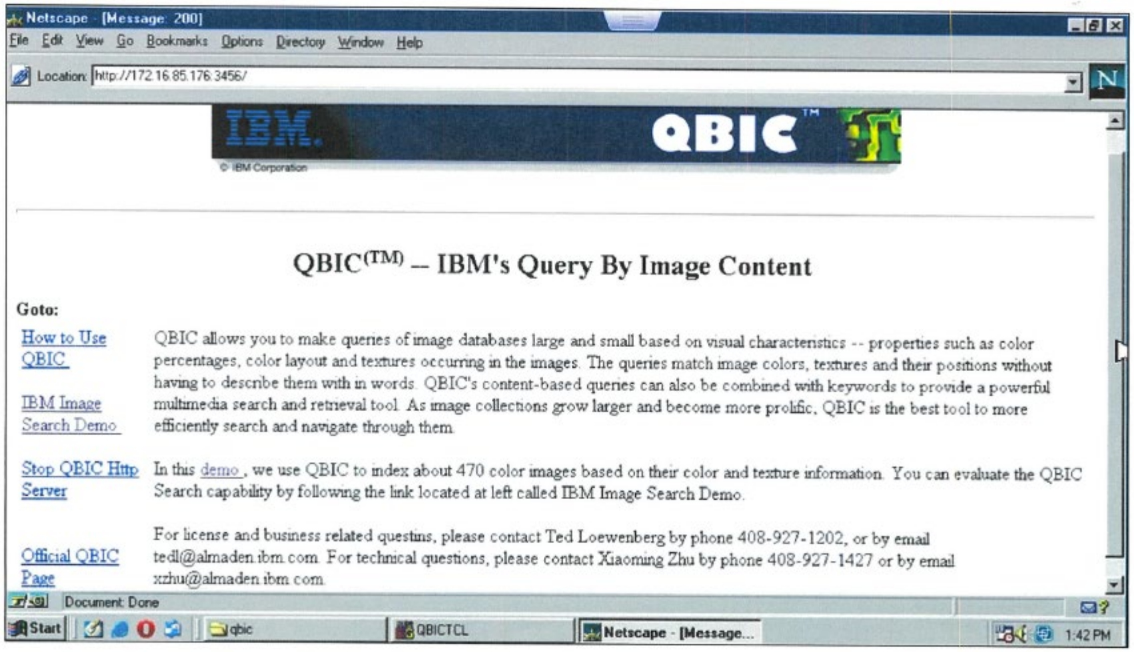 QBIC$^{(TM)}$ -- IBM's Query By Image Content **Goto:** How to Use QBIC IBM Image Search Demo Stop QBIC Http Server Official QBIC Page QBIC allows you to make queries of image databases large and small based on visual characteristics -- properties such as color percentages, color layout and textures occurring in the images. The queries match image colors, textures and their positions without having to describe them with in words. QBIC's content-based queries can also be combined with keywords to provide a powerful multimedia search and retrieval tool. As image collections grow larger and become more prolific, QBIC is the best tool to more efficiently search and navigate through them. In this demo , we use QBIC to index about 470 color images based on their color and texture information. You can evaluate the QBIC Search capability by following the link located at left called IBM Image Search Demo. For license and business related questins, please contact Ted Loewenberg by phone 408-927-1202, or by email tedl@almaden.ibm.com. For technical questions, please contact Xiaoming Zhu by phone 408-927-1427 or by email xzhu@almaden.ibm.com. Flickner 8/29/23 Depo at 79–80:

		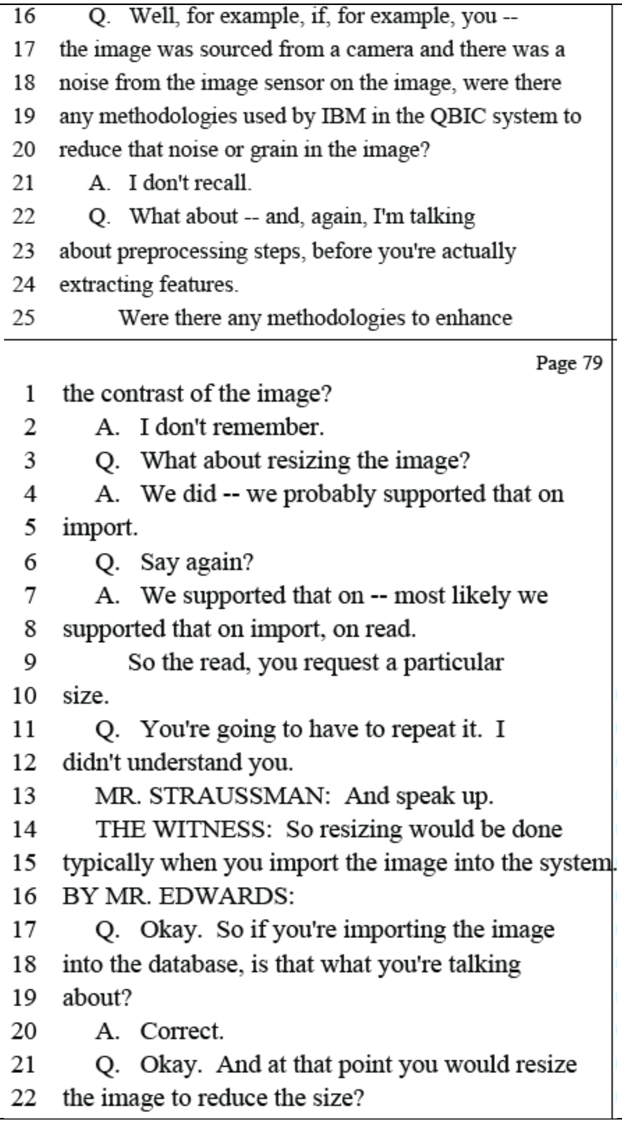 Q. Well, for example, if, for example, you -- the image was sourced from a camera and there was a noise from the image sensor on the image, were there any methodologies used by IBM in the QBIC system to reduce that noise or grain in the image? A. I don't recall. Q. What about -- and, again, I'm talking about preprocessing steps, before you're actually extracting features. Were there any methodologies to enhance Page 79 the contrast of the image? A. I don't remember. Q. What about resizing the image? A. We did -- we probably supported that on import. Q. Say again? A. We supported that on -- most likely we supported that on import, on read. So the read, you request a particular size. Q. You're going to have to repeat it. I didn't understand you. MR. STRAUSSMAN: And speak up. THE WITNESS: So resizing would be done typically when you import the image into the system. BY MR. EDWARDS: Q. Okay. So if you're importing the image into the database, is that what you're talking about? A. Correct. Q. Okay. And at that point you would resize the image to reduce the size?	

		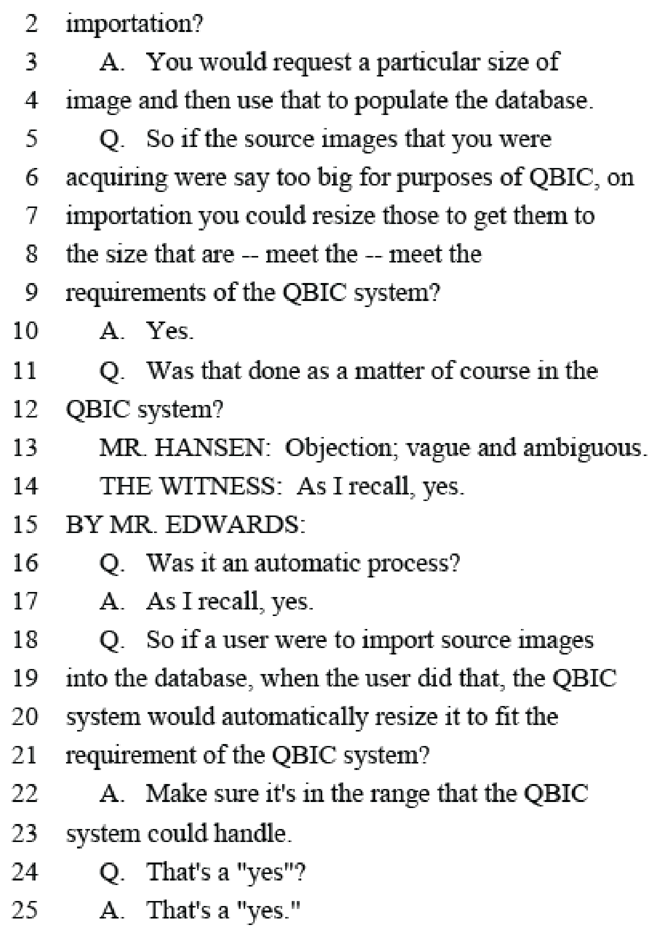 How would you resize the image on importation? A. You would request a particular size of image and then use that to populate the database. Q. So if the source images that you were acquiring were say too big for purposes of QBIC, on importation you could resize those to get them to the size that are -- meet the -- meet the requirements of the QBIC system? A. Yes. Q. Was that done as a matter of course in the QBIC system? MR. HANSEN: Objection; vague and ambiguous. THE WITNESS: As I recall, yes. BY MR. EDWARDS: Q. Was it an automatic process? A. As I recall, yes. Q. So if a user were to import source images into the database, when the user did that, the QBIC system would automatically resize it to fit the requirement of the QBIC system? A. Make sure it's in the range that the QBIC system could handle. Q. That's a "yes"? A. That's a "yes." *Id.* at 80–81:

		Q. How about converting the images to binary? A. Threshold. There probably was a threshold control. Q. And what do you mean by "threshold control"? A. Just change the value of the threshold. You make a binary decision, is a pixel bigger than a -- is a pixel bigger than a particular number or is it smaller and you put a 01 based on that. Q. And that would have been done prior to feature extraction? A. For some features it would be done that way. Q. So, for example, for shape, was it required to convert the image to binary before you did your shape analysis? A. No. Q. Were there any other features that, as part of the -- as part of the extraction process, the image had to be converted to binary before?

		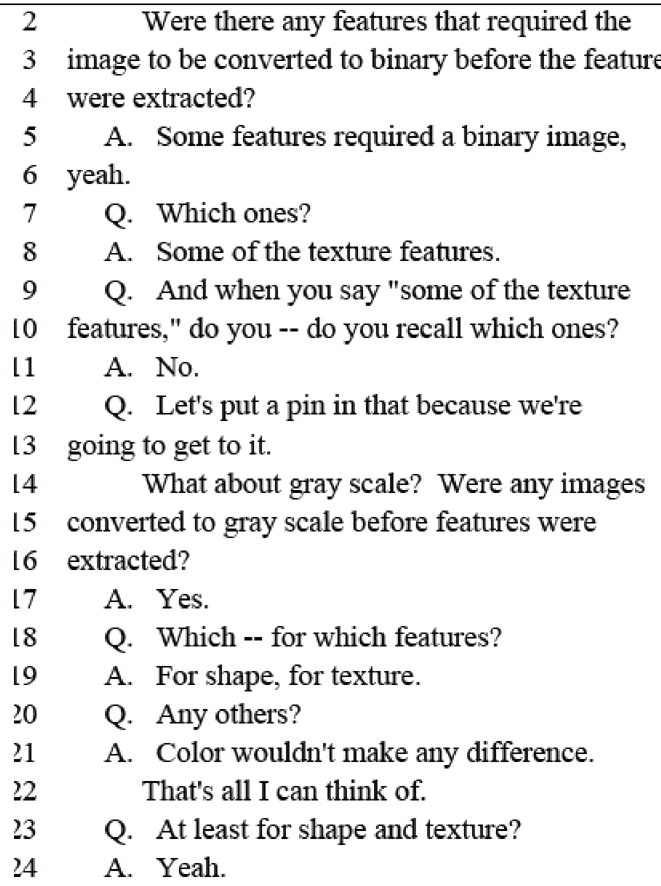 Were there any features that required the image to be converted to binary before the features were extracted? A. Some features required a binary image, yeah. Q. Which ones? A. Some of the texture features. Q. And when you say "some of the texture features," do you -- do you recall which ones? A. No. Q. Let's put a pin in that because we're going to get to it. What about gray scale? Were any images converted to gray scale before features were extracted? A. Yes. Q. Which -- for which features? A. For shape, for texture. Q. Any others? A. Color wouldn't make any difference. That's all I can think of. Q. At least for shape and texture? A. Yeah. *Id.* at 94–97: 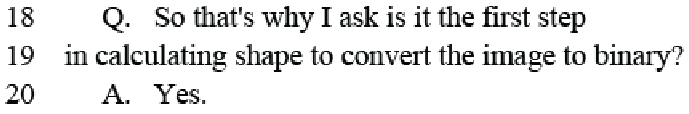 Q. So that's why I ask is it the first step in calculating shape to convert the image to binary? A. Yes. *Id.* at 101–102:

		Okay. So first sketch used something called an edge map of each image? A. Yeah. Q. What is an edge map? A. It's an estimate of the edges, the major edges of the image. Q. Okay. Sort of the outline; is that fair? A. The outline is slightly different, but outlines would be closed, whereas sketch would -- you didn't have to always do closed objects, you could do -- you could draw a circle or you could also draw a U. The circle would find sun-like objects and the U would find things that had a dominant curve -- turn to it. Q. So you could draw -- so, for example, you could draw a flower -- A. Right. Q. -- have the circle part for the flower and then for the stem it would just be a line? And that's where you're saying -- A. Right. Q. -- it doesn't necessarily -- it's broader

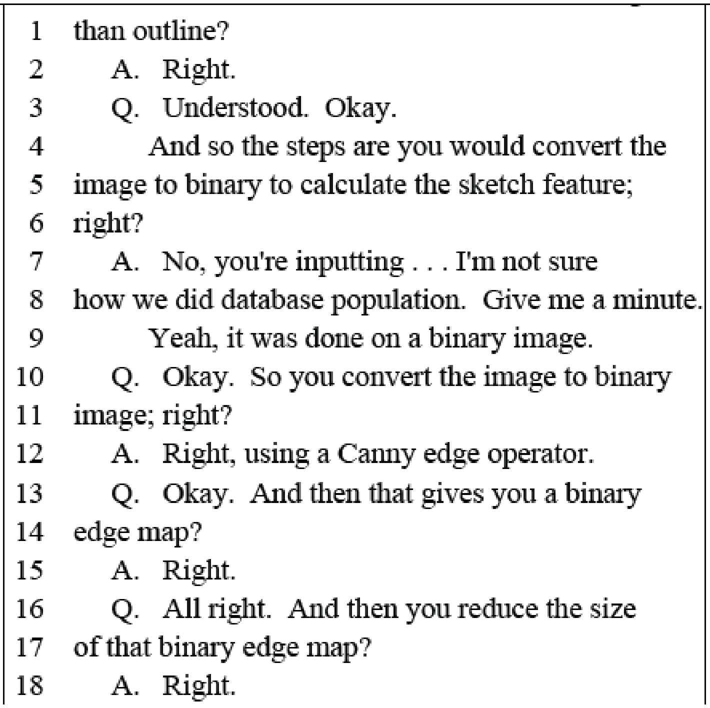
than outline?
A. Right.
Q. Understood. Okay.
And so the steps are you would convert the image to binary to calculate the sketch feature; right?
A. No, you're inputting . . . I'm not sure how we did database population. Give me a minute.
Yeah, it was done on a binary image.
Q. Okay. So you convert the image to binary image; right?
A. Right, using a Canny edge operator.
Q. Okay. And then that gives you a binary edge map?
A. Right.
Q. All right. And then you reduce the size of that binary edge map?
A. Right.

Id. at 140–141:

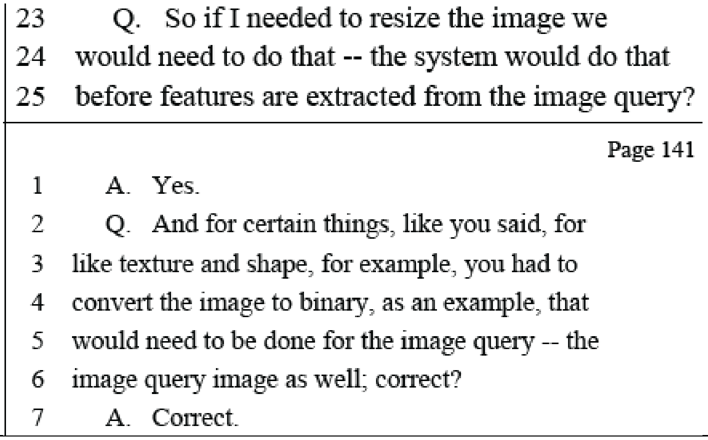
Q. So if I needed to resize the image we would need to do that -- the system would do that before features are extracted from the image query?

Page 141

A. Yes.
Q. And for certain things, like you said, for like texture and shape, for example, you had to convert the image to binary, as an example, that would need to be done for the image query -- the image query image as well; correct?
A. Correct.

Id. at 164–165:

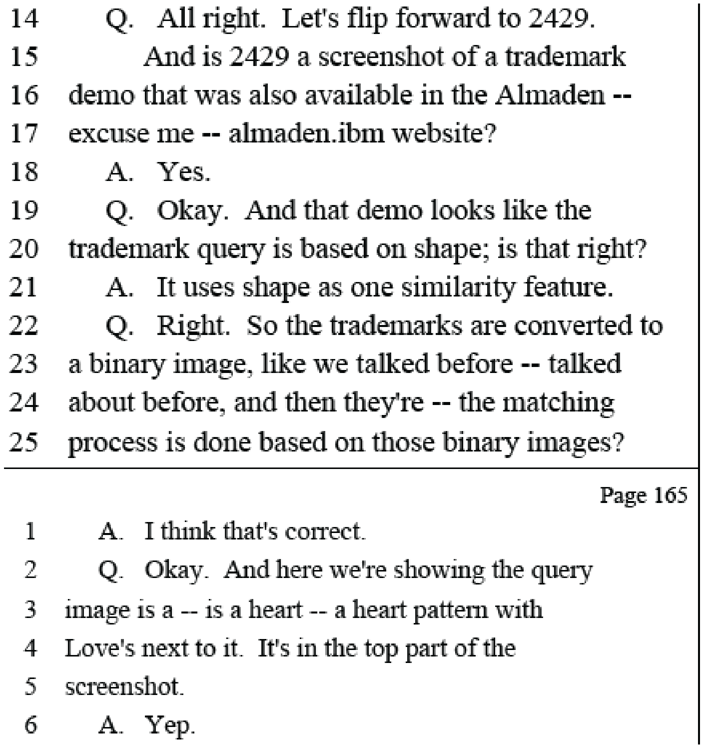

Q. All right. Let's flip forward to 2429.

And is 2429 a screenshot of a trademark demo that was also available in the Almaden -- excuse me -- almaden.ibm website?

A. Yes.

Q. Okay. And that demo looks like the trademark query is based on shape; is that right?

A. It uses shape as one similarity feature.

Q. Right. So the trademarks are converted to a binary image, like we talked before -- talked about before, and then they're -- the matching process is done based on those binary images?

Page 165

A. I think that's correct.

Q. Okay. And here we're showing the query image is a -- is a heart -- a heart pattern with Love's next to it. It's in the top part of the screenshot.

A. Yep.

Id. at 82:

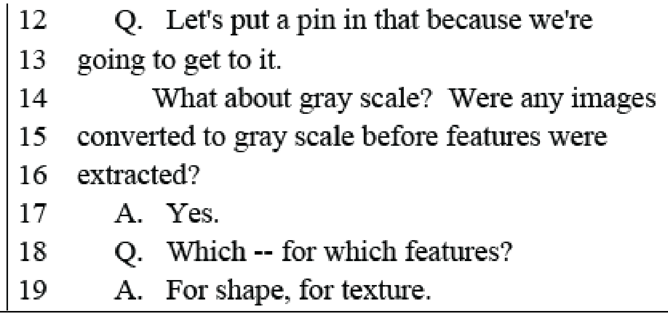

Q. Let's put a pin in that because we're going to get to it.

What about gray scale? Were any images converted to gray scale before features were extracted?

A. Yes.

Q. Which -- for which features?

A. For shape, for texture.

		Id. at 83-85: 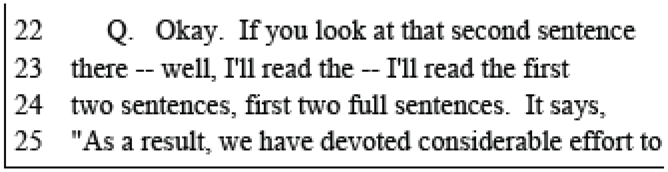 Q. Okay. If you look at that second sentence there -- well, I'll read the -- I'll read the first two sentences, first two full sentences. It says, "As a result, we have devoted considerable effort to Flickner 8/29/23 Depo at 55–61: 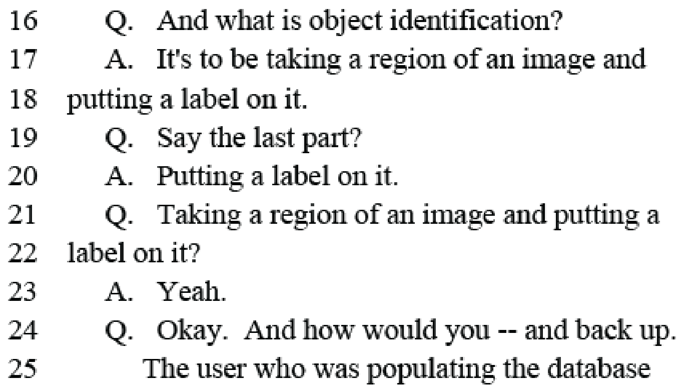 Q. And what is object identification? A. It's to be taking a region of an image and putting a label on it. Q. Say the last part? A. Putting a label on it. Q. Taking a region of an image and putting a label on it? A. Yeah. Q. Okay. And how would you -- and back up. The user who was populating the database

		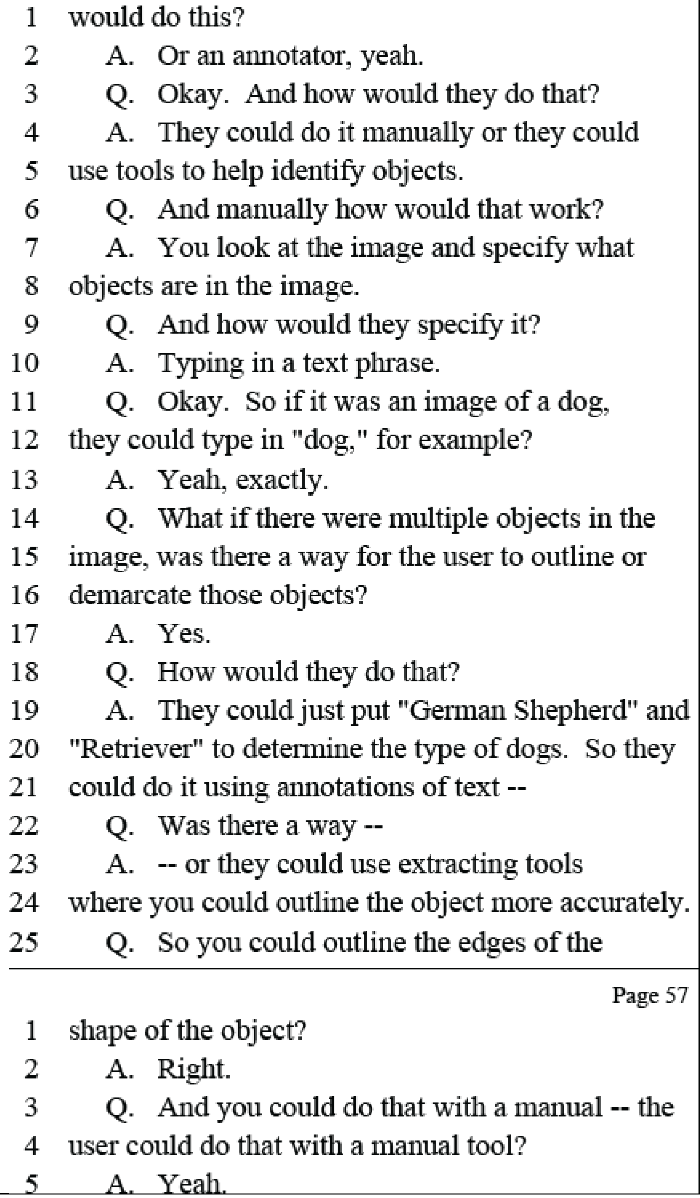 would do this? A. Or an annotator, yeah. Q. Okay. And how would they do that? A. They could do it manually or they could use tools to help identify objects. Q. And manually how would that work? A. You look at the image and specify what objects are in the image. Q. And how would they specify it? A. Typing in a text phrase. Q. Okay. So if it was an image of a dog, they could type in "dog," for example? A. Yeah, exactly. Q. What if there were multiple objects in the image, was there a way for the user to outline or demarcate those objects? A. Yes. Q. How would they do that? A. They could just put "German Shepherd" and "Retriever" to determine the type of dogs. So they could do it using annotations of text -- Q. Was there a way -- A. -- or they could use extracting tools where you could outline the object more accurately. Q. So you could outline the edges of the Page 57 shape of the object? A. Right. Q. And you could do that with a manual -- the user could do that with a manual tool? A. Yeah.	

		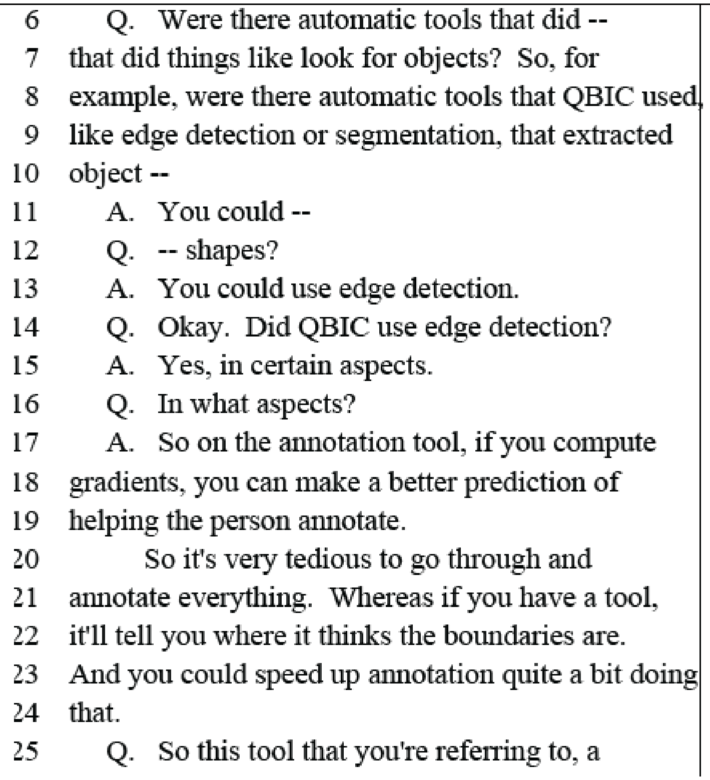 Q. Were there automatic tools that did -- that did things like look for objects? So, for example, were there automatic tools that QBIC used, like edge detection or segmentation, that extracted object -- A. You could -- Q. -- shapes? A. You could use edge detection. Q. Okay. Did QBIC use edge detection? A. Yes, in certain aspects. Q. In what aspects? A. So on the annotation tool, if you compute gradients, you can make a better prediction of helping the person annotate. So it's very tedious to go through and annotate everything. Whereas if you have a tool, it'll tell you where it thinks the boundaries are. And you could speed up annotation quite a bit doing that. Q. So this tool that you're referring to, a	

Page 58

user would provide a rough outline of the object and then the tool would essentially conform that outline better to the actual object in the image?
A. That was one way to do it. The other way to do it is to have the next prediction of the outline to be driven by the tool with the user input.
Q. And what was the -- I assume an algorithm did that?
A. Yeah.
Q. Do you know the -- what algorithm was used?
A. We were doing a lot of stuff with splines at the time.
(Reporter seeks clarification.)
A. Splines. S-p-l-i-n-e-s.
Q. Okay. And what are splines?
A. Splines are curves, mathematical representation of curves -- or a mathematical representation of curves.
Q. And so was also any automatic tools, like Canny edge detection or any other edge detection used that didn't require assistance from the user?
A. There were -- yeah, there were some that were more automatic. Some tools were more

Page 59

automatic.
Q. And those were in the form of algorithms?
A. Yeah.
Q. Do you recall any of the algorithms, the names?
A. We did various edge detection stuff.
Q. Okay.
(Reporter seeks clarification.)
A. Just various types. There were things like Sobel operators that we used --
(Reporter seeks clarification.)
A. Sobel.
Q. How do you spell Sobel?
A. S-o-b-e-l, I think.
Q. Okay.
A. That's one form of an edge detector.
Q. Any others that you can recall?
A. This is kind of image processing 101. There were Gabor filters.
(Reporter seeks clarification.)
Q. How do you spell that?
A. G-a-b-o-r.
There was . . . it's been so long. I don't remember what they call it. Very simple filters, but there's a simple convolution filter

Page 60

that was --
(Reporter seeks clarification.)
A. Simple convolution filter that was used for edge detection. Sobel was an advanced version of it.
Q. And what about segmentation, do you know what segmentation is?
A. Yeah.
Q. Okay. Segmentation, as I understand it, essentially separates foreground objects from the background. Is that a high level of explaining it?
A. That's a two-class segmentation, yeah.
Q. Okay. Did you do that type of segmentation to differentiate foreground objects from the background in this database population step?
A. We did research prototypes of it. I don't know what all went out.
Q. You don't know whether those made it into a commercial product?
A. I would say they probably didn't.
Q. Okay. Were they in the -- were they used in the -- in demo applications that you referred to?
A. They might have been.
Q. Okay. And so let's say that objects were

Page 61

identified within a particular image, what were done with those objects? Were they stored, for example, as a separate image or were they somehow linked to or associated with the image from which they were extracted?
A. Typically they were linked.
Q. They were linked to the source image?
A. Yeah.
Q. Okay. And this was done in the database?
A. Could be done in the database.
Q. How was it done?
A. It was probably done in a text database or some sort of key index pair database. So database has a lot of different meanings.
MR. HANSEN: Text database.
BY MR. EDWARDS:
Q. And, Mr. Flickner, you said probably done in a text database, is that what you said?
A. Yeah.
MR. STRAUSSMAN: Make sure to speak up.
THE WITNESS: Okay.
BY MR. EDWARDS:
Q. Or some sort of key index database, is that the other one?
A. Yes.

		Id. at 305: 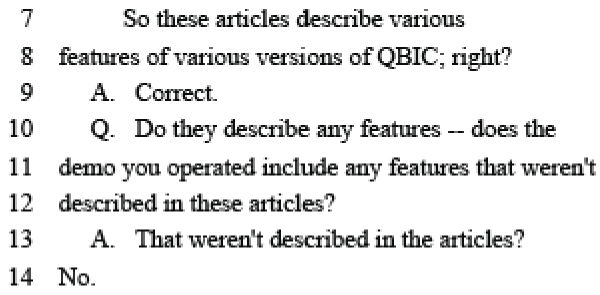 So these articles describe various features of various versions of QBIC; right? A. Correct. Q. Do they describe any features -- does the demo you operated include any features that weren't described in these articles? A. That weren't described in the articles? No. *Id.* at 35:

		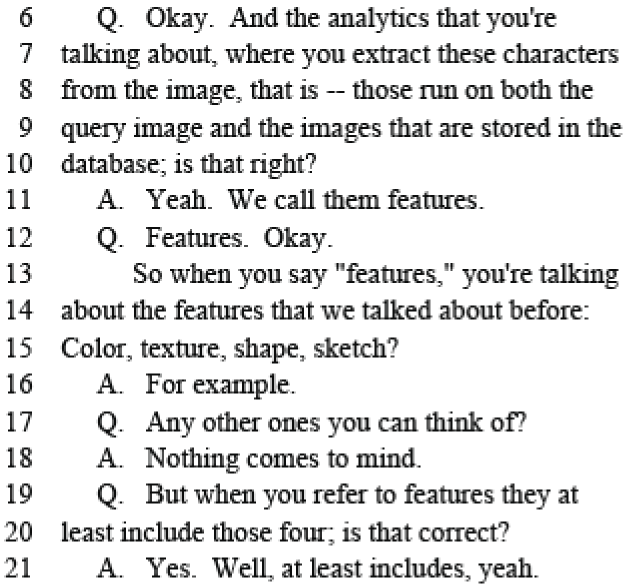 Q. Okay. And the analytics that you're talking about, where you extract these characters from the image, that is -- those run on both the query image and the images that are stored in the database; is that right? A. Yeah. We call them features. Q. Features. Okay. So when you say "features," you're talking about the features that we talked about before: Color, texture, shape, sketch? A. For example. Q. Any other ones you can think of? A. Nothing comes to mind. Q. But when you refer to features they at least include those four; is that correct? A. Yes. Well, at least includes, yeah. *Id.* at 311: 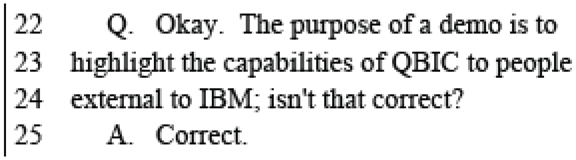 Q. Okay. The purpose of a demo is to highlight the capabilities of QBIC to people external to IBM; isn't that correct? A. Correct. *Id.* at 312-313:

		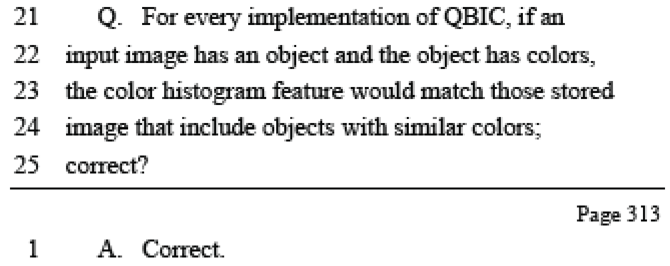 Q. For every implementation of QBIC, if an input image has an object and the object has colors, the color histogram feature would match those stored image that include objects with similar colors; correct? Page 313 A. Correct. *Id.* at 313-317: 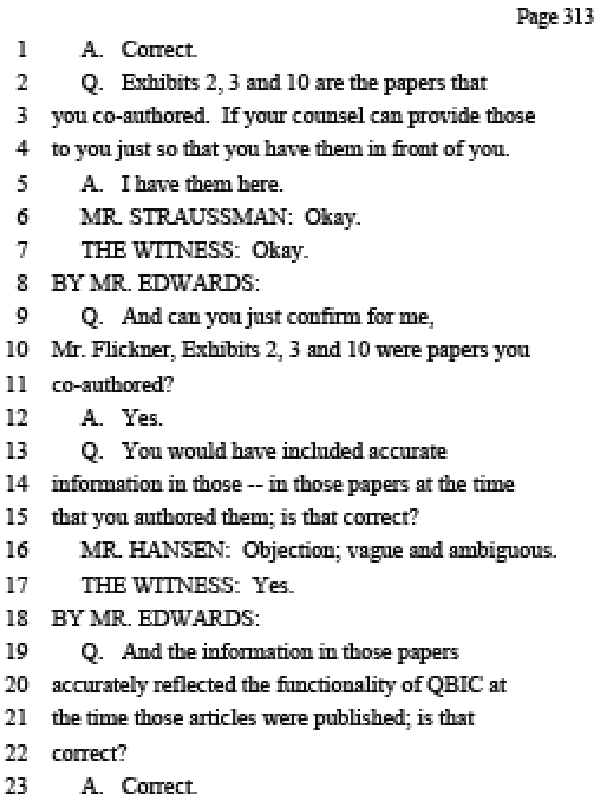 Page 313 A. Correct. Q. Exhibits 2, 3 and 10 are the papers that you co-authored. If your counsel can provide those to you just so that you have them in front of you. A. I have them here. MR. STRAUSSMAN: Okay. THE WITNESS: Okay. BY MR. EDWARDS: Q. And can you just confirm for me, Mr. Flickner, Exhibits 2, 3 and 10 were papers you co-authored? A. Yes. Q. You would have included accurate information in those -- in those papers at the time that you authored them; is that correct? MR. HANSEN: Objection; vague and ambiguous. THE WITNESS: Yes. BY MR. EDWARDS: Q. And the information in those papers accurately reflected the functionality of QBIC at the time those articles were published; is that correct? A. Correct.

		Page 314 BY MR. EDWARDS: Q. And each of those papers would have been peer reviewed before being published; correct? A. Yes. There was peer review. Q. Okay. Let's pick up Exhibit 3, please. A. Okay. Q. You have it -- A. Yeah. Q. -- handy? Okay. On the first page that is Bates-labeled IBM 893, in the right-hand column -- and as you're going down towards the bottom of the right-hand column, before you get to the line that crosses the middle of the page, there are three bullet points. Do you see those bullet points? A. Yeah. Q. It says, "QBIC allows queries on large image and video databases based on," bullet point one, "example images." Do you see that? A. Yeah. Q. Okay. And then also if you go to page 898. A. Okay. Page 315 Q. On the right side it says, "Representative Frame Generation." You see that? A. Yeah. Q. And for database population you can use r-frames of videos just as you can still image for both database population and for feature extraction for a input image; is that fair? MR. HANSEN: Objection; vague and ambiguous, lacks foundation. THE WITNESS: Yes. BY MR. EDWARDS: Q. Let's go to IBM 900. The bottom of the left-hand column to the top of the right-hand column. The bottom of the left-hand column starts, "We have described a prototype system." Do you see that? A. Yeah. Q. Okay. You keep going down, that paragraph refers to the technology described in this article [as read], "has been moved into a commercial standalone product, IBM's Ultimedia Manager, and is part of IBM's Digital Library and DB2 series of products." Do you agree with that?

		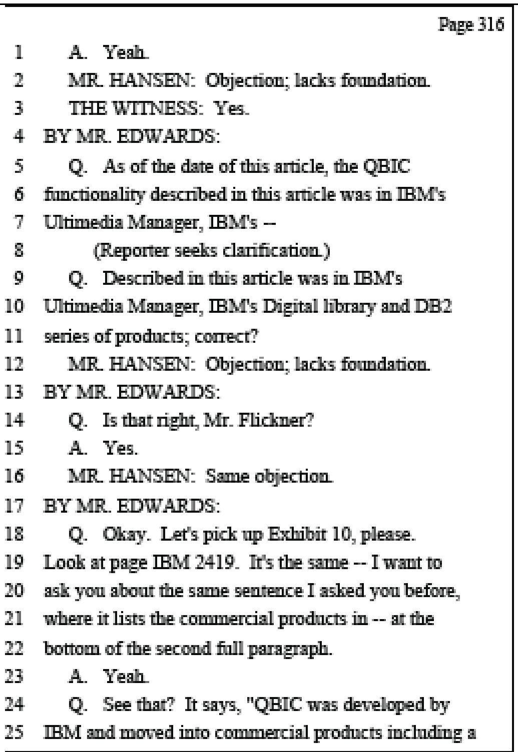 Page 316 A. Yeah. MR. HANSEN: Objection; lacks foundation. THE WITNESS: Yes. BY MR. EDWARDS: Q. As of the date of this article, the QBIC functionality described in this article was in IBM's Ultimedia Manager, IBM's -- (Reporter seeks clarification.) Q. Described in this article was in IBM's Ultimedia Manager, IBM's Digital library and DB2 series of products; correct? MR. HANSEN: Objection; lacks foundation. BY MR. EDWARDS: Q. Is that right, Mr. Flickner? A. Yes. MR. HANSEN: Same objection. BY MR. EDWARDS: Q. Okay. Let's pick up Exhibit 10, please. Look at page IBM 2419. It's the same -- I want to ask you about the same sentence I asked you before, where it lists the commercial products in -- at the bottom of the second full paragraph. A. Yeah. Q. See that? It says, "QBIC was developed by IBM and moved into commercial products including a 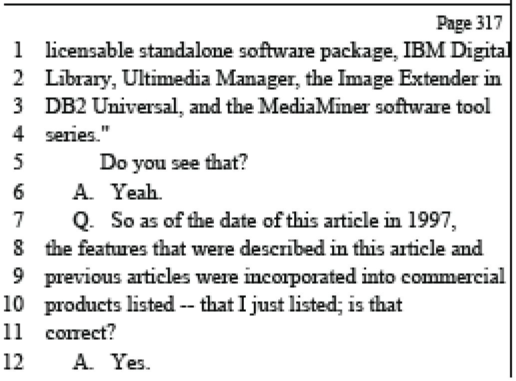 Page 317 licensable standalone software package, IBM Digital Library, Ultimedia Manager, the Image Extender in DB2 Universal, and the MediaMiner software tool series." Do you see that? A. Yeah. Q. So as of the date of this article in 1997, the features that were described in this article and previous articles were incorporated into commercial products listed -- that I just listed; is that correct? A. Yes. Flickner 8/29/23 Depo at 33–37:	

		Q. And the second thing you mentioned was you were responsible for the architecture interface for running analytics. What did that detail? A. You wanted -- you're writing programs to determine similarity. And to do that you would take the image and you'd transform -- you'd extract features from the image, then you would put the features in the database and you'd query the database for the features. Q. And when you say you are determining similarity, you mean similarity between features from one image to another image; is that right? A. Correct. Q. Okay. And what features were you looking at? A. We had features in color, in texture, in shape. A few other categories. Q. Sketch? A. Yep. Q. What's sketch? A. You could draw a sketch -- a rough sketch of an image and query. And so you could draw a	

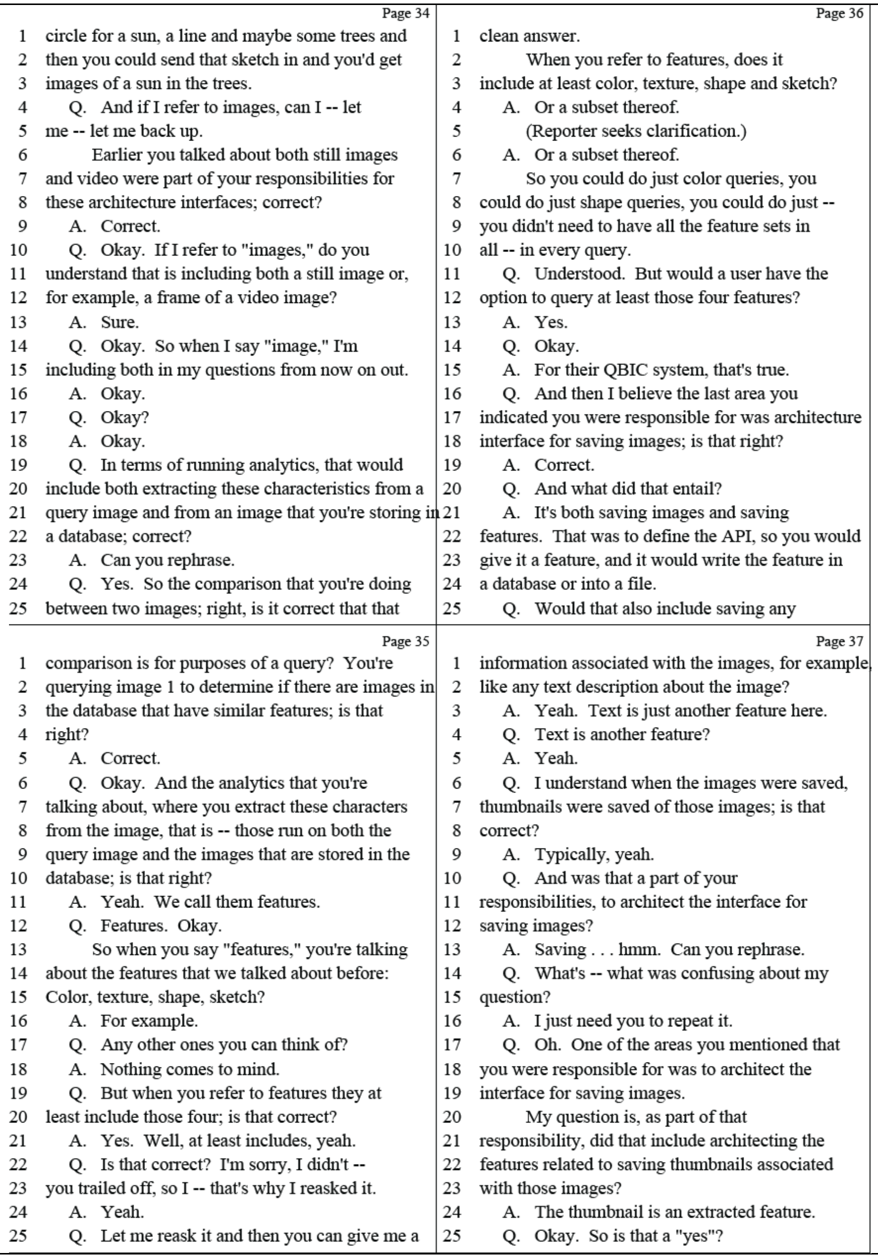

Page 34

circle for a sun, a line and maybe some trees and then you could send that sketch in and you'd get images of a sun in the trees.

Q. And if I refer to images, can I -- let me -- let me back up.

Earlier you talked about both still images and video were part of your responsibilities for these architecture interfaces; correct?

A. Correct.

Q. Okay. If I refer to "images," do you understand that is including both a still image or, for example, a frame of a video image?

A. Sure.

Q. Okay. So when I say "image," I'm including both in my questions from now on out.

A. Okay.

Q. Okay?

A. Okay.

Q. In terms of running analytics, that would include both extracting these characteristics from a query image and from an image that you're storing in a database; correct?

A. Can you rephrase.

Q. Yes. So the comparison that you're doing between two images; right, is it correct that that

Page 35

comparison is for purposes of a query? You're querying image 1 to determine if there are images in the database that have similar features; is that right?

A. Correct.

Q. Okay. And the analytics that you're talking about, where you extract these characters from the image, that is -- those run on both the query image and the images that are stored in the database; is that right?

A. Yeah. We call them features.

Q. Features. Okay.

So when you say "features," you're talking about the features that we talked about before: Color, texture, shape, sketch?

A. For example.

Q. Any other ones you can think of?

A. Nothing comes to mind.

Q. But when you refer to features they at least include those four; is that correct?

A. Yes. Well, at least includes, yeah.

Q. Is that correct? I'm sorry, I didn't -- you trailed off, so I -- that's why I reasked it.

A. Yeah.

Q. Let me reask it and then you can give me a

Page 36

clean answer.

When you refer to features, does it include at least color, texture, shape and sketch?

A. Or a subset thereof.

(Reporter seeks clarification.)

A. Or a subset thereof.

So you could do just color queries, you could do just shape queries, you could do just -- you didn't need to have all the feature sets in all -- in every query.

Q. Understood. But would a user have the option to query at least those four features?

A. Yes.

Q. Okay.

A. For their QBIC system, that's true.

Q. And then I believe the last area you indicated you were responsible for was architecture interface for saving images; is that right?

A. Correct.

Q. And what did that entail?

A. It's both saving images and saving features. That was to define the API, so you would give it a feature, and it would write the feature in a database or into a file.

Q. Would that also include saving any

Page 37

information associated with the images, for example, like any text description about the image?

A. Yeah. Text is just another feature here.

Q. Text is another feature?

A. Yeah.

Q. I understand when the images were saved, thumbnails were saved of those images; is that correct?

A. Typically, yeah.

Q. And was that a part of your responsibilities, to architect the interface for saving images?

A. Saving . . . hmm. Can you rephrase.

Q. What's -- what was confusing about my question?

A. I just need you to repeat it.

Q. Oh. One of the areas you mentioned that you were responsible for was to architect the interface for saving images.

My question is, as part of that responsibility, did that include architecting the features related to saving thumbnails associated with those images?

A. The thumbnail is an extracted feature.

Q. Okay. So is that a "yes"?

Id. at 52–53:

Q. Underneath there's a sentence that says, "There are three logical steps in a QBIC application: Database population, feature calculation, and image query."
Do you see that?
A. Yeah.
Q. You agree with that?
A. Yeah.
Q. All right. Well, let's talk about the first step, database population.
What is database population?
A. It's taking a set of images, extracting the features and putting them in a database.
Q. Okay. Is part of the database population

Page 53

step to create or prepare a thumbnail and store that in the database as well?
A. Could be.
Q. So if you look at -- under 2.1 Database Population, it says, "The first step in population is to simply load the images into the system."
Is that right?
A. Yep.
Q. And the next sentence says then you can also add the image to the database, preparing a reduced thumbnail and adding any available text information to the database.
A. Yep.
Q. Is that correct?
A. Yep.

Id. at 77–78:

		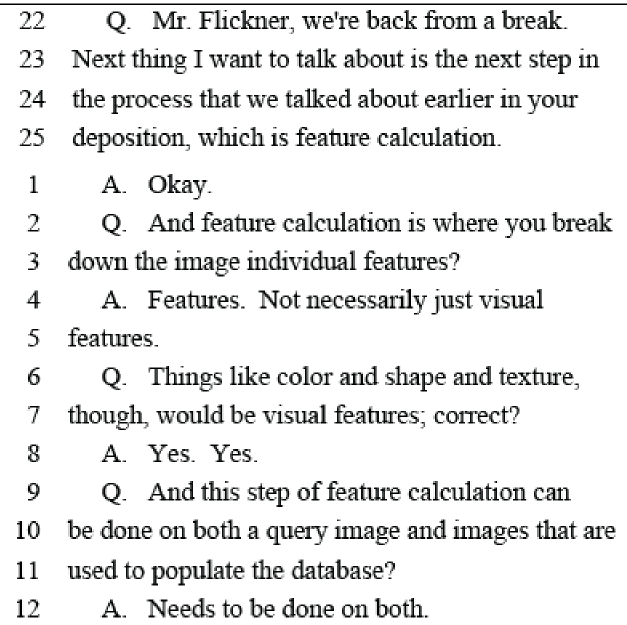 Q. Mr. Flickner, we're back from a break. Next thing I want to talk about is the next step in the process that we talked about earlier in your deposition, which is feature calculation. A. Okay. Q. And feature calculation is where you break down the image individual features? A. Features. Not necessarily just visual features. Q. Things like color and shape and texture, though, would be visual features; correct? A. Yes. Yes. Q. And this step of feature calculation can be done on both a query image and images that are used to populate the database? A. Needs to be done on both. *Id.* at 85–99: 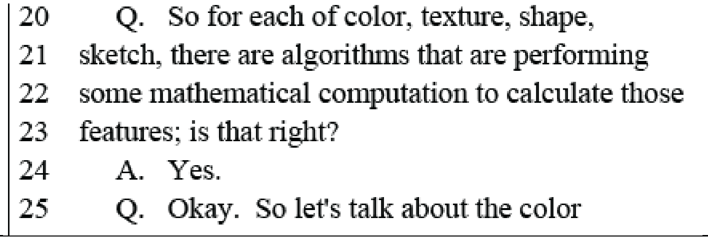 Q. So for each of color, texture, shape, sketch, there are algorithms that are performing some mathematical computation to calculate those features; is that right? A. Yes. Q. Okay. So let's talk about the color

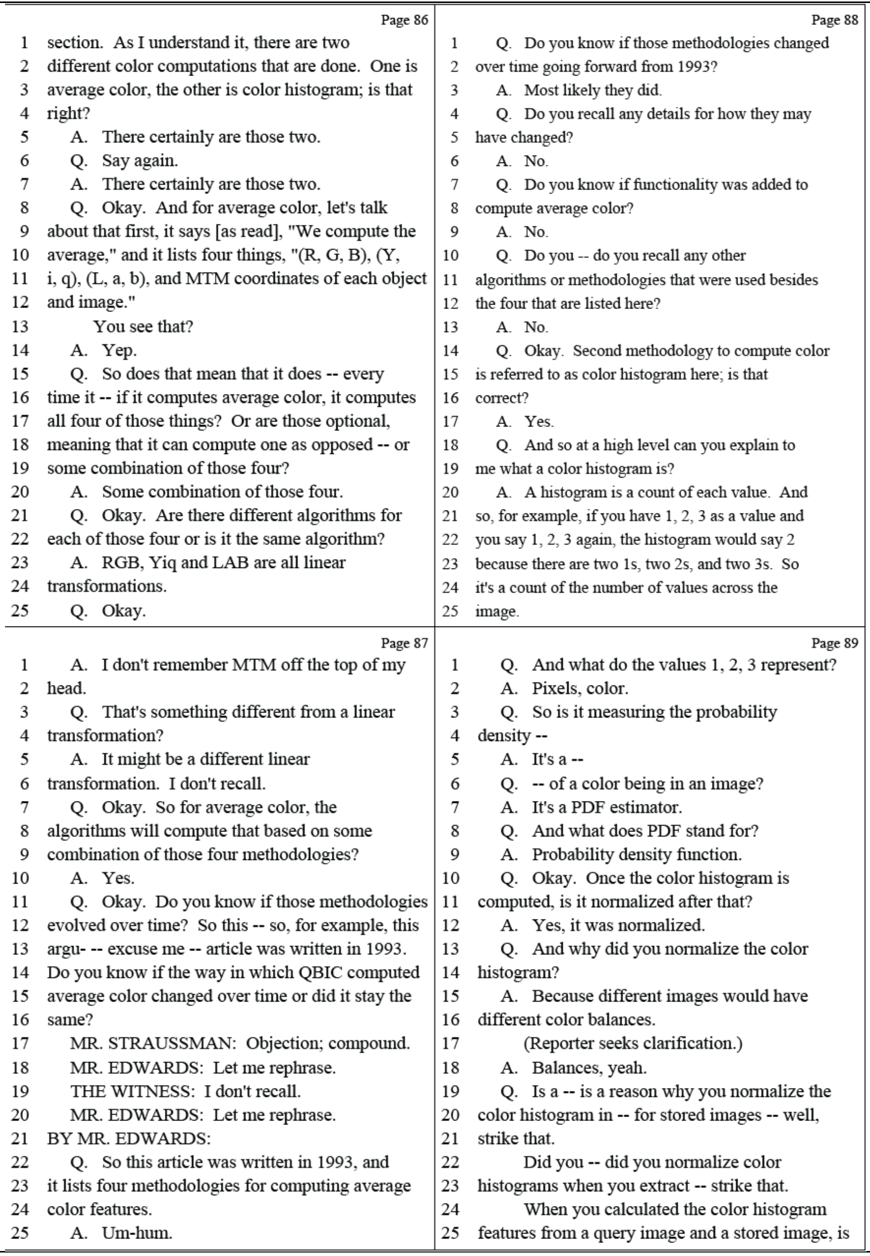

Page 86

section. As I understand it, there are two different color computations that are done. One is average color, the other is color histogram; is that right?
A. There certainly are those two.
Q. Say again.
A. There certainly are those two.
Q. Okay. And for average color, let's talk about that first, it says [as read], "We compute the average," and it lists four things, "(R, G, B), (Y, i, q), (L, a, b), and MTM coordinates of each object and image."
You see that?
A. Yep.
Q. So does that mean that it does -- every time it -- if it computes average color, it computes all four of those things? Or are those optional, meaning that it can compute one as opposed -- or some combination of those four?
A. Some combination of those four.
Q. Okay. Are there different algorithms for each of those four or is it the same algorithm?
A. RGB, Yiq and LAB are all linear transformations.
Q. Okay.

Page 87

A. I don't remember MTM off the top of my head.
Q. That's something different from a linear transformation?
A. It might be a different linear transformation. I don't recall.
Q. Okay. So for average color, the algorithms will compute that based on some combination of those four methodologies?
A. Yes.
Q. Okay. Do you know if those methodologies evolved over time? So this -- so, for example, this argu- -- excuse me -- article was written in 1993. Do you know if the way in which QBIC computed average color changed over time or did it stay the same?
MR. STRAUSSMAN: Objection; compound.
MR. EDWARDS: Let me rephrase.
THE WITNESS: I don't recall.
MR. EDWARDS: Let me rephrase.
BY MR. EDWARDS:
Q. So this article was written in 1993, and it lists four methodologies for computing average color features.
A. Um-hum.

Page 88

Q. Do you know if those methodologies changed over time going forward from 1993?
A. Most likely they did.
Q. Do you recall any details for how they may have changed?
A. No.
Q. Do you know if functionality was added to compute average color?
A. No.
Q. Do you -- do you recall any other algorithms or methodologies that were used besides the four that are listed here?
A. No.
Q. Okay. Second methodology to compute color is referred to as color histogram here; is that correct?
A. Yes.
Q. And so at a high level can you explain to me what a color histogram is?
A. A histogram is a count of each value. And so, for example, if you have 1, 2, 3 as a value and you say 1, 2, 3 again, the histogram would say 2 because there are two 1s, two 2s, and two 3s. So it's a count of the number of values across the image.

Page 89

Q. And what do the values 1, 2, 3 represent?
A. Pixels, color.
Q. So is it measuring the probability density --
A. It's a --
Q. -- of a color being in an image?
A. It's a PDF estimator.
Q. And what does PDF stand for?
A. Probability density function.
Q. Okay. Once the color histogram is computed, is it normalized after that?
A. Yes, it was normalized.
Q. And why did you normalize the color histogram?
A. Because different images would have different color balances.
(Reporter seeks clarification.)
A. Balances, yeah.
Q. Is a -- is a reason why you normalize the color histogram in -- for stored images -- well, strike that.
Did you -- did you normalize color histograms when you extract -- strike that.
When you calculated the color histogram features from a query image and a stored image, is

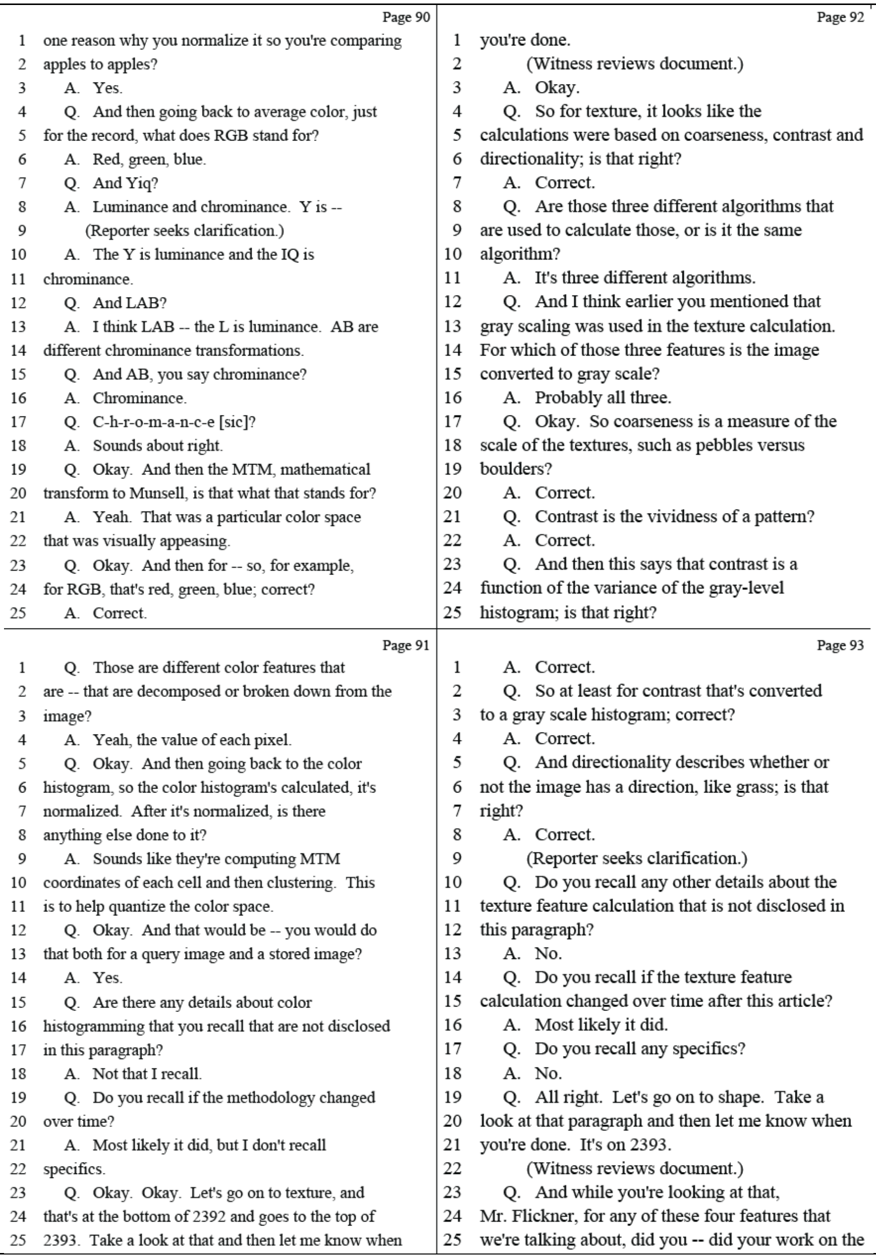

Page 90

one reason why you normalize it so you're comparing apples to apples?
A. Yes.
Q. And then going back to average color, just for the record, what does RGB stand for?
A. Red, green, blue.
Q. And Yiq?
A. Luminance and chrominance. Y is --
(Reporter seeks clarification.)
A. The Y is luminance and the IQ is chrominance.
Q. And LAB?
A. I think LAB -- the L is luminance. AB are different chrominance transformations.
Q. And AB, you say chrominance?
A. Chrominance.
Q. C-h-r-o-m-a-n-c-e [sic]?
A. Sounds about right.
Q. Okay. And then the MTM, mathematical transform to Munsell, is that what that stands for?
A. Yeah. That was a particular color space that was visually appeasing.
Q. Okay. And then for -- so, for example, for RGB, that's red, green, blue; correct?
A. Correct.

Page 91

Q. Those are different color features that are -- that are decomposed or broken down from the image?
A. Yeah, the value of each pixel.
Q. Okay. And then going back to the color histogram, so the color histogram's calculated, it's normalized. After it's normalized, is there anything else done to it?
A. Sounds like they're computing MTM coordinates of each cell and then clustering. This is to help quantize the color space.
Q. Okay. And that would be -- you would do that both for a query image and a stored image?
A. Yes.
Q. Are there any details about color histogramming that you recall that are not disclosed in this paragraph?
A. Not that I recall.
Q. Do you recall if the methodology changed over time?
A. Most likely it did, but I don't recall specifics.
Q. Okay. Okay. Let's go on to texture, and that's at the bottom of 2392 and goes to the top of 2393. Take a look at that and then let me know when

Page 92

you're done.
(Witness reviews document.)
A. Okay.
Q. So for texture, it looks like the calculations were based on coarseness, contrast and directionality; is that right?
A. Correct.
Q. Are those three different algorithms that are used to calculate those, or is it the same algorithm?
A. It's three different algorithms.
Q. And I think earlier you mentioned that gray scaling was used in the texture calculation. For which of those three features is the image converted to gray scale?
A. Probably all three.
Q. Okay. So coarseness is a measure of the scale of the textures, such as pebbles versus boulders?
A. Correct.
Q. Contrast is the vividness of a pattern?
A. Correct.
Q. And then this says that contrast is a function of the variance of the gray-level histogram; is that right?

Page 93

A. Correct.
Q. So at least for contrast that's converted to a gray scale histogram; correct?
A. Correct.
Q. And directionality describes whether or not the image has a direction, like grass; is that right?
A. Correct.
(Reporter seeks clarification.)
Q. Do you recall any other details about the texture feature calculation that is not disclosed in this paragraph?
A. No.
Q. Do you recall if the texture feature calculation changed over time after this article?
A. Most likely it did.
Q. Do you recall any specifics?
A. No.
Q. All right. Let's go on to shape. Take a look at that paragraph and then let me know when you're done. It's on 2393.
(Witness reviews document.)
Q. And while you're looking at that, Mr. Flickner, for any of these four features that we're talking about, did you -- did your work on the

Page 94

QBIC project -- were you more involved in one or more of these features than the other?
A. Yes.
Q. Which one?
A. Shape.
Q. Okay.
(Witness reviews document.)
A. So one thing I can state is originally I couldn't remember the name of the edge operator -- and the name is Canny, the Canny edge operator -- when we were talking about finding edges.
Q. Okay. So this -- the paragraph on IBM 2393 refreshed your memory that shape calculation used Canny edge detection?
A. No, that's sketch.
Q. Oh, I'm sorry.
A. I'm just stating that originally I was asked a question about standard edge operators. I couldn't remember the name. The name was Canny.
Q. Canny. And Canny is used in the sketch calculation?
A. Yeah, it's used in the sketch calculation.
Q. Understood. Okay. Let's go back to shape.
And for shape, is it a first step -- well,

Page 95

back up, strike that.
To calculate the shape feature in a query image or a stored database image, is the first step to convert that image into a binary image?
A. I don't think so.
Q. So if you look at the second paragraph -- and the reason why I ask is -- it's the second sentence in the second paragraph under "Shape features." It says [as read], "All shapes are assumed to be non-occluded planar shapes allowing for each shape to be represented as a binary image."
See that?
A. Yeah.
Q. And then the next sentence says [as read], "The area is computed as the number of pixels set in the binary image."
A. Right.
Q. So that's why I ask is it the first step in calculating shape to convert the image to binary?
A. Yes.
Q. And then once you convert it to binary image, then what do you do next in order to calculate the shape?
A. You're computing something called perround, which is perimeter squared over the area,

Page 96

of the object, of the blob.
Q. Can you say the name of the word again?
A. I think it's called perround.
Q. Per round?
A. Yeah.
Q. Those -- are those two words, p-e-r, r-o-u-n-d?
A. P-e-r -- one word.
Q. Okay. Perround, one word.
A. It's basically the perimeter squared over area.
(Reporter seeks clarification.)
A. Perimeter squared over area.
(Reporter seeks clarification.)
A. Length of the -- over area.
Q. Okay. So let's step back a little bit. So it says -- the second paragraph says [as read], "Our shape features are based on a combination of heuristic" -- h-e-u-r-i-s-t-i-c -- "shape features area, circularity, eccentricity, major access orientation and a set of algebraic moment invariants," i-n-v-a-r-i-a-n-t-s.
Did I read that correctly?
A. Um-hum.
Q. Okay. For circularity, that is computed

Page 97

as perimeter squared over area; correct?
A. Right.
Q. All right. Area is the number of pixels in the binary image?
A. In the binary blob.
Q. In the binary what?
A. Blob.
Q. Blob. And when you say "blob," what does that refer to?
A. The outline of the object that you're doing a shape query on.
Q. And so are each of these, area, circularity, eccentricity, major axis orientation and algebraic moment invariants, are those each different algorithms or is it all the same algorithm?
A. It's different algorithms.
Q. Okay. And since you worked on the shape feature the most, do you recall any other details that are not disclosed in this paragraph?
A. Well, this evolved and we ended up doing an implementation of something that's not in here.
Q. Okay. And what was that implementation?
A. Something called turning angle per round, turning angle, which is -- it looks at the tangent

		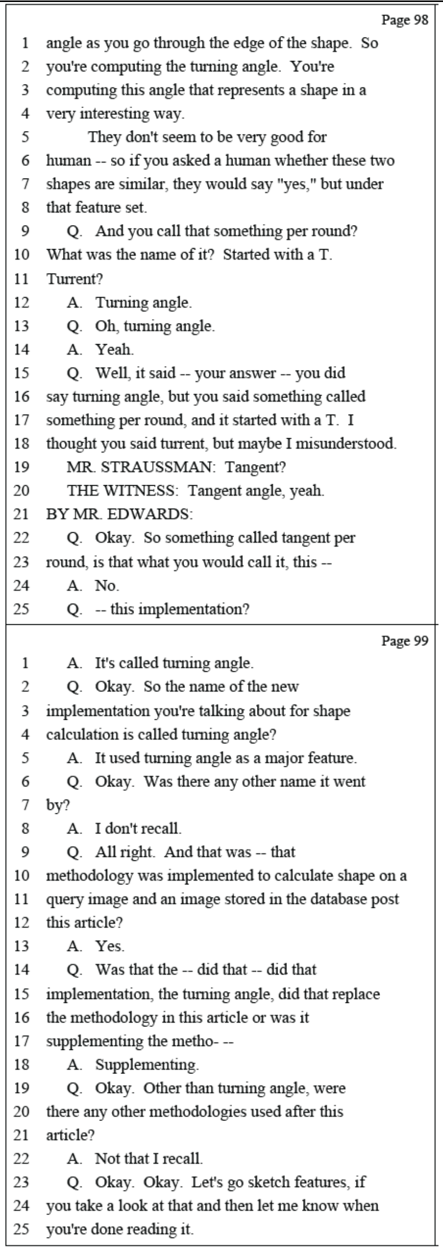 Page 98 angle as you go through the edge of the shape. So you're computing the turning angle. You're computing this angle that represents a shape in a very interesting way. They don't seem to be very good for human -- so if you asked a human whether these two shapes are similar, they would say "yes," but under that feature set. Q. And you call that something per round? What was the name of it? Started with a T. Turrent? A. Turning angle. Q. Oh, turning angle. A. Yeah. Q. Well, it said -- your answer -- you did say turning angle, but you said something called something per round, and it started with a T. I thought you said turrent, but maybe I misunderstood. MR. STRAUSSMAN: Tangent? THE WITNESS: Tangent angle, yeah. BY MR. EDWARDS: Q. Okay. So something called tangent per round, is that what you would call it, this -- A. No. Q. -- this implementation? Page 99 A. It's called turning angle. Q. Okay. So the name of the new implementation you're talking about for shape calculation is called turning angle? A. It used turning angle as a major feature. Q. Okay. Was there any other name it went by? A. I don't recall. Q. All right. And that was -- that methodology was implemented to calculate shape on a query image and an image stored in the database post this article? A. Yes. Q. Was that the -- did that -- did that implementation, the turning angle, did that replace the methodology in this article or was it supplementing the metho- -- A. Supplementing. Q. Okay. Other than turning angle, were there any other methodologies used after this article? A. Not that I recall. Q. Okay. Okay. Let's go sketch features, if you take a look at that and then let me know when you're done reading it.

Id. at 119–128:

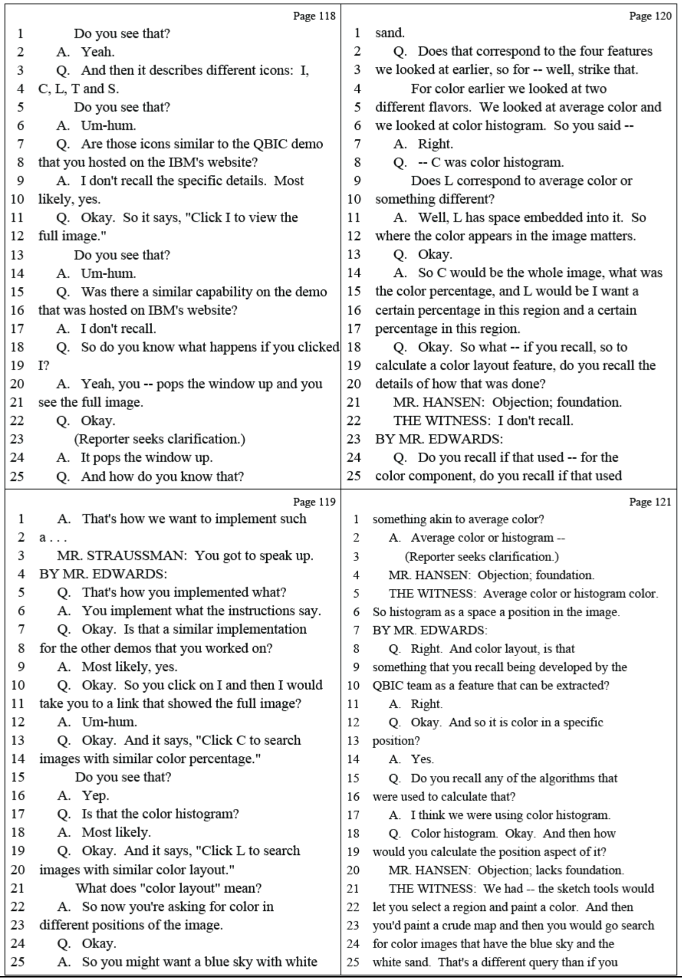

Page 118

Do you see that?
A. Yeah.
Q. And then it describes different icons: I, C, L, T and S.
Do you see that?
A. Um-hum.
Q. Are those icons similar to the QBIC demo that you hosted on the IBM's website?
A. I don't recall the specific details. Most likely, yes.
Q. Okay. So it says, "Click I to view the full image."
Do you see that?
A. Um-hum.
Q. Was there a similar capability on the demo that was hosted on IBM's website?
A. I don't recall.
Q. So do you know what happens if you clicked I?
A. Yeah, you -- pops the window up and you see the full image.
Q. Okay.
(Reporter seeks clarification.)
A. It pops the window up.
Q. And how do you know that?

Page 119

A. That's how we want to implement such a . . .
MR. STRAUSSMAN: You got to speak up.
BY MR. EDWARDS:
Q. That's how you implemented what?
A. You implement what the instructions say.
Q. Okay. Is that a similar implementation for the other demos that you worked on?
A. Most likely, yes.
Q. Okay. So you click on I and then I would take you to a link that showed the full image?
A. Um-hum.
Q. Okay. And it says, "Click C to search images with similar color percentage."
Do you see that?
A. Yep.
Q. Is that the color histogram?
A. Most likely.
Q. Okay. And it says, "Click L to search images with similar color layout."
What does "color layout" mean?
A. So now you're asking for color in different positions of the image.
Q. Okay.
A. So you might want a blue sky with white

Page 120

sand.
Q. Does that correspond to the four features we looked at earlier, so for -- well, strike that.
For color earlier we looked at two different flavors. We looked at average color and we looked at color histogram. So you said --
A. Right.
Q. -- C was color histogram.
Does L correspond to average color or something different?
A. Well, L has space embedded into it. So where the color appears in the image matters.
Q. Okay.
A. So C would be the whole image, what was the color percentage, and L would be I want a certain percentage in this region and a certain percentage in this region.
Q. Okay. So what -- if you recall, so to calculate a color layout feature, do you recall the details of how that was done?
MR. HANSEN: Objection; foundation.
THE WITNESS: I don't recall.
BY MR. EDWARDS:
Q. Do you recall if that used -- for the color component, do you recall if that used

Page 121

something akin to average color?
A. Average color or histogram --
(Reporter seeks clarification.)
MR. HANSEN: Objection; foundation.
THE WITNESS: Average color or histogram color. So histogram as a space a position in the image.
BY MR. EDWARDS:
Q. Right. And color layout, is that something that you recall being developed by the QBIC team as a feature that can be extracted?
A. Right.
Q. Okay. And so it is color in a specific position?
A. Yes.
Q. Do you recall any of the algorithms that were used to calculate that?
A. I think we were using color histogram.
Q. Color histogram. Okay. And then how would you calculate the position aspect of it?
MR. HANSEN: Objection; lacks foundation.
THE WITNESS: We had -- the sketch tools would let you select a region and paint a color. And then you'd paint a crude map and then you would go search for color images that have the blue sky and the white sand. That's a different query than if you

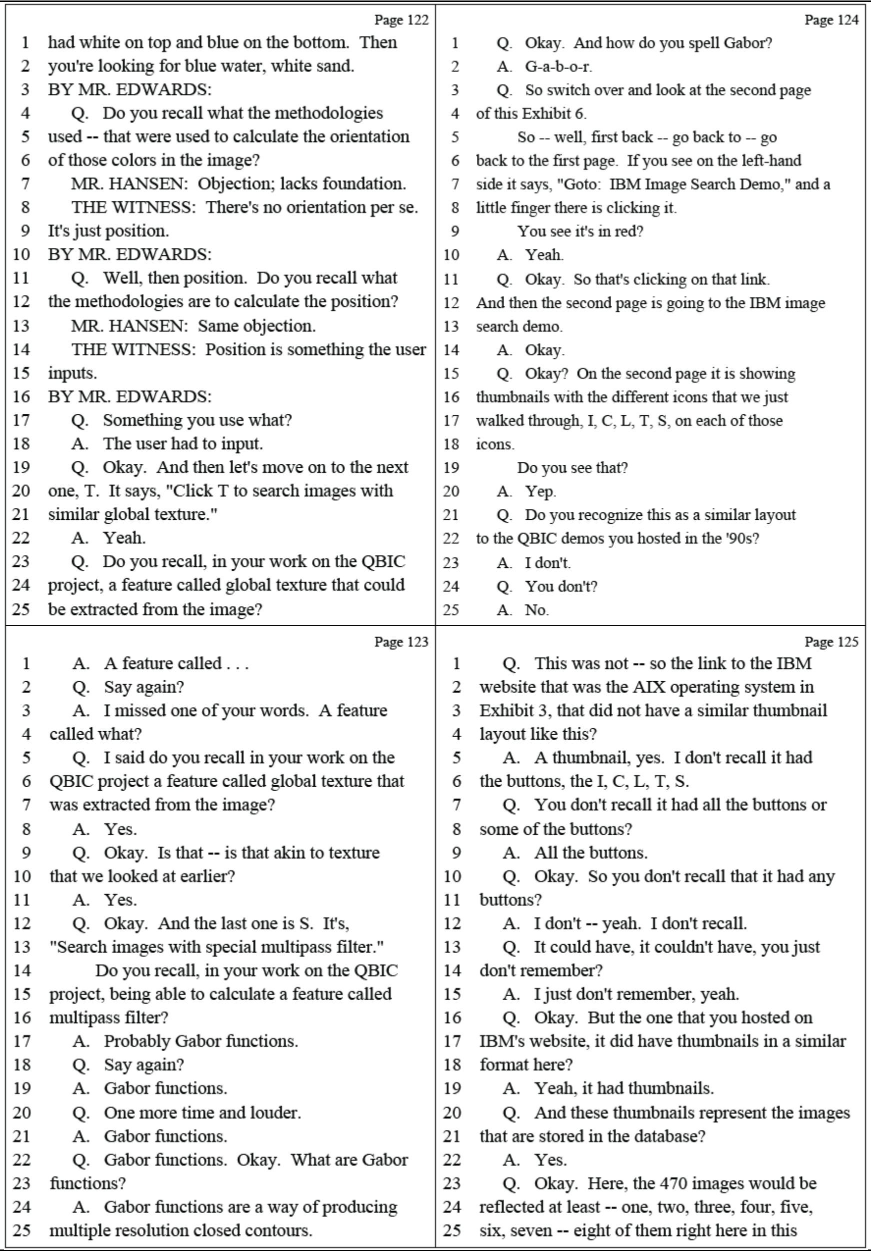

had white on top and blue on the bottom. Then you're looking for blue water, white sand.

BY MR. EDWARDS:

Q. Do you recall what the methodologies used -- that were used to calculate the orientation of those colors in the image?

MR. HANSEN: Objection; lacks foundation.

THE WITNESS: There's no orientation per se. It's just position.

BY MR. EDWARDS:

Q. Well, then position. Do you recall what the methodologies are to calculate the position?

MR. HANSEN: Same objection.

THE WITNESS: Position is something the user inputs.

BY MR. EDWARDS:

Q. Something you use what?

A. The user had to input.

Q. Okay. And then let's move on to the next one, T. It says, "Click T to search images with similar global texture."

A. Yeah.

Q. Do you recall, in your work on the QBIC project, a feature called global texture that could be extracted from the image?

A. A feature called . . .

Q. Say again?

A. I missed one of your words. A feature called what?

Q. I said do you recall in your work on the QBIC project a feature called global texture that was extracted from the image?

A. Yes.

Q. Okay. Is that -- is that akin to texture that we looked at earlier?

A. Yes.

Q. Okay. And the last one is S. It's, "Search images with special multipass filter."

Do you recall, in your work on the QBIC project, being able to calculate a feature called multipass filter?

A. Probably Gabor functions.

Q. Say again?

A. Gabor functions.

Q. One more time and louder.

A. Gabor functions.

Q. Gabor functions. Okay. What are Gabor functions?

A. Gabor functions are a way of producing multiple resolution closed contours.

Q. Okay. And how do you spell Gabor?

A. G-a-b-o-r.

Q. So switch over and look at the second page of this Exhibit 6.

So -- well, first back -- go back to -- go back to the first page. If you see on the left-hand side it says, "Goto: IBM Image Search Demo," and a little finger there is clicking it.

You see it's in red?

A. Yeah.

Q. Okay. So that's clicking on that link. And then the second page is going to the IBM image search demo.

A. Okay.

Q. Okay? On the second page it is showing thumbnails with the different icons that we just walked through, I, C, L, T, S, on each of those icons.

Do you see that?

A. Yep.

Q. Do you recognize this as a similar layout to the QBIC demos you hosted in the '90s?

A. I don't.

Q. You don't?

A. No.

Q. This was not -- so the link to the IBM website that was the AIX operating system in Exhibit 3, that did not have a similar thumbnail layout like this?

A. A thumbnail, yes. I don't recall it had the buttons, the I, C, L, T, S.

Q. You don't recall it had all the buttons or some of the buttons?

A. All the buttons.

Q. Okay. So you don't recall that it had any buttons?

A. I don't -- yeah. I don't recall.

Q. It could have, it couldn't have, you just don't remember?

A. I just don't remember, yeah.

Q. Okay. But the one that you hosted on IBM's website, it did have thumbnails in a similar format here?

A. Yeah, it had thumbnails.

Q. And these thumbnails represent the images that are stored in the database?

A. Yes.

Q. Okay. Here, the 470 images would be reflected at least -- one, two, three, four, five, six, seven -- eight of them right here in this

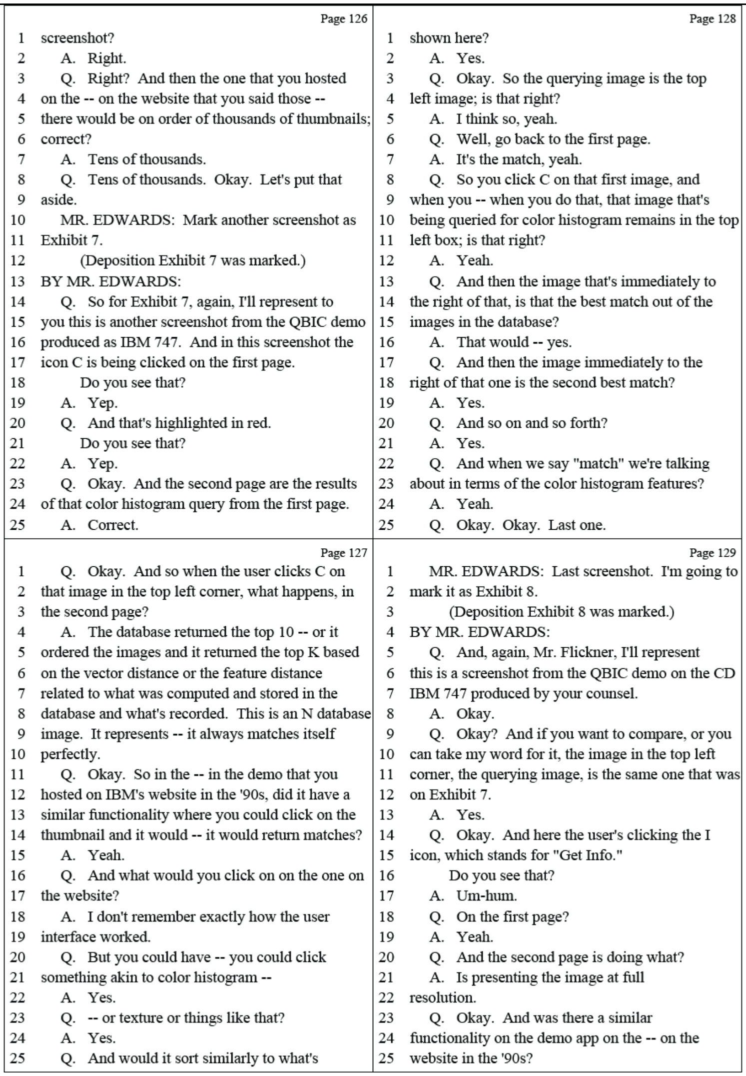

Page 126

screenshot?
A. Right.
Q. Right? And then the one that you hosted on the -- on the website that you said those -- there would be on order of thousands of thumbnails; correct?
A. Tens of thousands.
Q. Tens of thousands. Okay. Let's put that aside.
MR. EDWARDS: Mark another screenshot as Exhibit 7.
(Deposition Exhibit 7 was marked.)
BY MR. EDWARDS:
Q. So for Exhibit 7, again, I'll represent to you this is another screenshot from the QBIC demo produced as IBM 747. And in this screenshot the icon C is being clicked on the first page.
Do you see that?
A. Yep.
Q. And that's highlighted in red.
Do you see that?
A. Yep.
Q. Okay. And the second page are the results of that color histogram query from the first page.
A. Correct.

Page 127

Q. Okay. And so when the user clicks C on that image in the top left corner, what happens, in the second page?
A. The database returned the top 10 -- or it ordered the images and it returned the top K based on the vector distance or the feature distance related to what was computed and stored in the database and what's recorded. This is an N database image. It represents -- it always matches itself perfectly.
Q. Okay. So in the -- in the demo that you hosted on IBM's website in the '90s, did it have a similar functionality where you could click on the thumbnail and it would -- it would return matches?
A. Yeah.
Q. And what would you click on on the one on the website?
A. I don't remember exactly how the user interface worked.
Q. But you could have -- you could click something akin to color histogram --
A. Yes.
Q. -- or texture or things like that?
A. Yes.
Q. And would it sort similarly to what's

Page 128

shown here?
A. Yes.
Q. Okay. So the querying image is the top left image; is that right?
A. I think so, yeah.
Q. Well, go back to the first page.
A. It's the match, yeah.
Q. So you click C on that first image, and when you -- when you do that, that image that's being queried for color histogram remains in the top left box; is that right?
A. Yeah.
Q. And then the image that's immediately to the right of that, is that the best match out of the images in the database?
A. That would -- yes.
Q. And then the image immediately to the right of that one is the second best match?
A. Yes.
Q. And so on and so forth?
A. Yes.
Q. And when we say "match" we're talking about in terms of the color histogram features?
A. Yeah.
Q. Okay. Okay. Last one.

Page 129

MR. EDWARDS: Last screenshot. I'm going to mark it as Exhibit 8.
(Deposition Exhibit 8 was marked.)
BY MR. EDWARDS:
Q. And, again, Mr. Flickner, I'll represent this is a screenshot from the QBIC demo on the CD IBM 747 produced by your counsel.
A. Okay.
Q. Okay? And if you want to compare, or you can take my word for it, the image in the top left corner, the querying image, is the same one that was on Exhibit 7.
A. Yes.
Q. Okay. And here the user's clicking the I icon, which stands for "Get Info."
Do you see that?
A. Um-hum.
Q. On the first page?
A. Yeah.
Q. And the second page is doing what?
A. Is presenting the image at full resolution.
Q. Okay. And was there a similar functionality on the demo app on the -- on the website in the '90s?

Id. at 135 – 136:

		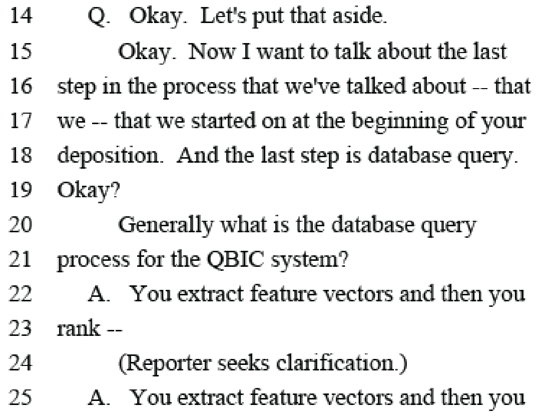 Q. Okay. Let's put that aside. Okay. Now I want to talk about the last step in the process that we've talked about -- that we -- that we started on at the beginning of your deposition. And the last step is database query. Okay? Generally what is the database query process for the QBIC system? A. You extract feature vectors and then you rank -- (Reporter seeks clarification.) A. You extract feature vectors and then you 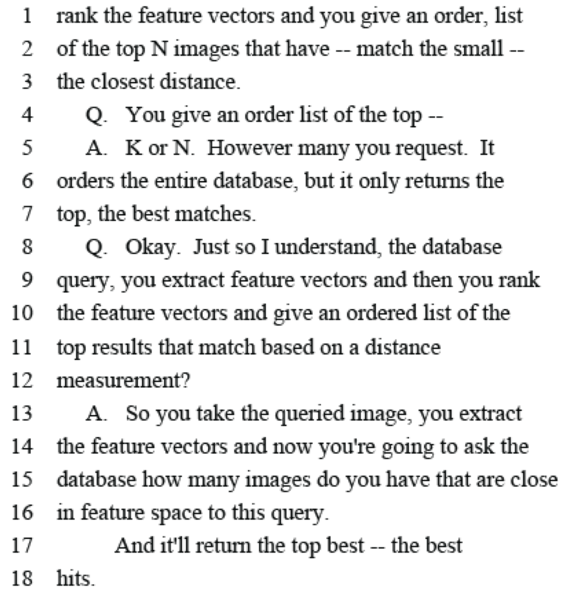 rank the feature vectors and you give an order, list of the top N images that have -- match the small -- the closest distance. Q. You give an order list of the top -- A. K or N. However many you request. It orders the entire database, but it only returns the top, the best matches. Q. Okay. Just so I understand, the database query, you extract feature vectors and then you rank the feature vectors and give an ordered list of the top results that match based on a distance measurement? A. So you take the queried image, you extract the feature vectors and now you're going to ask the database how many images do you have that are close in feature space to this query. And it'll return the top best -- the best hits. *Id.* at 139–165:

		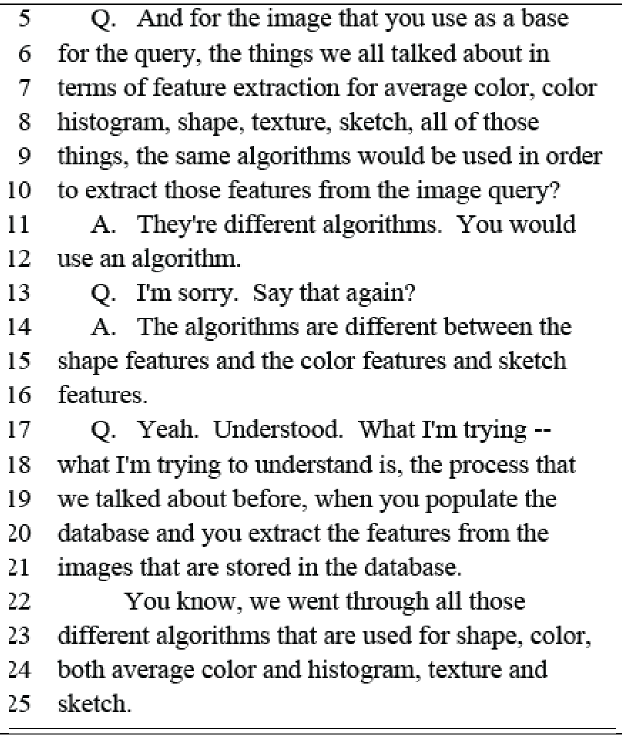 Q. And for the image that you use as a base for the query, the things we all talked about in terms of feature extraction for average color, color histogram, shape, texture, sketch, all of those things, the same algorithms would be used in order to extract those features from the image query? A. They're different algorithms. You would use an algorithm. Q. I'm sorry. Say that again? A. The algorithms are different between the shape features and the color features and sketch features. Q. Yeah. Understood. What I'm trying -- what I'm trying to understand is, the process that we talked about before, when you populate the database and you extract the features from the images that are stored in the database. You know, we went through all those different algorithms that are used for shape, color, both average color and histogram, texture and sketch.

		A. Um-hum. Q. Do you remember that? A. Yep. Q. Okay. That same process -- that exact same process for each of those features is used to extract those features from an image query? A. Correct. Q. Okay. And those similar -- so the algorithm, for example, for average color, that algorithm that's used to extract features for a database image is the same algorithm that's used to extract image for an image query? A. Yes. Q. Same for color histogram; correct? Same algorithm's used for both? A. It should be, yeah. Q. Same for shape and texture as well; correct? A. Right.

			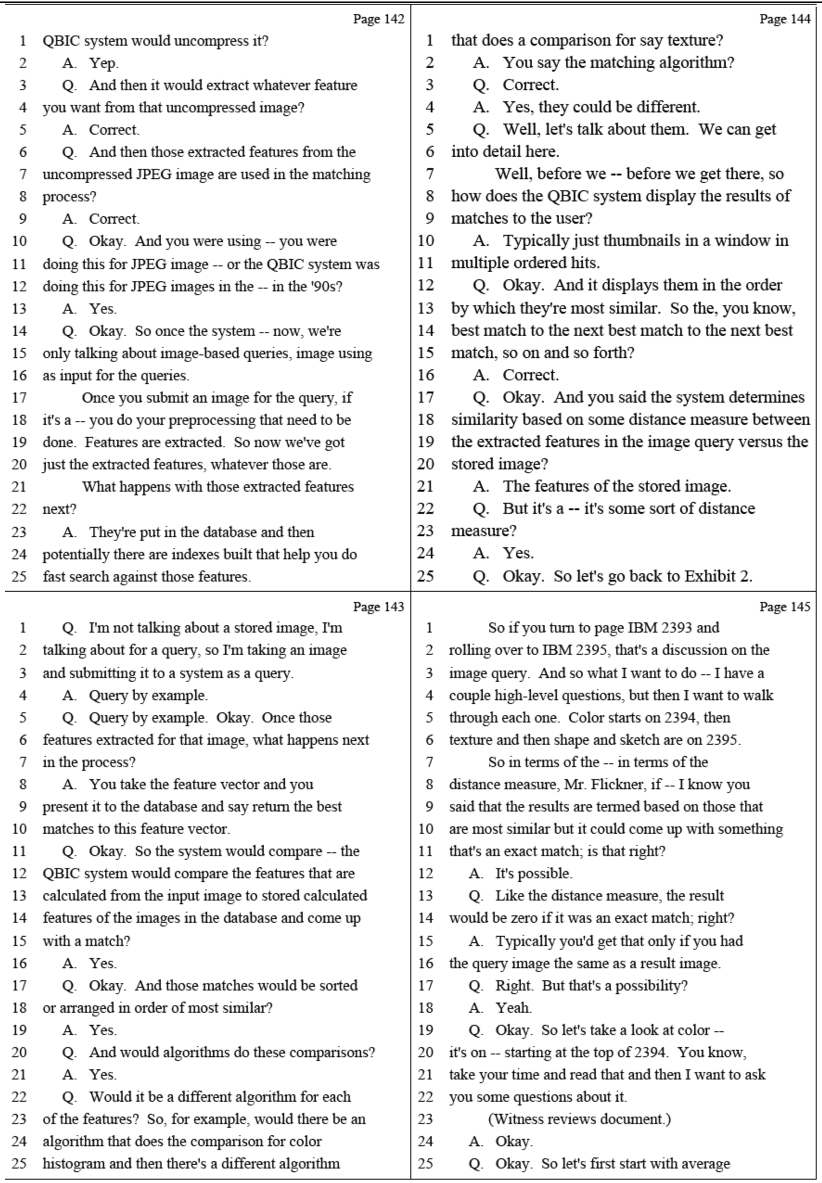

Page 142

QBIC system would uncompress it?
A. Yep.
Q. And then it would extract whatever feature you want from that uncompressed image?
A. Correct.
Q. And then those extracted features from the uncompressed JPEG image are used in the matching process?
A. Correct.
Q. Okay. And you were using -- you were doing this for JPEG image -- or the QBIC system was doing this for JPEG images in the -- in the '90s?
A. Yes.
Q. Okay. So once the system -- now, we're only talking about image-based queries, image using as input for the queries.
Once you submit an image for the query, if it's a -- you do your preprocessing that need to be done. Features are extracted. So now we've got just the extracted features, whatever those are.
What happens with those extracted features next?
A. They're put in the database and then potentially there are indexes built that help you do fast search against those features.

Page 143

Q. I'm not talking about a stored image, I'm talking about for a query, so I'm taking an image and submitting it to a system as a query.
A. Query by example.
Q. Query by example. Okay. Once those features extracted for that image, what happens next in the process?
A. You take the feature vector and you present it to the database and say return the best matches to this feature vector.
Q. Okay. So the system would compare -- the QBIC system would compare the features that are calculated from the input image to stored calculated features of the images in the database and come up with a match?
A. Yes.
Q. Okay. And those matches would be sorted or arranged in order of most similar?
A. Yes.
Q. And would algorithms do these comparisons?
A. Yes.
Q. Would it be a different algorithm for each of the features? So, for example, would there be an algorithm that does the comparison for color histogram and then there's a different algorithm

Page 144

that does a comparison for say texture?
A. You say the matching algorithm?
Q. Correct.
A. Yes, they could be different.
Q. Well, let's talk about them. We can get into detail here.
Well, before we -- before we get there, so how does the QBIC system display the results of matches to the user?
A. Typically just thumbnails in a window in multiple ordered hits.
Q. Okay. And it displays them in the order by which they're most similar. So the, you know, best match to the next best match to the next best match, so on and so forth?
A. Correct.
Q. Okay. And you said the system determines similarity based on some distance measure between the extracted features in the image query versus the stored image?
A. The features of the stored image.
Q. But it's a -- it's some sort of distance measure?
A. Yes.
Q. Okay. So let's go back to Exhibit 2.

Page 145

So if you turn to page IBM 2393 and rolling over to IBM 2395, that's a discussion on the image query. And so what I want to do -- I have a couple high-level questions, but then I want to walk through each one. Color starts on 2394, then texture and then shape and sketch are on 2395.
So in terms of the -- in terms of the distance measure, Mr. Flickner, if -- I know you said that the results are termed based on those that are most similar but it could come up with something that's an exact match; is that right?
A. It's possible.
Q. Like the distance measure, the result would be zero if it was an exact match; right?
A. Typically you'd get that only if you had the query image the same as a result image.
Q. Right. But that's a possibility?
A. Yeah.
Q. Okay. So let's take a look at color -- it's on -- starting at the top of 2394. You know, take your time and read that and then I want to ask you some questions about it.
(Witness reviews document.)
A. Okay.
Q. Okay. So let's first start with average

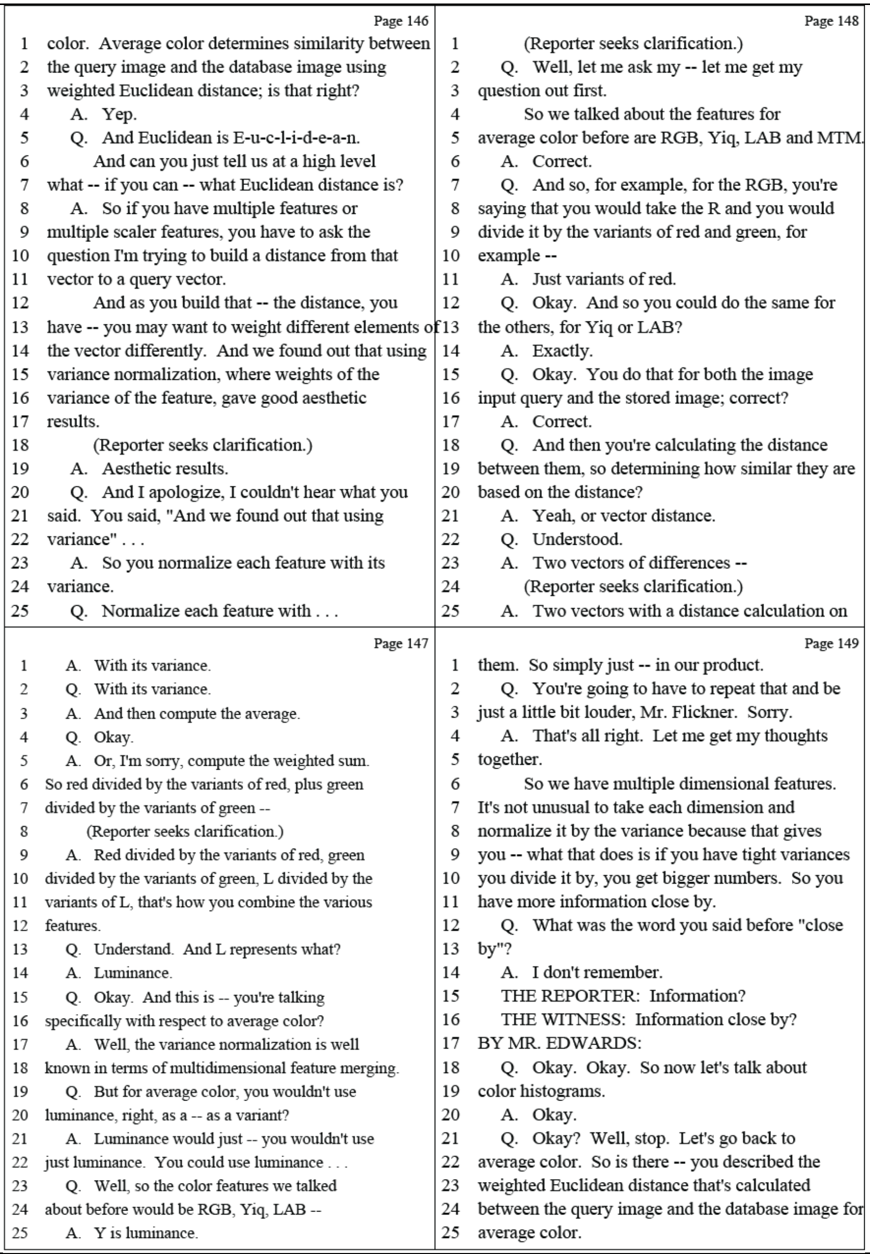

Page 146

color. Average color determines similarity between the query image and the database image using weighted Euclidean distance; is that right?

A. Yep.

Q. And Euclidean is E-u-c-l-i-d-e-a-n.

And can you just tell us at a high level what -- if you can -- what Euclidean distance is?

A. So if you have multiple features or multiple scaler features, you have to ask the question I'm trying to build a distance from that vector to a query vector.

And as you build that -- the distance, you have -- you may want to weight different elements of the vector differently. And we found out that using variance normalization, where weights of the variance of the feature, gave good aesthetic results.

(Reporter seeks clarification.)

A. Aesthetic results.

Q. And I apologize, I couldn't hear what you said. You said, "And we found out that using variance" . . .

A. So you normalize each feature with its variance.

Q. Normalize each feature with . . .

Page 147

A. With its variance.

Q. With its variance.

A. And then compute the average.

Q. Okay.

A. Or, I'm sorry, compute the weighted sum. So red divided by the variants of red, plus green divided by the variants of green --

(Reporter seeks clarification.)

A. Red divided by the variants of red, green divided by the variants of green, L divided by the variants of L, that's how you combine the various features.

Q. Understand. And L represents what?

A. Luminance.

Q. Okay. And this is -- you're talking specifically with respect to average color?

A. Well, the variance normalization is well known in terms of multidimensional feature merging.

Q. But for average color, you wouldn't use luminance, right, as a -- as a variant?

A. Luminance would just -- you wouldn't use just luminance. You could use luminance . . .

Q. Well, so the color features we talked about before would be RGB, Yiq, LAB --

A. Y is luminance.

Page 148

(Reporter seeks clarification.)

Q. Well, let me ask my -- let me get my question out first.

So we talked about the features for average color before are RGB, Yiq, LAB and MTM.

A. Correct.

Q. And so, for example, for the RGB, you're saying that you would take the R and you would divide it by the variants of red and green, for example --

A. Just variants of red.

Q. Okay. And so you could do the same for the others, for Yiq or LAB?

A. Exactly.

Q. Okay. You do that for both the image input query and the stored image; correct?

A. Correct.

Q. And then you're calculating the distance between them, so determining how similar they are based on the distance?

A. Yeah, or vector distance.

Q. Understood.

A. Two vectors of differences --

(Reporter seeks clarification.)

A. Two vectors with a distance calculation on

Page 149

them. So simply just -- in our product.

Q. You're going to have to repeat that and be just a little bit louder, Mr. Flickner. Sorry.

A. That's all right. Let me get my thoughts together.

So we have multiple dimensional features. It's not unusual to take each dimension and normalize it by the variance because that gives you -- what that does is if you have tight variances you divide it by, you get bigger numbers. So you have more information close by.

Q. What was the word you said before "close by"?

A. I don't remember.

THE REPORTER: Information?

THE WITNESS: Information close by?

BY MR. EDWARDS:

Q. Okay. Okay. So now let's talk about color histograms.

A. Okay.

Q. Okay? Well, stop. Let's go back to average color. So is there -- you described the weighted Euclidean distance that's calculated between the query image and the database image for average color.

Page 150

Is there any other detail that you recall that's not disclosed in the paragraph at IBM 2394?
A. Not that I recall.
Q. Okay. Do you recall whether this distance methodology changed after this article or it stayed the same?
A. It changed.
Q. Are you certain or do you know?
A. I'm not certain, but I'm pretty certain.
Q. And how did it change?
A. Different features were shaped like an inch in color. The algorithm evolved. It wasn't static.
Q. And how did it evolve?
A. I don't recall all the details.
Q. Do you recall whether it supplemented the weighted Euclidean distance or replaced it?
A. I don't recall.
Q. Is it fair to say there was always some distance measurement to determine similarity for average color in the QBIC system?
A. There's some measurement, yeah.
Q. Okay. Let's talk about color histogram.
So in order to determine the similarity between a query image and a database image, the QBIC

Page 151

system would use an algorithm that determines distance between normalized histograms?
A. Correct.
Q. Do you recall any details of how that worked?
A. Not off the top of my head.
Q. Do you recall whether that methodology evolved for matching color histograms after this article?
A. Not that I recall.
Q. Okay. All right. Let's talk about texture next. And I don't know if you've read that, but that's at the bottom of 2394 and it goes to just barely at the top of 2395.
(Witness reviews document.)
A. Okay.
Q. So for texture, the similarity is determined by the distance or the weighted Euclidean distance between the query object and -- the query object image and the database object image; is that right?
MR. HANSEN: Objection; vague and ambiguous.
THE WITNESS: Can you repeat.
BY MR. EDWARDS:
Q. Yeah. For texture, the similarity between

Page 152

the query image and the database image was determined by weighted Euclidean distance; is that right?
MR. HANSEN: Same objection.
THE WITNESS: Correct.
BY MR. EDWARDS:
Q. And how did -- how was that weighted Euclidean distance determined?
A. It was inverse variance.
(Reporter seeks clarification.)
A. Variance.
Q. And it was inverse variance for components like coarseness, contrast and directionality?
A. Correct.
Q. Okay. And do you recall if that methodology changed or stayed the same after this article for the QBIC system?
A. Most likely it evolved.
Q. Do you know for sure?
A. No.
Q. Do you recall any details?
A. No.
Q. All right. Let's talk about shape next.
The similarity between the query image and the database image also used weighted Euclidean

Page 153

distance for shape; correct?
A. Correct.
Q. And do you recall how that worked?
A. There's computing moments. And we used the shape measure, which I mentioned earlier, curvature and turning angle.
(Reporter seeks clarification.)
A. Curvature and turning angle.
Q. And the curvature and turning angle would be compared between the image query and the database image; correct?
A. Correct.
Q. And do you recall if that methodology changed after this article came out in '93?
A. Yes, it did.
Q. How so?
A. Well, in the IEEE computer paper, we described a different way of doing shape measures. And they included the moment calculations.
Q. And you use the moment calculations to determine a distance between the image query and the -- and the database image?
A. Yes.
Q. Sorry?
A. Yeah. Yes.

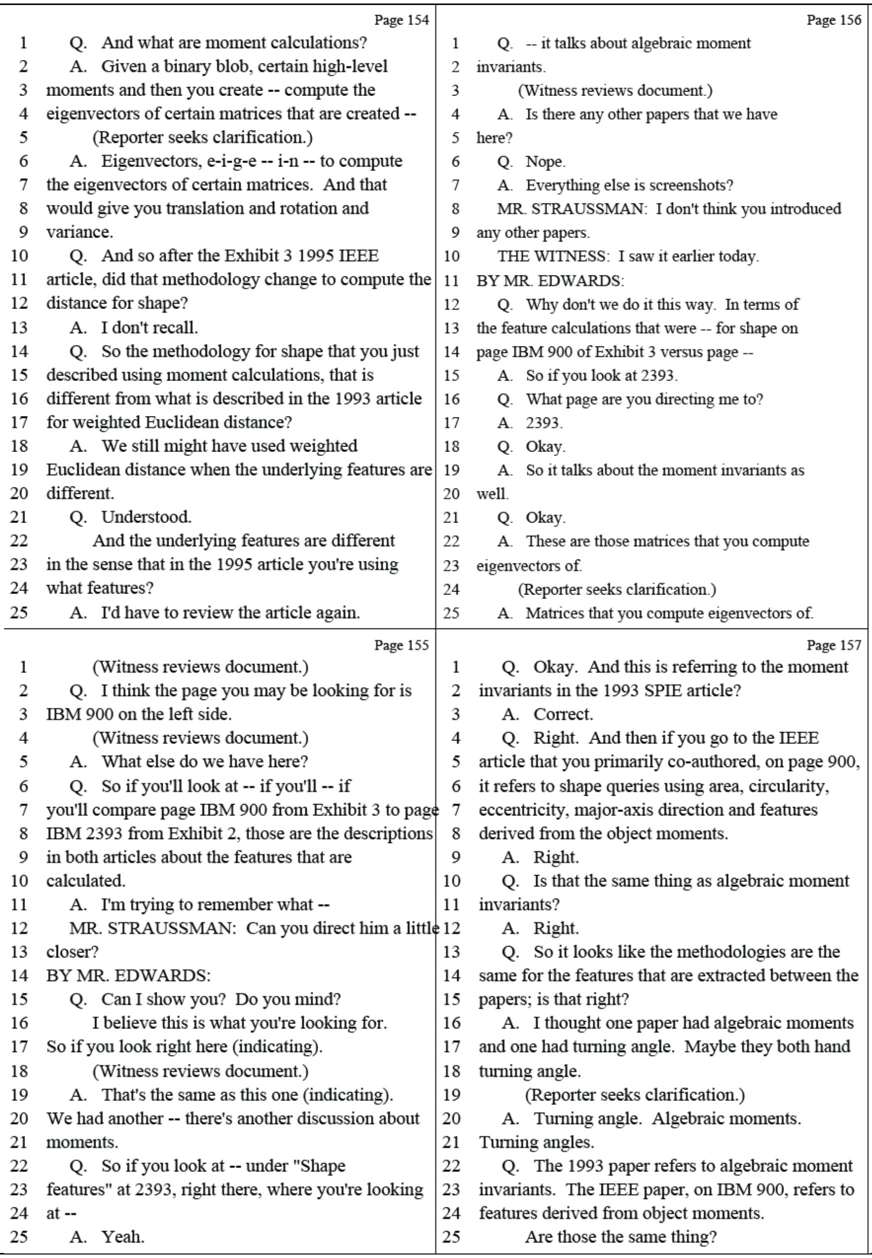

Page 154

Q. And what are moment calculations?
A. Given a binary blob, certain high-level moments and then you create -- compute the eigenvectors of certain matrices that are created --
(Reporter seeks clarification.)
A. Eigenvectors, e-i-g-e -- i-n -- to compute the eigenvectors of certain matrices. And that would give you translation and rotation and variance.
Q. And so after the Exhibit 3 1995 IEEE article, did that methodology change to compute the distance for shape?
A. I don't recall.
Q. So the methodology for shape that you just described using moment calculations, that is different from what is described in the 1993 article for weighted Euclidean distance?
A. We still might have used weighted Euclidean distance when the underlying features are different.
Q. Understood.
And the underlying features are different in the sense that in the 1995 article you're using what features?
A. I'd have to review the article again.

Page 155

(Witness reviews document.)
Q. I think the page you may be looking for is IBM 900 on the left side.
(Witness reviews document.)
A. What else do we have here?
Q. So if you'll look at -- if you'll -- if you'll compare page IBM 900 from Exhibit 3 to page IBM 2393 from Exhibit 2, those are the descriptions in both articles about the features that are calculated.
A. I'm trying to remember what --
MR. STRAUSSMAN: Can you direct him a little closer?
BY MR. EDWARDS:
Q. Can I show you? Do you mind?
I believe this is what you're looking for. So if you look right here (indicating).
(Witness reviews document.)
A. That's the same as this one (indicating). We had another -- there's another discussion about moments.
Q. So if you look at -- under "Shape features" at 2393, right there, where you're looking at --
A. Yeah.

Page 156

Q. -- it talks about algebraic moment invariants.
(Witness reviews document.)
A. Is there any other papers that we have here?
Q. Nope.
A. Everything else is screenshots?
MR. STRAUSSMAN: I don't think you introduced any other papers.
THE WITNESS: I saw it earlier today.
BY MR. EDWARDS:
Q. Why don't we do it this way. In terms of the feature calculations that were -- for shape on page IBM 900 of Exhibit 3 versus page --
A. So if you look at 2393.
Q. What page are you directing me to?
A. 2393.
Q. Okay.
A. So it talks about the moment invariants as well.
Q. Okay.
A. These are those matrices that you compute eigenvectors of.
(Reporter seeks clarification.)
A. Matrices that you compute eigenvectors of.

Page 157

Q. Okay. And this is referring to the moment invariants in the 1993 SPIE article?
A. Correct.
Q. Right. And then if you go to the IEEE article that you primarily co-authored, on page 900, it refers to shape queries using area, circularity, eccentricity, major-axis direction and features derived from the object moments.
A. Right.
Q. Is that the same thing as algebraic moment invariants?
A. Right.
Q. So it looks like the methodologies are the same for the features that are extracted between the papers; is that right?
A. I thought one paper had algebraic moments and one had turning angle. Maybe they both hand turning angle.
(Reporter seeks clarification.)
A. Turning angle. Algebraic moments. Turning angles.
Q. The 1993 paper refers to algebraic moment invariants. The IEEE paper, on IBM 900, refers to features derived from object moments.
Are those the same thing?

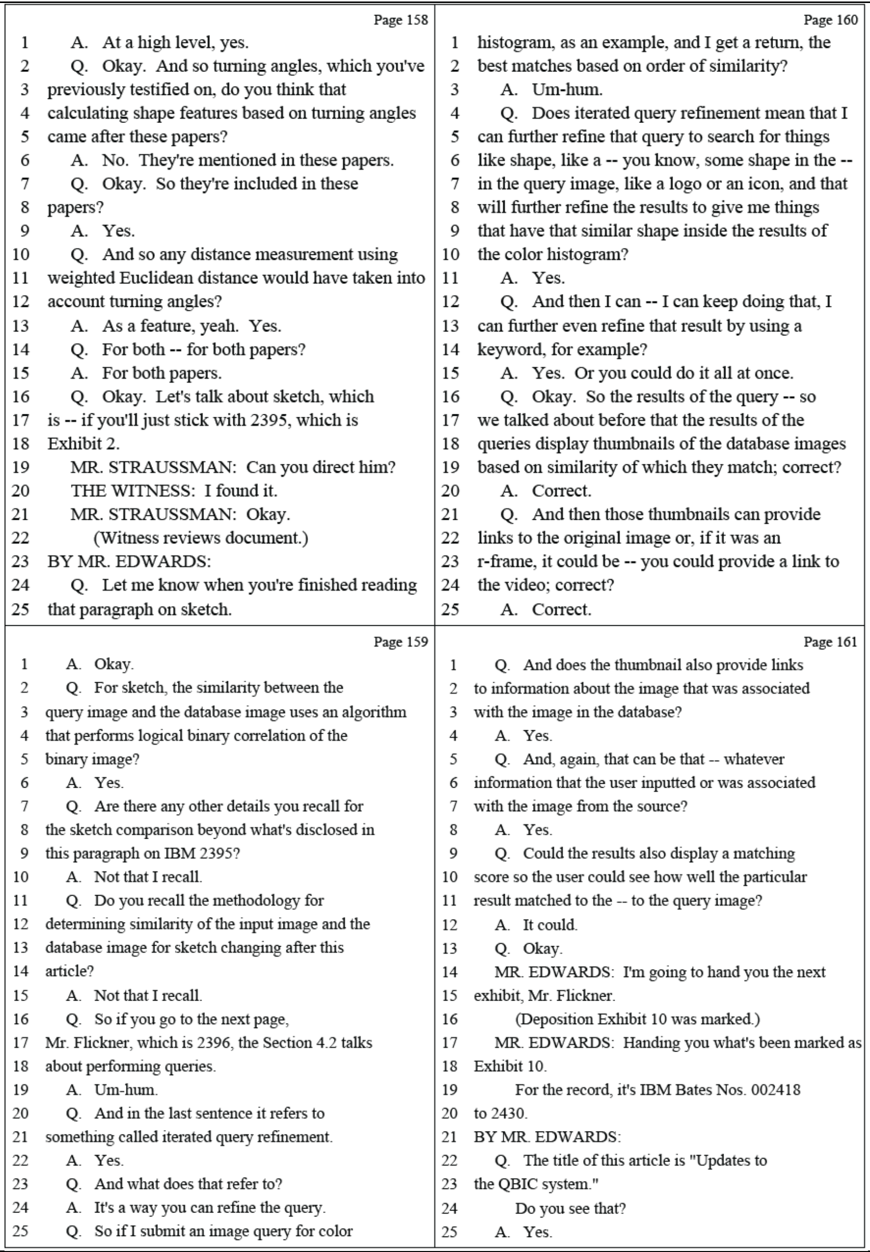

Page 158

A. At a high level, yes.

Q. Okay. And so turning angles, which you've previously testified on, do you think that calculating shape features based on turning angles came after these papers?

A. No. They're mentioned in these papers.

Q. Okay. So they're included in these papers?

A. Yes.

Q. And so any distance measurement using weighted Euclidean distance would have taken into account turning angles?

A. As a feature, yeah. Yes.

Q. For both -- for both papers?

A. For both papers.

Q. Okay. Let's talk about sketch, which is -- if you'll just stick with 2395, which is Exhibit 2.

MR. STRAUSSMAN: Can you direct him?

THE WITNESS: I found it.

MR. STRAUSSMAN: Okay.

(Witness reviews document.)

BY MR. EDWARDS:

Q. Let me know when you're finished reading that paragraph on sketch.

Page 159

A. Okay.

Q. For sketch, the similarity between the query image and the database image uses an algorithm that performs logical binary correlation of the binary image?

A. Yes.

Q. Are there any other details you recall for the sketch comparison beyond what's disclosed in this paragraph on IBM 2395?

A. Not that I recall.

Q. Do you recall the methodology for determining similarity of the input image and the database image for sketch changing after this article?

A. Not that I recall.

Q. So if you go to the next page, Mr. Flickner, which is 2396, the Section 4.2 talks about performing queries.

A. Um-hum.

Q. And in the last sentence it refers to something called iterated query refinement.

A. Yes.

Q. And what does that refer to?

A. It's a way you can refine the query.

Q. So if I submit an image query for color histogram, as an example, and I get a return, the best matches based on order of similarity?

Page 160

A. Um-hum.

Q. Does iterated query refinement mean that I can further refine that query to search for things like shape, like a -- you know, some shape in the -- in the query image, like a logo or an icon, and that will further refine the results to give me things that have that similar shape inside the results of the color histogram?

A. Yes.

Q. And then I can -- I can keep doing that, I can further even refine that result by using a keyword, for example?

A. Yes. Or you could do it all at once.

Q. Okay. So the results of the query -- so we talked about before that the results of the queries display thumbnails of the database images based on similarity of which they match; correct?

A. Correct.

Q. And then those thumbnails can provide links to the original image or, if it was an r-frame, it could be -- you could provide a link to the video; correct?

A. Correct.

Page 161

Q. And does the thumbnail also provide links to information about the image that was associated with the image in the database?

A. Yes.

Q. And, again, that can be that -- whatever information that the user inputted or was associated with the image from the source?

A. Yes.

Q. Could the results also display a matching score so the user could see how well the particular result matched to the -- to the query image?

A. It could.

Q. Okay.

MR. EDWARDS: I'm going to hand you the next exhibit, Mr. Flickner.

(Deposition Exhibit 10 was marked.)

MR. EDWARDS: Handing you what's been marked as Exhibit 10.

For the record, it's IBM Bates Nos. 002418 to 2430.

BY MR. EDWARDS:

Q. The title of this article is "Updates to the QBIC system."

Do you see that?

A. Yes.

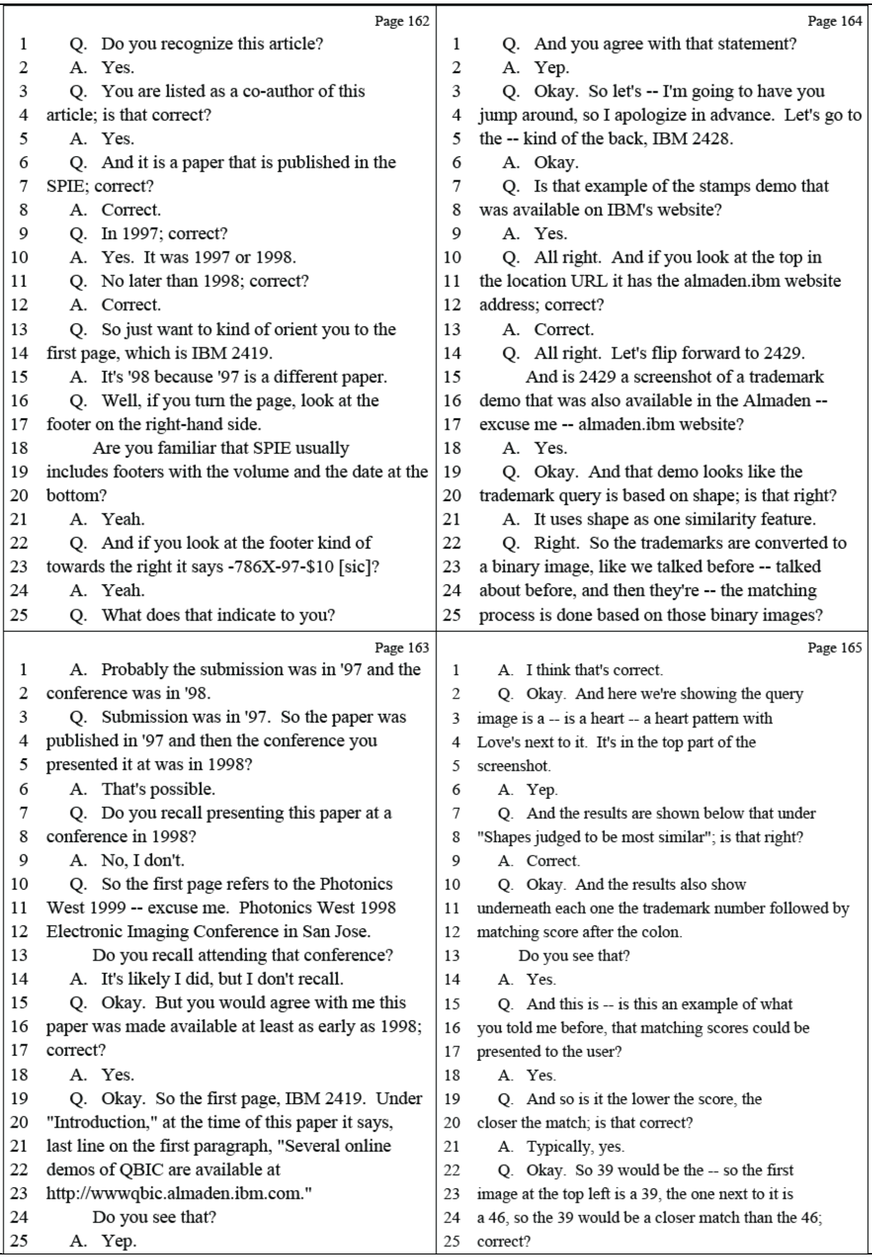

Page 162

Q. Do you recognize this article?
A. Yes.
Q. You are listed as a co-author of this article; is that correct?
A. Yes.
Q. And it is a paper that is published in the SPIE; correct?
A. Correct.
Q. In 1997; correct?
A. Yes. It was 1997 or 1998.
Q. No later than 1998; correct?
A. Correct.
Q. So just want to kind of orient you to the first page, which is IBM 2419.
A. It's '98 because '97 is a different paper.
Q. Well, if you turn the page, look at the footer on the right-hand side.
Are you familiar that SPIE usually includes footers with the volume and the date at the bottom?
A. Yeah.
Q. And if you look at the footer kind of towards the right it says -786X-97-$10 [sic]?
A. Yeah.
Q. What does that indicate to you?

Page 163

A. Probably the submission was in '97 and the conference was in '98.
Q. Submission was in '97. So the paper was published in '97 and then the conference you presented it at was in 1998?
A. That's possible.
Q. Do you recall presenting this paper at a conference in 1998?
A. No, I don't.
Q. So the first page refers to the Photonics West 1999 -- excuse me. Photonics West 1998 Electronic Imaging Conference in San Jose.
Do you recall attending that conference?
A. It's likely I did, but I don't recall.
Q. Okay. But you would agree with me this paper was made available at least as early as 1998; correct?
A. Yes.
Q. Okay. So the first page, IBM 2419. Under "Introduction," at the time of this paper it says, last line on the first paragraph, "Several online demos of QBIC are available at http://wwwqbic.almaden.ibm.com."
Do you see that?
A. Yep.

Page 164

Q. And you agree with that statement?
A. Yep.
Q. Okay. So let's -- I'm going to have you jump around, so I apologize in advance. Let's go to the -- kind of the back, IBM 2428.
A. Okay.
Q. Is that example of the stamps demo that was available on IBM's website?
A. Yes.
Q. All right. And if you look at the top in the location URL it has the almaden.ibm website address; correct?
A. Correct.
Q. All right. Let's flip forward to 2429.
And is 2429 a screenshot of a trademark demo that was also available in the Almaden -- excuse me -- almaden.ibm website?
A. Yes.
Q. Okay. And that demo looks like the trademark query is based on shape; is that right?
A. It uses shape as one similarity feature.
Q. Right. So the trademarks are converted to a binary image, like we talked before -- talked about before, and then they're -- the matching process is done based on those binary images?

Page 165

A. I think that's correct.
Q. Okay. And here we're showing the query image is a -- is a heart -- a heart pattern with Love's next to it. It's in the top part of the screenshot.
A. Yep.
Q. And the results are shown below that under "Shapes judged to be most similar"; is that right?
A. Correct.
Q. Okay. And the results also show underneath each one the trademark number followed by matching score after the colon.
Do you see that?
A. Yes.
Q. And this is -- is this an example of what you told me before, that matching scores could be presented to the user?
A. Yes.
Q. And so is it the lower the score, the closer the match; is that correct?
A. Typically, yes.
Q. Okay. So 39 would be the -- so the first image at the top left is a 39, the one next to it is a 46, so the 39 would be a closer match than the 46; correct?

Id. at 174–182:

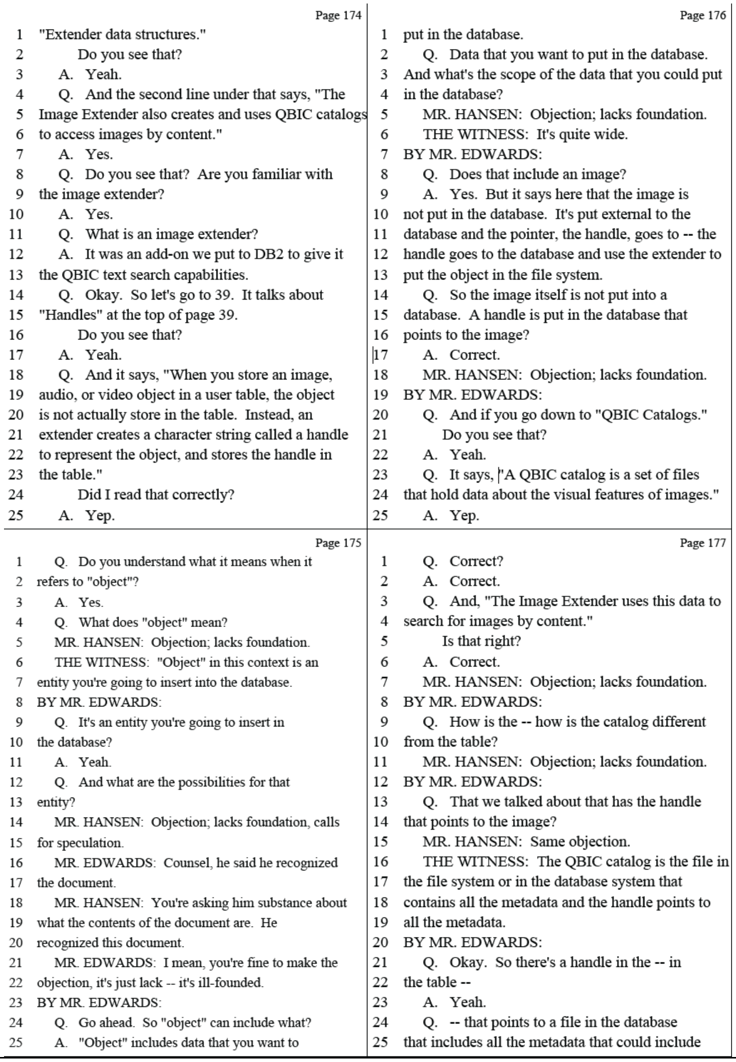

Page 174

"Extender data structures."
Do you see that?
A. Yeah.
Q. And the second line under that says, "The Image Extender also creates and uses QBIC catalogs to access images by content."
A. Yes.
Q. Do you see that? Are you familiar with the image extender?
A. Yes.
Q. What is an image extender?
A. It was an add-on we put to DB2 to give it the QBIC text search capabilities.
Q. Okay. So let's go to 39. It talks about "Handles" at the top of page 39.
Do you see that?
A. Yeah.
Q. And it says, "When you store an image, audio, or video object in a user table, the object is not actually store in the table. Instead, an extender creates a character string called a handle to represent the object, and stores the handle in the table."
Did I read that correctly?
A. Yep.

Page 175

Q. Do you understand what it means when it refers to "object"?
A. Yes.
Q. What does "object" mean?
MR. HANSEN: Objection; lacks foundation.
THE WITNESS: "Object" in this context is an entity you're going to insert into the database.
BY MR. EDWARDS:
Q. It's an entity you're going to insert in the database?
A. Yeah.
Q. And what are the possibilities for that entity?
MR. HANSEN: Objection; lacks foundation, calls for speculation.
MR. EDWARDS: Counsel, he said he recognized the document.
MR. HANSEN: You're asking him substance about what the contents of the document are. He recognized this document.
MR. EDWARDS: I mean, you're fine to make the objection, it's just lack -- it's ill-founded.
BY MR. EDWARDS:
Q. Go ahead. So "object" can include what?
A. "Object" includes data that you want to

Page 176

put in the database.
Q. Data that you want to put in the database. And what's the scope of the data that you could put in the database?
MR. HANSEN: Objection; lacks foundation.
THE WITNESS: It's quite wide.
BY MR. EDWARDS:
Q. Does that include an image?
A. Yes. But it says here that the image is not put in the database. It's put external to the database and the pointer, the handle, goes to -- the handle goes to the database and use the extender to put the object in the file system.
Q. So the image itself is not put into a database. A handle is put in the database that points to the image?
A. Correct.
MR. HANSEN: Objection; lacks foundation.
BY MR. EDWARDS:
Q. And if you go down to "QBIC Catalogs."
Do you see that?
A. Yeah.
Q. It says, "A QBIC catalog is a set of files that hold data about the visual features of images."
A. Yep.

Page 177

Q. Correct?
A. Correct.
Q. And, "The Image Extender uses this data to search for images by content."
Is that right?
A. Correct.
MR. HANSEN: Objection; lacks foundation.
BY MR. EDWARDS:
Q. How is the -- how is the catalog different from the table?
MR. HANSEN: Objection; lacks foundation.
BY MR. EDWARDS:
Q. That we talked about that has the handle that points to the image?
MR. HANSEN: Same objection.
THE WITNESS: The QBIC catalog is the file in the file system or in the database system that contains all the metadata and the handle points to all the metadata.
BY MR. EDWARDS:
Q. Okay. So there's a handle in the -- in the table --
A. Yeah.
Q. -- that points to a file in the database that includes all the metadata that could include

			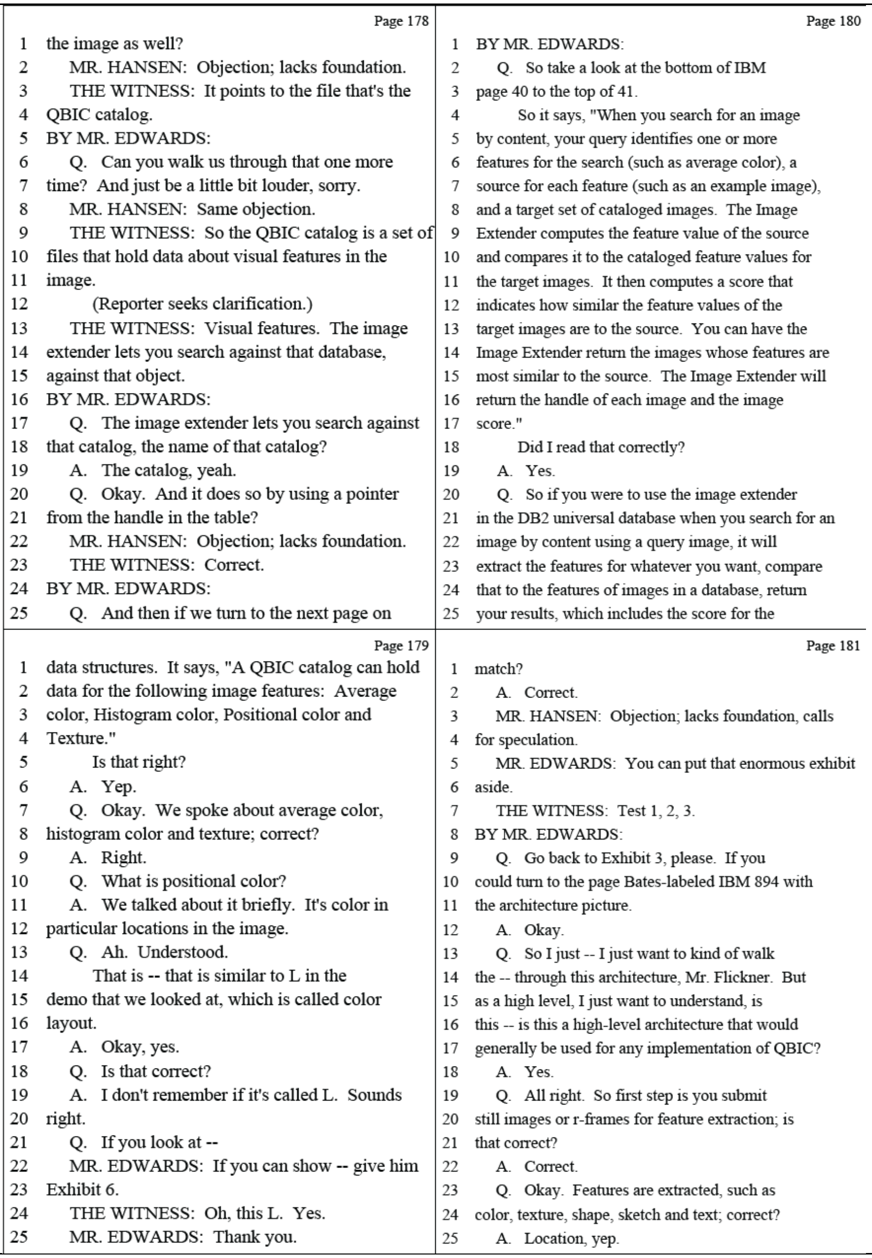

Page 178

the image as well?
MR. HANSEN: Objection; lacks foundation.
THE WITNESS: It points to the file that's the QBIC catalog.
BY MR. EDWARDS:
Q. Can you walk us through that one more time? And just be a little bit louder, sorry.
MR. HANSEN: Same objection.
THE WITNESS: So the QBIC catalog is a set of files that hold data about visual features in the image.
(Reporter seeks clarification.)
THE WITNESS: Visual features. The image extender lets you search against that database, against that object.
BY MR. EDWARDS:
Q. The image extender lets you search against that catalog, the name of that catalog?
A. The catalog, yeah.
Q. Okay. And it does so by using a pointer from the handle in the table?
MR. HANSEN: Objection; lacks foundation.
THE WITNESS: Correct.
BY MR. EDWARDS:
Q. And then if we turn to the next page on

Page 179

data structures. It says, "A QBIC catalog can hold data for the following image features: Average color, Histogram color, Positional color and Texture."
Is that right?
A. Yep.
Q. Okay. We spoke about average color, histogram color and texture; correct?
A. Right.
Q. What is positional color?
A. We talked about it briefly. It's color in particular locations in the image.
Q. Ah. Understood.
That is -- that is similar to L in the demo that we looked at, which is called color layout.
A. Okay, yes.
Q. Is that correct?
A. I don't remember if it's called L. Sounds right.
Q. If you look at --
MR. EDWARDS: If you can show -- give him Exhibit 6.
THE WITNESS: Oh, this L. Yes.
MR. EDWARDS: Thank you.

Page 180

BY MR. EDWARDS:
Q. So take a look at the bottom of IBM page 40 to the top of 41.
So it says, "When you search for an image by content, your query identifies one or more features for the search (such as average color), a source for each feature (such as an example image), and a target set of cataloged images. The Image Extender computes the feature value of the source and compares it to the cataloged feature values for the target images. It then computes a score that indicates how similar the feature values of the target images are to the source. You can have the Image Extender return the images whose features are most similar to the source. The Image Extender will return the handle of each image and the image score."
Did I read that correctly?
A. Yes.
Q. So if you were to use the image extender in the DB2 universal database when you search for an image by content using a query image, it will extract the features for whatever you want, compare that to the features of images in a database, return your results, which includes the score for the

Page 181

match?
A. Correct.
MR. HANSEN: Objection; lacks foundation, calls for speculation.
MR. EDWARDS: You can put that enormous exhibit aside.
THE WITNESS: Test 1, 2, 3.
BY MR. EDWARDS:
Q. Go back to Exhibit 3, please. If you could turn to the page Bates-labeled IBM 894 with the architecture picture.
A. Okay.
Q. So I just -- I just want to kind of walk the -- through this architecture, Mr. Flickner. But as a high level, I just want to understand, is this -- is this a high-level architecture that would generally be used for any implementation of QBIC?
A. Yes.
Q. All right. So first step is you submit still images or r-frames for feature extraction; is that correct?
A. Correct.
Q. Okay. Features are extracted, such as color, texture, shape, sketch and text; correct?
A. Location, yep.

		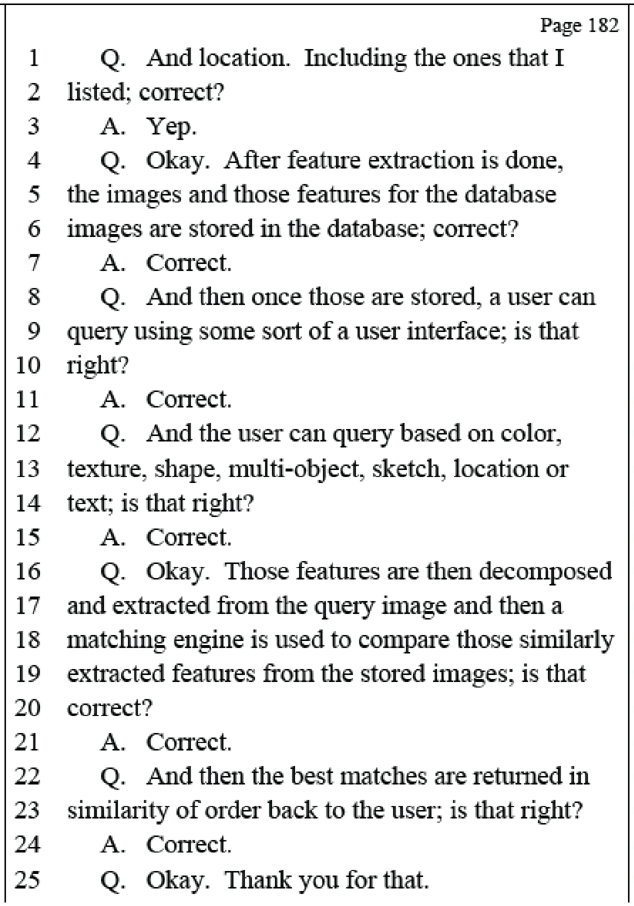 Page 182 Q. And location. Including the ones that I listed; correct? A. Yep. Q. Okay. After feature extraction is done, the images and those features for the database images are stored in the database; correct? A. Correct. Q. And then once those are stored, a user can query using some sort of a user interface; is that right? A. Correct. Q. And the user can query based on color, texture, shape, multi-object, sketch, location or text; is that right? A. Correct. Q. Okay. Those features are then decomposed and extracted from the query image and then a matching engine is used to compare those similarly extracted features from the stored images; is that correct? A. Correct. Q. And then the best matches are returned in similarity of order back to the user; is that right? A. Correct. Q. Okay. Thank you for that. *Id.* at 190–191: 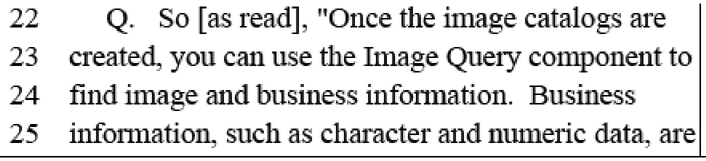 Q. So [as read], "Once the image catalogs are created, you can use the Image Query component to find image and business information. Business information, such as character and numeric data, are

		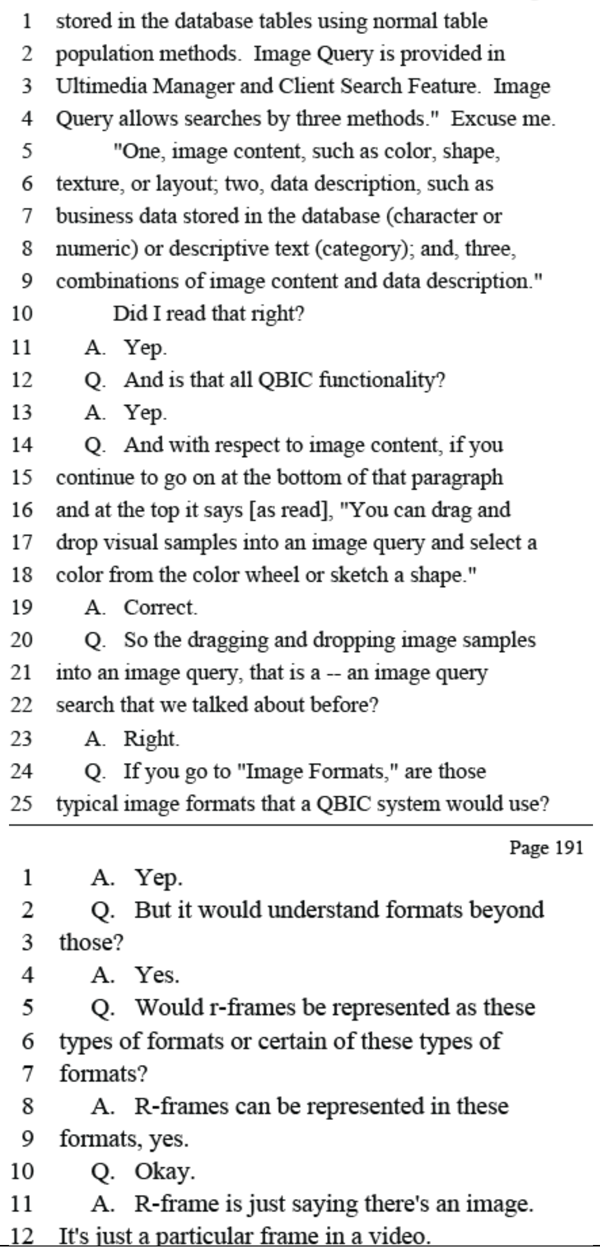 stored in the database tables using normal table population methods. Image Query is provided in Ultimedia Manager and Client Search Feature. Image Query allows searches by three methods." Excuse me. "One, image content, such as color, shape, texture, or layout; two, data description, such as business data stored in the database (character or numeric) or descriptive text (category); and, three, combinations of image content and data description." Did I read that right? A. Yep. Q. And is that all QBIC functionality? A. Yep. Q. And with respect to image content, if you continue to go on at the bottom of that paragraph and at the top it says [as read], "You can drag and drop visual samples into an image query and select a color from the color wheel or sketch a shape." A. Correct. Q. So the dragging and dropping image samples into an image query, that is a -- an image query search that we talked about before? A. Right. Q. If you go to "Image Formats," are those typical image formats that a QBIC system would use? Page 191 A. Yep. Q. But it would understand formats beyond those? A. Yes. Q. Would r-frames be represented as these types of formats or certain of these types of formats? A. R-frames can be represented in these formats, yes. Q. Okay. A. R-frame is just saying there's an image. It's just a particular frame in a video.

		Id. at 198 – 199: 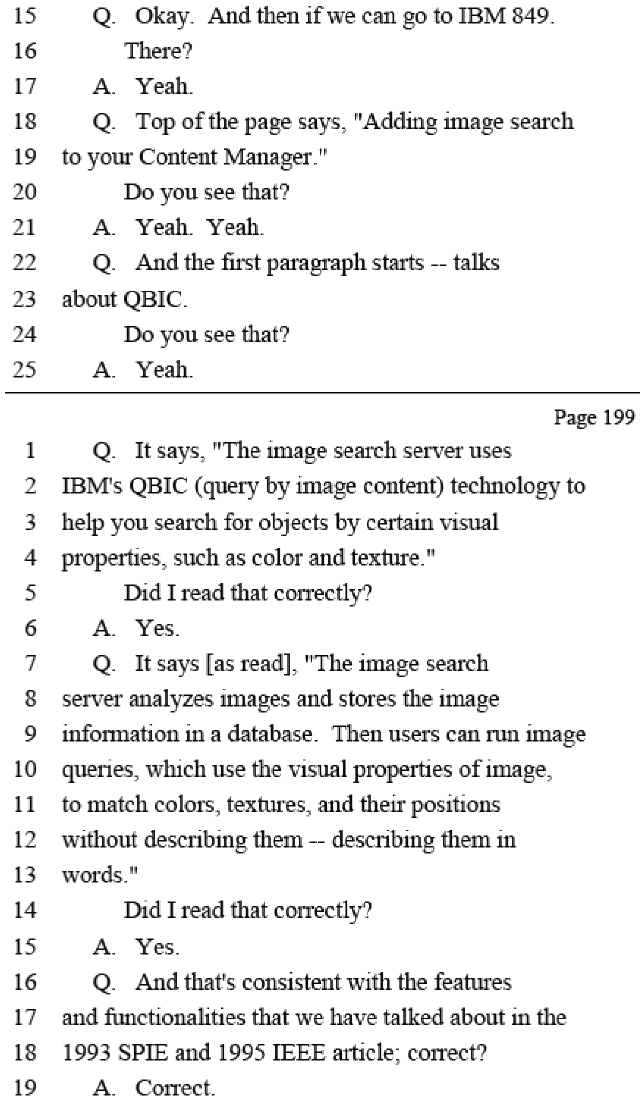 Q. Okay. And then if we can go to IBM 849. There? A. Yeah. Q. Top of the page says, "Adding image search to your Content Manager." Do you see that? A. Yeah. Yeah. Q. And the first paragraph starts -- talks about QBIC. Do you see that? A. Yeah. Page 199 Q. It says, "The image search server uses IBM's QBIC (query by image content) technology to help you search for objects by certain visual properties, such as color and texture." Did I read that correctly? A. Yes. Q. It says [as read], "The image search server analyzes images and stores the image information in a database. Then users can run image queries, which use the visual properties of image, to match colors, textures, and their positions without describing them -- describing them in words." Did I read that correctly? A. Yes. Q. And that's consistent with the features and functionalities that we have talked about in the 1993 SPIE and 1995 IEEE article; correct? A. Correct. *Id.* at 305:

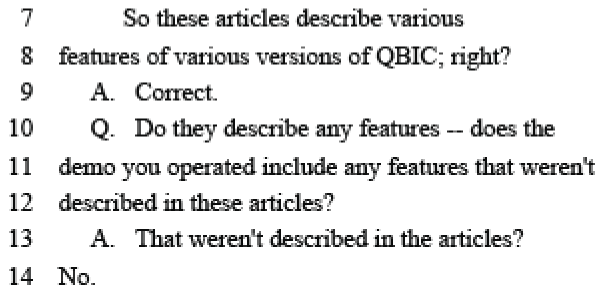
So these articles describe various features of various versions of QBIC; right?
A. Correct.
Q. Do they describe any features -- does the demo you operated include any features that weren't described in these articles?
A. That weren't described in the articles? No.

Id. at 35:

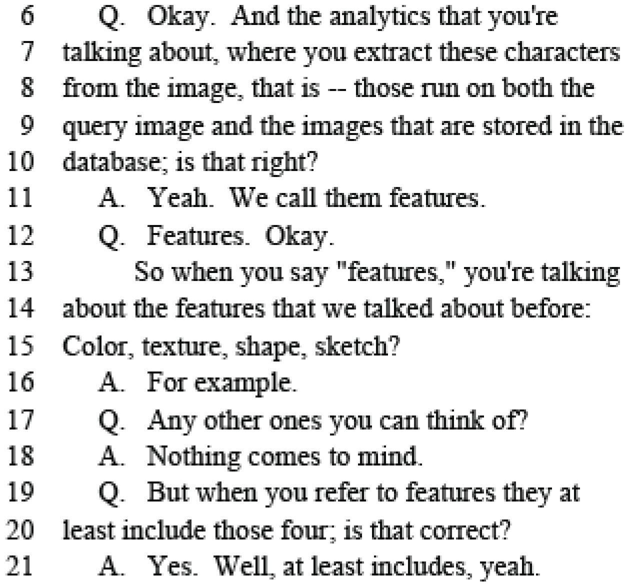
Q. Okay. And the analytics that you're talking about, where you extract these characters from the image, that is -- those run on both the query image and the images that are stored in the database; is that right?
A. Yeah. We call them features.
Q. Features. Okay.
So when you say "features," you're talking about the features that we talked about before: Color, texture, shape, sketch?
A. For example.
Q. Any other ones you can think of?
A. Nothing comes to mind.
Q. But when you refer to features they at least include those four; is that correct?
A. Yes. Well, at least includes, yeah.

Id. at 109-110:

		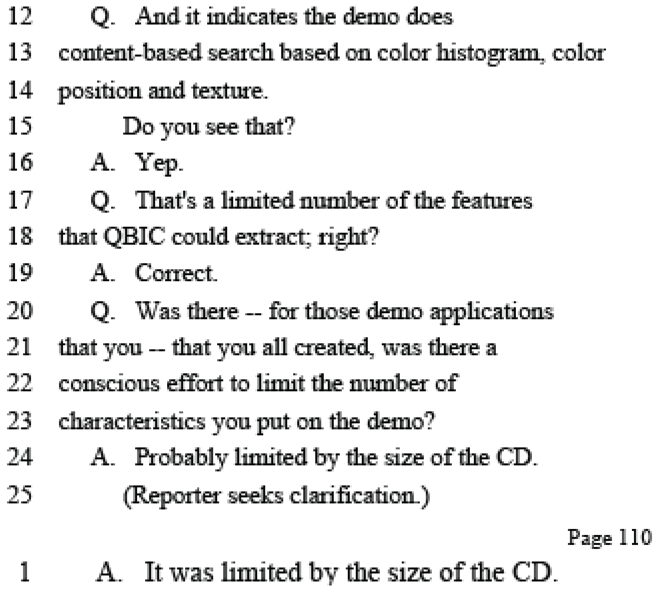 Q. And it indicates the demo does content-based search based on color histogram, color position and texture. Do you see that? A. Yep. Q. That's a limited number of the features that QBIC could extract; right? A. Correct. Q. Was there -- for those demo applications that you -- that you all created, was there a conscious effort to limit the number of characteristics you put on the demo? A. Probably limited by the size of the CD. (Reporter seeks clarification.) Page 110 A. It was limited by the size of the CD. *Id.* at 311: 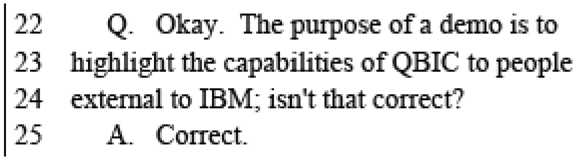 Q. Okay. The purpose of a demo is to highlight the capabilities of QBIC to people external to IBM; isn't that correct? A. Correct. *Id.* at 312 – 313:

		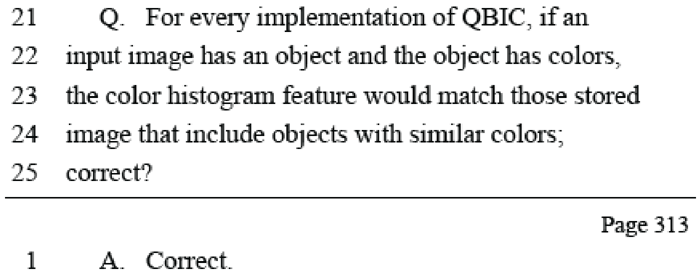 Q. For every implementation of QBIC, if an input image has an object and the object has colors, the color histogram feature would match those stored image that include objects with similar colors; correct? Page 313 A. Correct. *Id.* at 313-317: 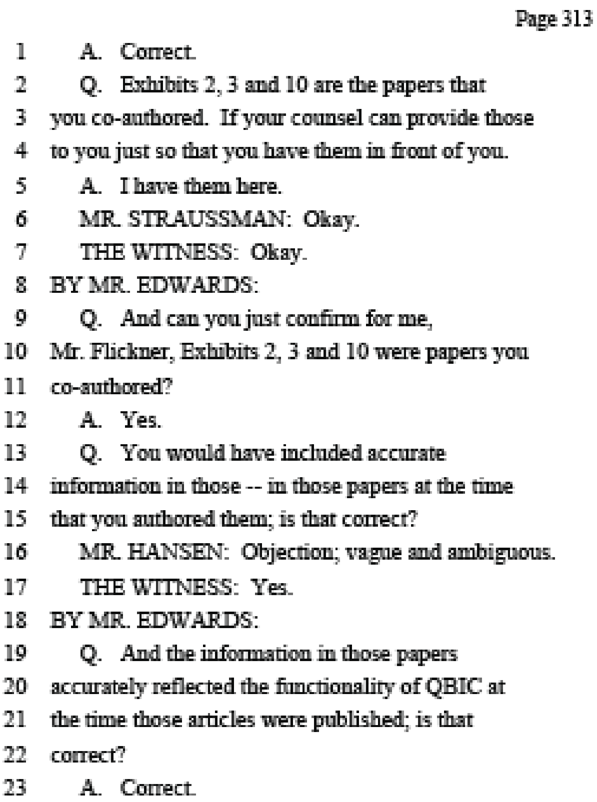 Page 313 A. Correct. Q. Exhibits 2, 3 and 10 are the papers that you co-authored. If your counsel can provide those to you just so that you have them in front of you. A. I have them here. MR. STRAUSSMAN: Okay. THE WITNESS: Okay. BY MR. EDWARDS: Q. And can you just confirm for me, Mr. Flickner, Exhibits 2, 3 and 10 were papers you co-authored? A. Yes. Q. You would have included accurate information in those -- in those papers at the time that you authored them; is that correct? MR. HANSEN: Objection; vague and ambiguous. THE WITNESS: Yes. BY MR. EDWARDS: Q. And the information in those papers accurately reflected the functionality of QBIC at the time those articles were published; is that correct? A. Correct.

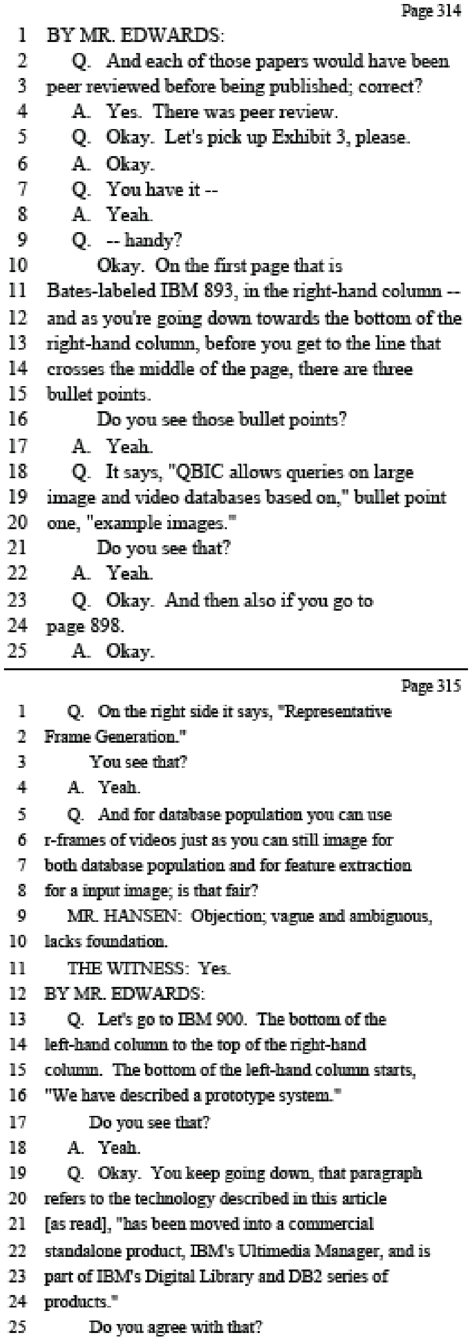

Page 314

BY MR. EDWARDS:
Q. And each of those papers would have been peer reviewed before being published; correct?
A. Yes. There was peer review.
Q. Okay. Let's pick up Exhibit 3, please.
A. Okay.
Q. You have it --
A. Yeah.
Q. -- handy?
Okay. On the first page that is Bates-labeled IBM 893, in the right-hand column -- and as you're going down towards the bottom of the right-hand column, before you get to the line that crosses the middle of the page, there are three bullet points.
Do you see those bullet points?
A. Yeah.
Q. It says, "QBIC allows queries on large image and video databases based on," bullet point one, "example images."
Do you see that?
A. Yeah.
Q. Okay. And then also if you go to page 898.
A. Okay.

Page 315

Q. On the right side it says, "Representative Frame Generation."
You see that?
A. Yeah.
Q. And for database population you can use r-frames of videos just as you can still image for both database population and for feature extraction for a input image; is that fair?
MR. HANSEN: Objection; vague and ambiguous, lacks foundation.
THE WITNESS: Yes.
BY MR. EDWARDS:
Q. Let's go to IBM 900. The bottom of the left-hand column to the top of the right-hand column. The bottom of the left-hand column starts, "We have described a prototype system."
Do you see that?
A. Yeah.
Q. Okay. You keep going down, that paragraph refers to the technology described in this article [as read], "has been moved into a commercial standalone product, IBM's Ultimedia Manager, and is part of IBM's Digital Library and DB2 series of products."
Do you agree with that?

		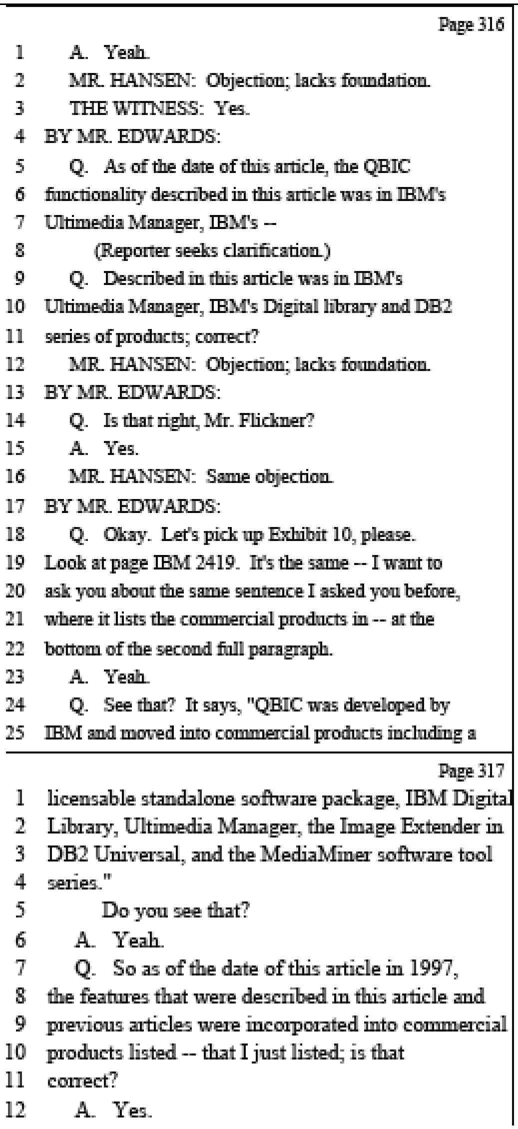 Page 316 A. Yeah. MR. HANSEN: Objection; lacks foundation. THE WITNESS: Yes. BY MR. EDWARDS: Q. As of the date of this article, the QBIC functionality described in this article was in IBM's Ultimedia Manager, IBM's -- (Reporter seeks clarification.) Q. Described in this article was in IBM's Ultimedia Manager, IBM's Digital library and DB2 series of products; correct? MR. HANSEN: Objection; lacks foundation. BY MR. EDWARDS: Q. Is that right, Mr. Flickner? A. Yes. MR. HANSEN: Same objection. BY MR. EDWARDS: Q. Okay. Let's pick up Exhibit 10, please. Look at page IBM 2419. It's the same -- I want to ask you about the same sentence I asked you before, where it lists the commercial products in -- at the bottom of the second full paragraph. A. Yeah. Q. See that? It says, "QBIC was developed by IBM and moved into commercial products including a Page 317 licensable standalone software package, IBM Digital Library, Ultimedia Manager, the Image Extender in DB2 Universal, and the MediaMiner software tool series." Do you see that? A. Yeah. Q. So as of the date of this article in 1997, the features that were described in this article and previous articles were incorporated into commercial products listed -- that I just listed; is that correct? A. Yes.

		Id. at 317–322: 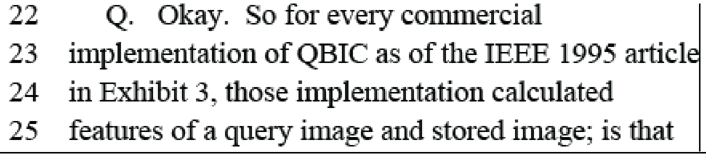 Q. Okay. So for every commercial implementation of QBIC as of the IEEE 1995 article in Exhibit 3, those implementation calculated features of a query image and stored image; is that

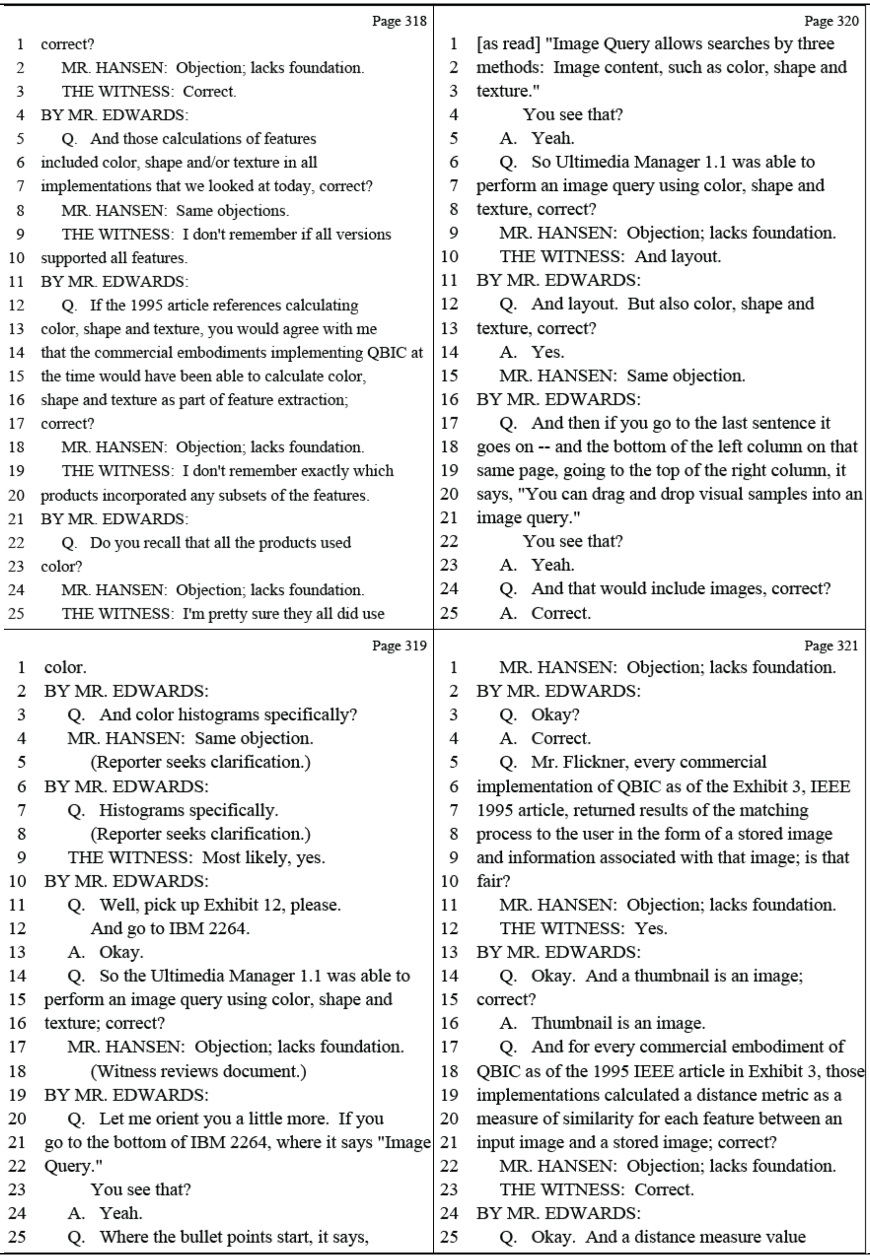

Page 318

correct?
MR. HANSEN: Objection; lacks foundation.
THE WITNESS: Correct.
BY MR. EDWARDS:
Q. And those calculations of features included color, shape and/or texture in all implementations that we looked at today, correct?
MR. HANSEN: Same objections.
THE WITNESS: I don't remember if all versions supported all features.
BY MR. EDWARDS:
Q. If the 1995 article references calculating color, shape and texture, you would agree with me that the commercial embodiments implementing QBIC at the time would have been able to calculate color, shape and texture as part of feature extraction; correct?
MR. HANSEN: Objection; lacks foundation.
THE WITNESS: I don't remember exactly which products incorporated any subsets of the features.
BY MR. EDWARDS:
Q. Do you recall that all the products used color?
MR. HANSEN: Objection; lacks foundation.
THE WITNESS: I'm pretty sure they all did use

Page 319

color.
BY MR. EDWARDS:
Q. And color histograms specifically?
MR. HANSEN: Same objection.
(Reporter seeks clarification.)
BY MR. EDWARDS:
Q. Histograms specifically.
(Reporter seeks clarification.)
THE WITNESS: Most likely, yes.
BY MR. EDWARDS:
Q. Well, pick up Exhibit 12, please.
And go to IBM 2264.
A. Okay.
Q. So the Ultimedia Manager 1.1 was able to perform an image query using color, shape and texture; correct?
MR. HANSEN: Objection; lacks foundation.
(Witness reviews document.)
BY MR. EDWARDS:
Q. Let me orient you a little more. If you go to the bottom of IBM 2264, where it says "Image Query."
You see that?
A. Yeah.
Q. Where the bullet points start, it says,

Page 320

[as read] "Image Query allows searches by three methods: Image content, such as color, shape and texture."
You see that?
A. Yeah.
Q. So Ultimedia Manager 1.1 was able to perform an image query using color, shape and texture, correct?
MR. HANSEN: Objection; lacks foundation.
THE WITNESS: And layout.
BY MR. EDWARDS:
Q. And layout. But also color, shape and texture, correct?
A. Yes.
MR. HANSEN: Same objection.
BY MR. EDWARDS:
Q. And then if you go to the last sentence it goes on -- and the bottom of the left column on that same page, going to the top of the right column, it says, "You can drag and drop visual samples into an image query."
You see that?
A. Yeah.
Q. And that would include images, correct?
A. Correct.

Page 321

MR. HANSEN: Objection; lacks foundation.
BY MR. EDWARDS:
Q. Okay?
A. Correct.
Q. Mr. Flickner, every commercial implementation of QBIC as of the Exhibit 3, IEEE 1995 article, returned results of the matching process to the user in the form of a stored image and information associated with that image; is that fair?
MR. HANSEN: Objection; lacks foundation.
THE WITNESS: Yes.
BY MR. EDWARDS:
Q. Okay. And a thumbnail is an image; correct?
A. Thumbnail is an image.
Q. And for every commercial embodiment of QBIC as of the 1995 IEEE article in Exhibit 3, those implementations calculated a distance metric as a measure of similarity for each feature between an input image and a stored image; correct?
MR. HANSEN: Objection; lacks foundation.
THE WITNESS: Correct.
BY MR. EDWARDS:
Q. Okay. And a distance measure value

		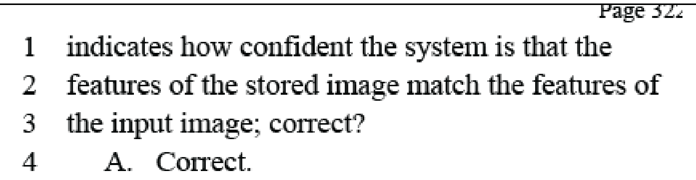Page 322 indicates how confident the system is that the features of the stored image match the features of the input image; correct? A. Correct. Flickner 8/29/23 Depo at 68–72:

		Q. Okay. And in what form would those results be displayed to me? For images would they be displayed as a thumbnail? A. It could be. Page 69 Q. And what about r-frames, how would those be displayed to you? A. It could be as a thumbnail or as a short video clip or as a single representative frame. Q. Okay. So they could be -- so for a video, there's a video that matched the image query of an r-frame, the results I could get back could include three things. It could be a still image, it could be a thumbnail, or it could be an r-frame representation? A. R-frame, I think, is representative frame. So it was a key frame in that video sequence, summary frame. Q. Once I ran a query on an r-frame; right? Well, let's stop there. If I run a query on still images, typically the results I would get back are thumbnails that are most similar to the characteristics which I queried from the query image; right? A. You could get thumbnails back. You could also get the full-sized image back. Q. Okay. A. Thumbnail is a tool for helping you return lots of results and quickly scan it.	

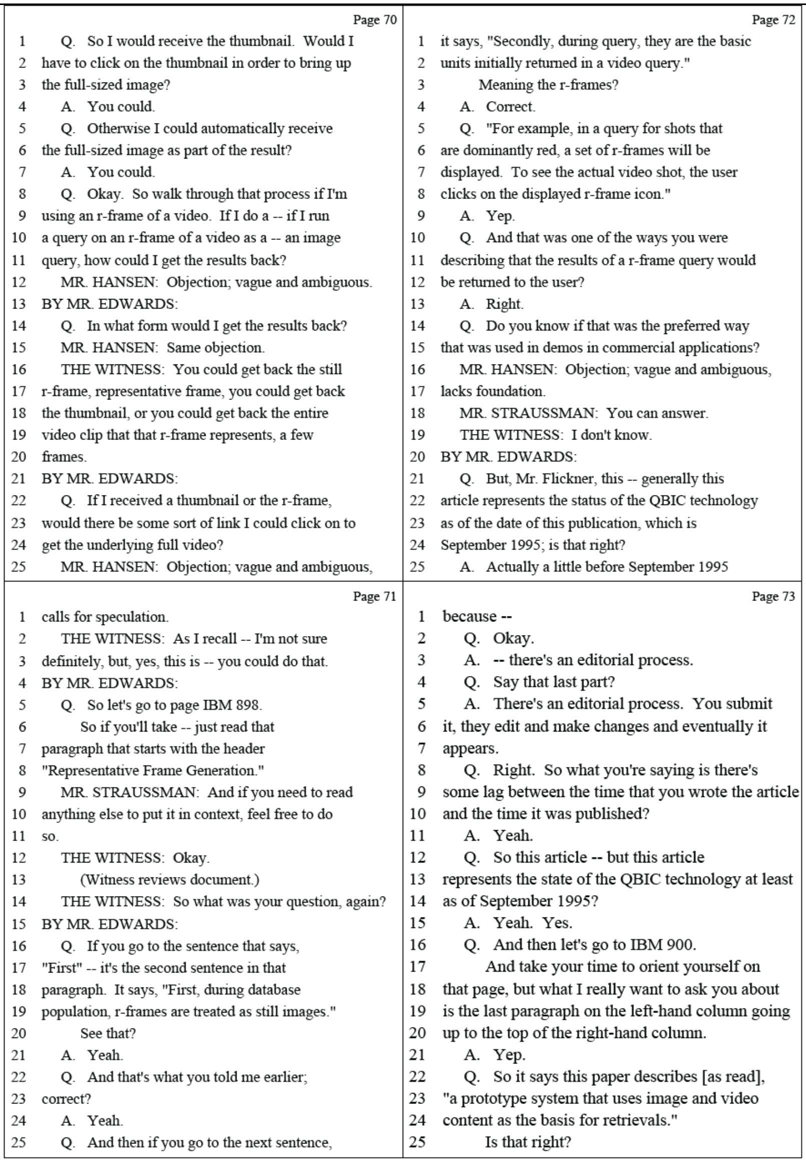

Page 70

Q. So I would receive the thumbnail. Would I have to click on the thumbnail in order to bring up the full-sized image?
A. You could.
Q. Otherwise I could automatically receive the full-sized image as part of the result?
A. You could.
Q. Okay. So walk through that process if I'm using an r-frame of a video. If I do a -- if I run a query on an r-frame of a video as a -- an image query, how could I get the results back?
MR. HANSEN: Objection; vague and ambiguous.
BY MR. EDWARDS:
Q. In what form would I get the results back?
MR. HANSEN: Same objection.
THE WITNESS: You could get back the still r-frame, representative frame, you could get back the thumbnail, or you could get back the entire video clip that that r-frame represents, a few frames.
BY MR. EDWARDS:
Q. If I received a thumbnail or the r-frame, would there be some sort of link I could click on to get the underlying full video?
MR. HANSEN: Objection; vague and ambiguous,

Page 71

calls for speculation.
THE WITNESS: As I recall -- I'm not sure definitely, but, yes, this is -- you could do that.
BY MR. EDWARDS:
Q. So let's go to page IBM 898.
So if you'll take -- just read that paragraph that starts with the header "Representative Frame Generation."
MR. STRAUSSMAN: And if you need to read anything else to put it in context, feel free to do so.
THE WITNESS: Okay.
(Witness reviews document.)
THE WITNESS: So what was your question, again?
BY MR. EDWARDS:
Q. If you go to the sentence that says, "First" -- it's the second sentence in that paragraph. It says, "First, during database population, r-frames are treated as still images."
See that?
A. Yeah.
Q. And that's what you told me earlier; correct?
A. Yeah.
Q. And then if you go to the next sentence,

Page 72

it says, "Secondly, during query, they are the basic units initially returned in a video query."
Meaning the r-frames?
A. Correct.
Q. "For example, in a query for shots that are dominantly red, a set of r-frames will be displayed. To see the actual video shot, the user clicks on the displayed r-frame icon."
A. Yep.
Q. And that was one of the ways you were describing that the results of a r-frame query would be returned to the user?
A. Right.
Q. Do you know if that was the preferred way that was used in demos in commercial applications?
MR. HANSEN: Objection; vague and ambiguous, lacks foundation.
MR. STRAUSSMAN: You can answer.
THE WITNESS: I don't know.
BY MR. EDWARDS:
Q. But, Mr. Flickner, this -- generally this article represents the status of the QBIC technology as of the date of this publication, which is September 1995; is that right?
A. Actually a little before September 1995

Page 73

because --
Q. Okay.
A. -- there's an editorial process.
Q. Say that last part?
A. There's an editorial process. You submit it, they edit and make changes and eventually it appears.
Q. Right. So what you're saying is there's some lag between the time that you wrote the article and the time it was published?
A. Yeah.
Q. So this article -- but this article represents the state of the QBIC technology at least as of September 1995?
A. Yeah. Yes.
Q. And then let's go to IBM 900.
And take your time to orient yourself on that page, but what I really want to ask you about is the last paragraph on the left-hand column going up to the top of the right-hand column.
A. Yep.
Q. So it says this paper describes [as read], "a prototype system that uses image and video content as the basis for retrievals."
Is that right?

		Id. at 129 – 131: 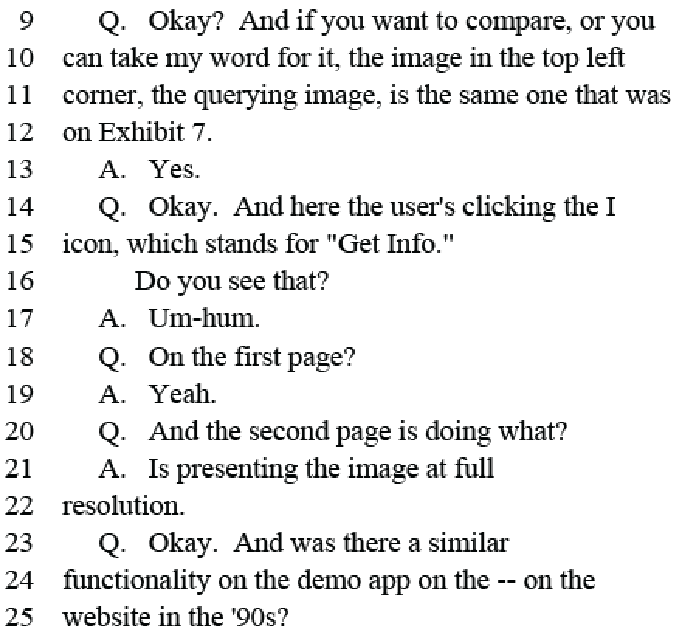 Q. Okay? And if you want to compare, or you can take my word for it, the image in the top left corner, the querying image, is the same one that was on Exhibit 7. A. Yes. Q. Okay. And here the user's clicking the I icon, which stands for "Get Info." Do you see that? A. Um-hum. Q. On the first page? A. Yeah. Q. And the second page is doing what? A. Is presenting the image at full resolution. Q. Okay. And was there a similar functionality on the demo app on the -- on the website in the '90s?

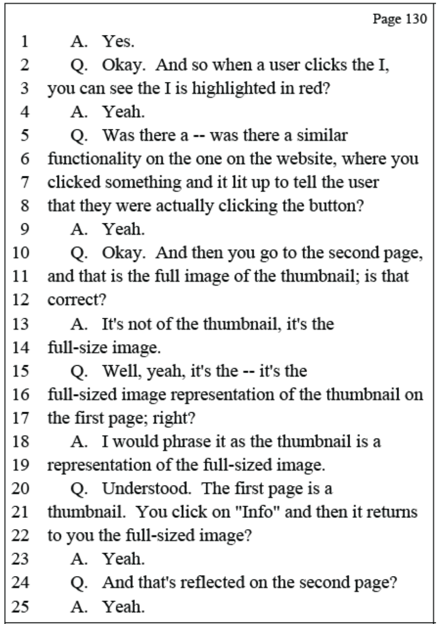

Page 130

A. Yes.
Q. Okay. And so when a user clicks the I, you can see the I is highlighted in red?
A. Yeah.
Q. Was there a -- was there a similar functionality on the one on the website, where you clicked something and it lit up to tell the user that they were actually clicking the button?
A. Yeah.
Q. Okay. And then you go to the second page, and that is the full image of the thumbnail; is that correct?
A. It's not of the thumbnail, it's the full-size image.
Q. Well, yeah, it's the -- it's the full-sized image representation of the thumbnail on the first page; right?
A. I would phrase it as the thumbnail is a representation of the full-sized image.
Q. Understood. The first page is a thumbnail. You click on "Info" and then it returns to you the full-sized image?
A. Yeah.
Q. And that's reflected on the second page?
A. Yeah.

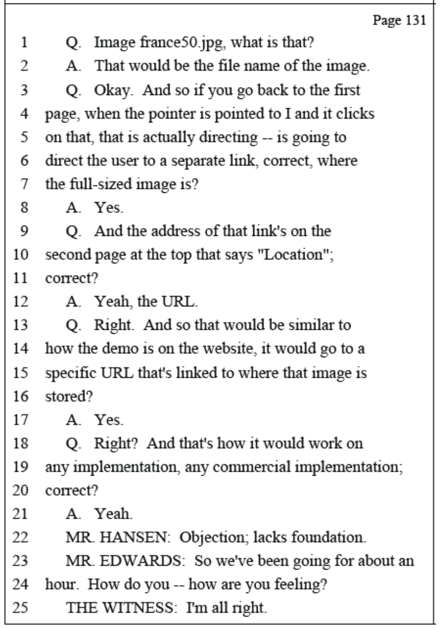

Page 131

Q. Image france50.jpg, what is that?
A. That would be the file name of the image.
Q. Okay. And so if you go back to the first page, when the pointer is pointed to I and it clicks on that, that is actually directing -- is going to direct the user to a separate link, correct, where the full-sized image is?
A. Yes.
Q. And the address of that link's on the second page at the top that says "Location"; correct?
A. Yeah, the URL.
Q. Right. And so that would be similar to how the demo is on the website, it would go to a specific URL that's linked to where that image is stored?
A. Yes.
Q. Right? And that's how it would work on any implementation, any commercial implementation; correct?
A. Yeah.
MR. HANSEN: Objection; lacks foundation.
MR. EDWARDS: So we've been going for about an hour. How do you -- how are you feeling?
THE WITNESS: I'm all right.

		Id. at 161: 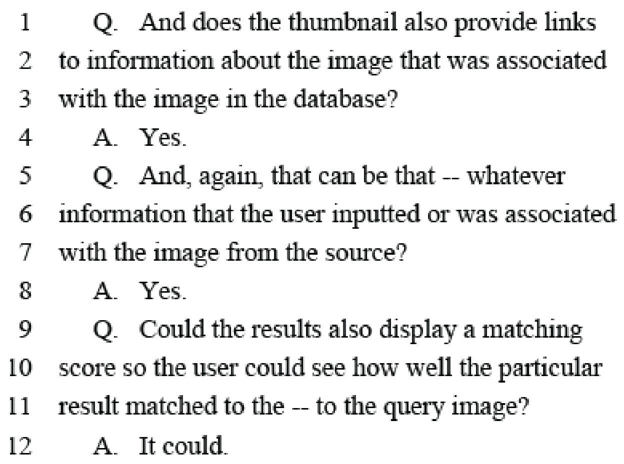 Q. And does the thumbnail also provide links to information about the image that was associated with the image in the database? A. Yes. Q. And, again, that can be that -- whatever information that the user inputted or was associated with the image from the source? A. Yes. Q. Could the results also display a matching score so the user could see how well the particular result matched to the -- to the query image? A. It could. *Id.* at 165:

		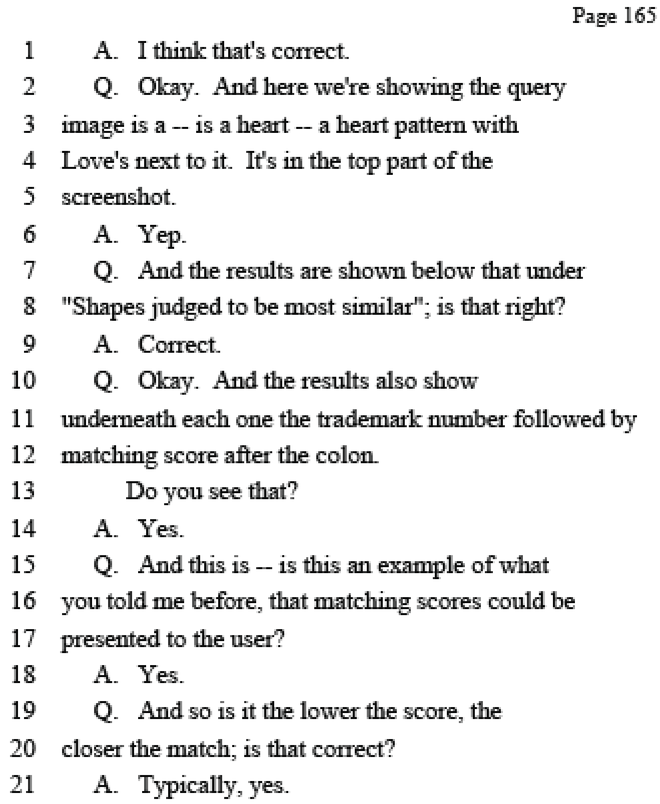 Page 165 A. I think that's correct. Q. Okay. And here we're showing the query image is a -- is a heart -- a heart pattern with Love's next to it. It's in the top part of the screenshot. A. Yep. Q. And the results are shown below that under "Shapes judged to be most similar"; is that right? A. Correct. Q. Okay. And the results also show underneath each one the trademark number followed by matching score after the colon. Do you see that? A. Yes. Q. And this is -- is this an example of what you told me before, that matching scores could be presented to the user? A. Yes. Q. And so is it the lower the score, the closer the match; is that correct? A. Typically, yes. *Id.* at 180-181:

		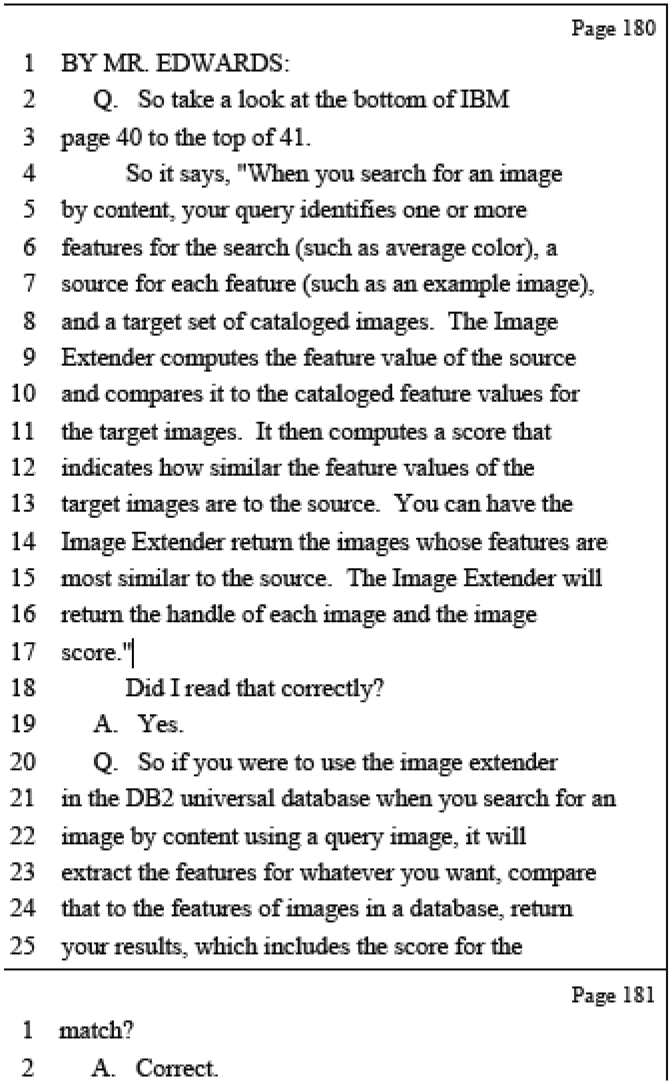 Page 180 BY MR. EDWARDS: Q. So take a look at the bottom of IBM page 40 to the top of 41. So it says, "When you search for an image by content, your query identifies one or more features for the search (such as average color), a source for each feature (such as an example image), and a target set of cataloged images. The Image Extender computes the feature value of the source and compares it to the cataloged feature values for the target images. It then computes a score that indicates how similar the feature values of the target images are to the source. You can have the Image Extender return the images whose features are most similar to the source. The Image Extender will return the handle of each image and the image score." Did I read that correctly? A. Yes. Q. So if you were to use the image extender in the DB2 universal database when you search for an image by content using a query image, it will extract the features for whatever you want, compare that to the features of images in a database, return your results, which includes the score for the Page 181 match? A. Correct.

Id. at 305:

So these articles describe various features of various versions of QBIC; right?

A. Correct.

Q. Do they describe any features -- does the demo you operated include any features that weren't described in these articles?

A. That weren't described in the articles? No.

Flickner 8/29/23 Depo Ex. 7 at 1:

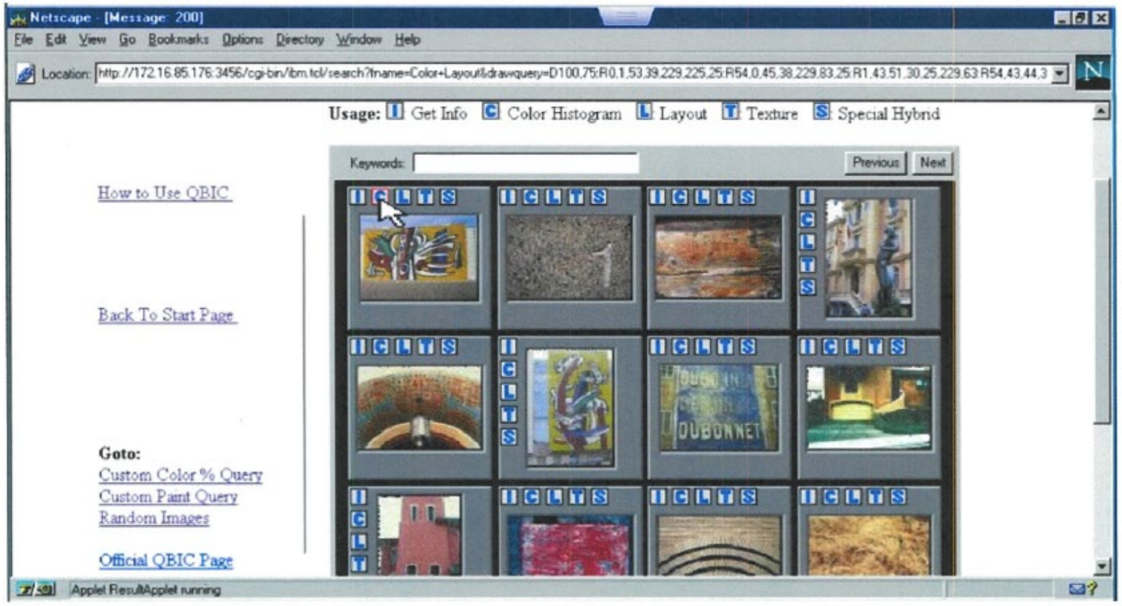

Location: http://172.16.85.176:3456/cgi-bin/ibm.tcl/search?fname=Color+Layout&drawquery=D100,75:R0,1,53,39,229,225,25:R54,0,45,38,229,83,25:R1,43,51,30,25,229,63:R54,43,44,3

		Id. at 2: 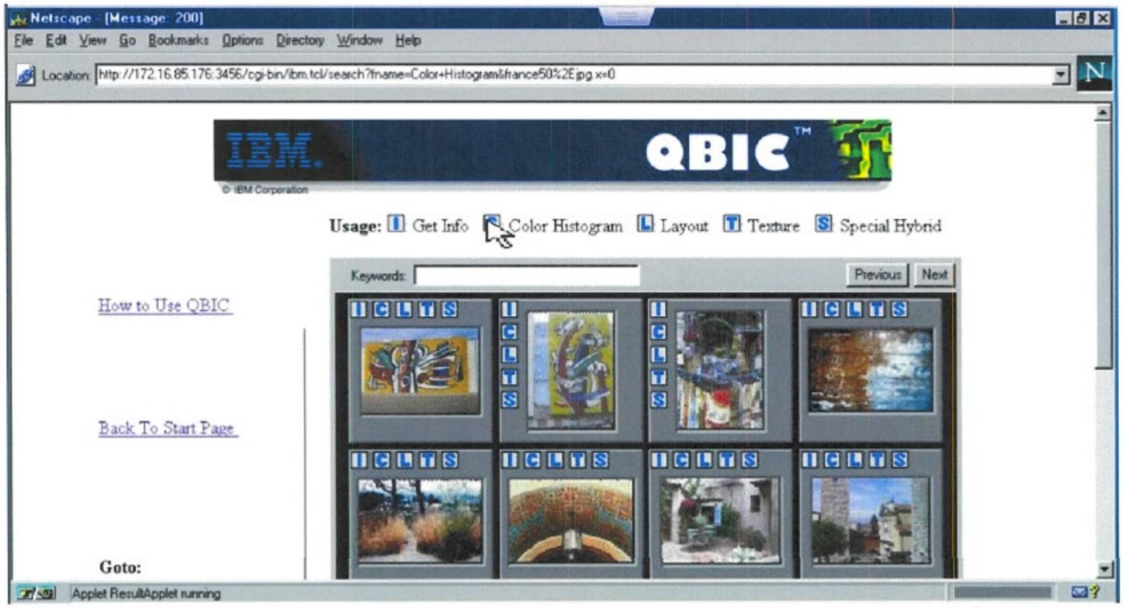 Location: http://172.16.85.176:3456/cgi-bin/ibm.tcl/search?fname=Color+Histogram&france50%2Ejpg.x=0 Flickner 8/29/23 Depo Ex. 8 at 1:

<table>
<tr>
<td></td>
<td></td>
<td>
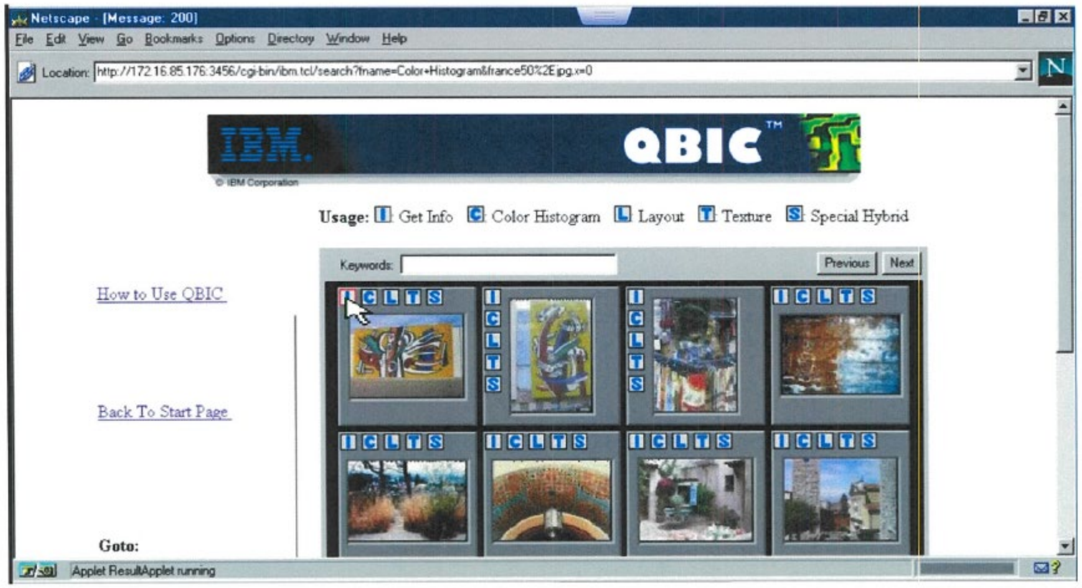

Location: http://172.16.85.176:3456/cgi-bin/ibm.tcl/search?fname=Color+Histogram&france50%2Ejpg.x=0

Id. at 2:
</td>
</tr>
</table>

		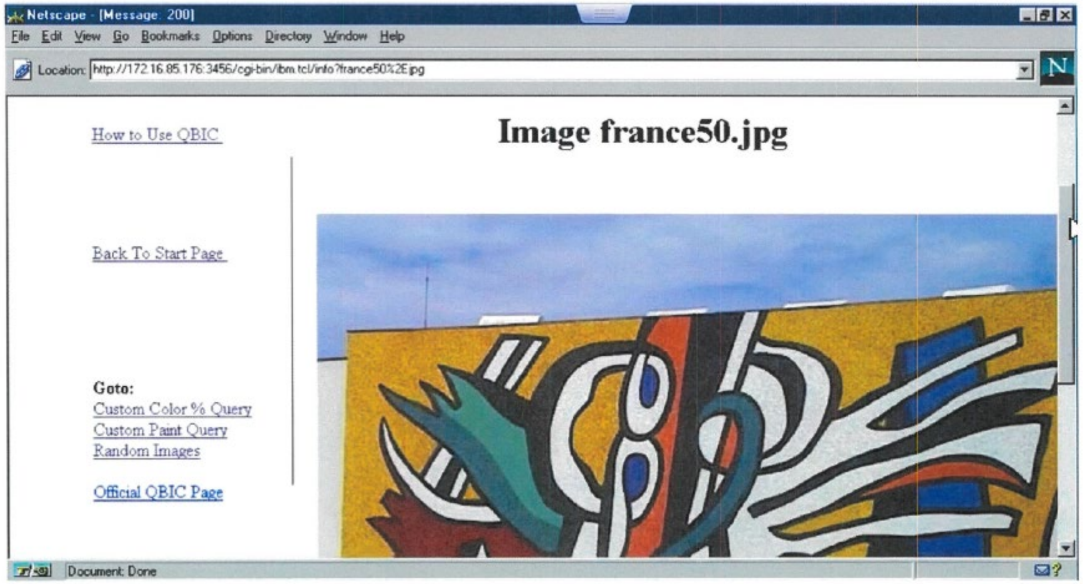 QBIC (Query by Image Content) (BOFA00001391–392) at 2 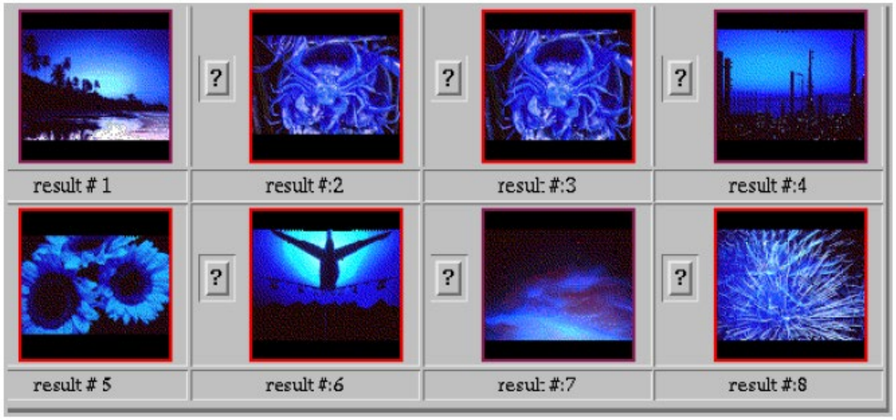 **Result of QBIC Search**

		To the extent it is found that QBIC does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that QBIC does not anticipate this claim, QBIC renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine QBIC with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

Exhibit E-31

Bank of America Corporation ("BofA" or "Defendant") provides this chart subject to all reservations, objections, statements, and disclaimers set forth herein and in BofA's Final Invalidity Contentions Cover Pleading, as well as any amendment, supplement, or modification thereof, which are incorporated herein by reference in their entirety.

On information and belief, CoolTown was made, known or used by others in the United States no later than July 10, 2000. *See generally* T. Kindberg et al., A Web-Based Nomadic Computing System; J. Barton et al., The Cooltown User Experience; cooltown developer's network, coolbase overview; cooltown developer's network, people, places, things: web presence for the real world. CoolTown is available as prior art at least under 35 U.S.C. § 102(a). On information and belief, HP was using CoolTown and made it available to the public in the United States by releasing the website on the Internet (and accessible to U.S. users) by at least March 2000, making the software available on the HP website by March 2000, and making available on the Internet (and accessible to U.S. users) the open source code, which reflects the functionality of CoolTown by at least March 2000. *See e.g.*, HP000043–044; HP000045–046; HP000060–061; *see also* https://www.zdnet.com/article/hp-brings-open-source-to-cooltown/. Therefore, HP Cooltown was at least known, used, and/or sold to the public by at least 2000. On information and belief, CoolTown is also available as prior art under 102(g). BofA is diligently investigating CoolTown and will supplement with additional information, if necessary.

As shown in the chart below, CoolTown anticipates asserted claims 18, 27, 29, and 31 of the '252 Patent. To the extent CoolTown is found not to expressly disclose certain limitations in the asserted claims, such limitations are inherent. To the extent CoolTown is found not to anticipate any asserted claims or claim elements of the '252 Patent, the reference nevertheless renders those claims or claim elements obvious under 35 U.S.C. § 103, either alone or in combination with other art identified herein or in the cover pleading of BofA's Invalidity Contentions, for the reasons identified in the cover pleading. These Final Invalidity Contentions are not an admission by BofA that the accused products, including any current or past versions of the accused products, are covered by, or infringe the asserted claims, but are based instead on the recognition that if the claims are interpreted to be broad enough to encompass the accused products, the claims must also be construed to have that same scope when considering whether they are invalid.

The following chart is based on the Court's claim construction rulings (Dkt. Nos. 145, 153, 236), and to the extent not contradictory, Plaintiffs' apparent interpretations of the claims in Plaintiff's "Preliminary" Final Infringement Contentions. BofA notes that in many instances, Plaintiff's "Preliminary" Final Infringement Contentions ignore and/or contradict the Court's claim construction rulings. BofA does not accept the assumptions concerning the scope and meaning implicit in Plaintiff's Preliminary Final Infringement Contentions, to the extent those assumptions are discernible, and reserves the right to challenge Plaintiff's proposed (or implied) constructions. Furthermore, Plaintiffs' "Preliminary" Final Contentions fail to comply with the level of specificity required by SPR 2.1.2 and 2.1.3, and include improper reservation of rights to supplement with "additional ways in which BoA Accused Products infringe." Any alleged deficient discovery is of Plaintiffs own making. Most if not all of the deficiencies in Plaintiffs' Preliminary Final Contentions have been documented by Defendants through multiple letters since at least May 18, 2021 and up to Bank of America's most recent letter of August 19, 2023. Plaintiffs' shifting infringement contentions, deficiencies in its infringement contentions, and changing claim scope have prejudiced BofA from being able to reasonably prepare its defenses, including its invalidity contentions. BofA also reserves the right to revise and supplement these charts if and when Plaintiff is permitted to provide revised or supplemental Infringement Contentions. Citations given in the chart below are merely representative of the respective elements to provide notice to Plaintiffs' of BofA's invalidity theories and are not meant to be exhaustive.

1. Claim 18

Element	'252 Claim Recitation	Exemplary Citations in Reference
[18Pre]	A method for retrieving information from image processing, the method comprising:	CoolTown discloses a method for retrieving information from image processing, the method comprising: The Cooltown User Experience (HP000088 – 093) at 1: 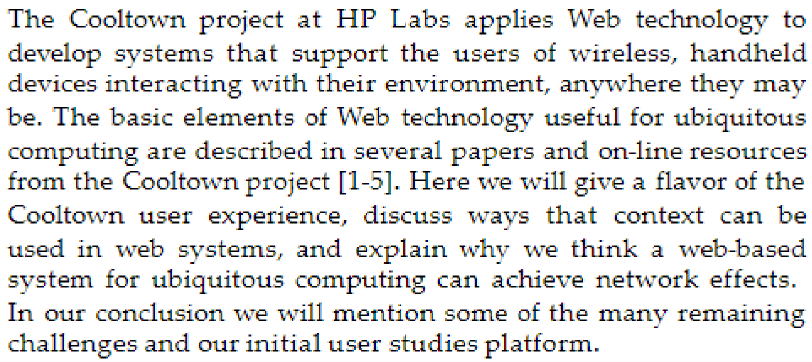 The Cooltown project at HP Labs applies Web technology to develop systems that support the users of wireless, handheld devices interacting with their environment, anywhere they may be. The basic elements of Web technology useful for ubiquitous computing are described in several papers and on-line resources from the Cooltown project [1-5]. Here we will give a flavor of the Cooltown user experience, discuss ways that context can be used in web systems, and explain why we think a web-based system for ubiquitous computing can achieve network effects. In our conclusion we will mention some of the many remaining challenges and our initial user studies platform. A Web-Based Nomadic Computing System (HP000074 – 87) at 1 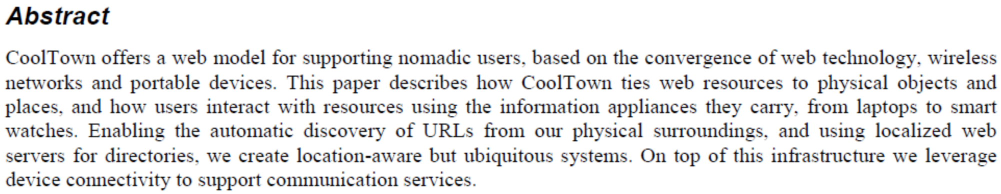 ***Abstract*** CoolTown offers a web model for supporting nomadic users, based on the convergence of web technology, wireless networks and portable devices. This paper describes how CoolTown ties web resources to physical objects and places, and how users interact with resources using the information appliances they carry, from laptops to smart watches. Enabling the automatic discovery of URLs from our physical surroundings, and using localized web servers for directories, we create location-aware but ubiquitous systems. On top of this infrastructure we leverage device connectivity to support communication services. Coolbase[1] Overview at 1

---

[1] *See* https://web.archive.org/web/20011030090650/http://cooltown.hp.com/dev/coolbase-overview.asp

Element	'252 Claim Recitation	Exemplary Citations in Reference
		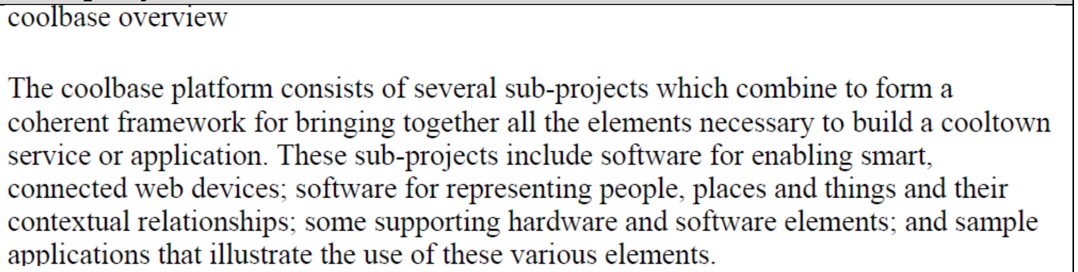 coolbase overview The coolbase platform consists of several sub-projects which combine to form a coherent framework for bringing together all the elements necessary to build a cooltown service or application. These sub-projects include software for enabling smart, connected web devices; software for representing people, places and things and their contextual relationships; some supporting hardware and software elements; and sample applications that illustrate the use of these various elements.
[18A]	providing a mobile device having a camera the mobile device configured to capture an image and configured to transmit data relating to the image an image processing platform;	CoolTown discloses a mobile device having a camera, the mobile device configured to capture an image, and configured to transmit data relating to the image an image processing platform. A Web-Based Nomadic Computing System (HP000074 – 87) at 1 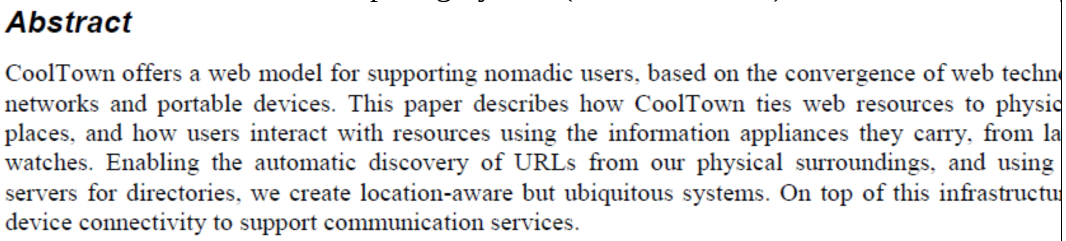 **Abstract** CoolTown offers a web model for supporting nomadic users, based on the convergence of web techn networks and portable devices. This paper describes how CoolTown ties web resources to physic places, and how users interact with resources using the information appliances they carry, from la watches. Enabling the automatic discovery of URLs from our physical surroundings, and using servers for directories, we create location-aware but ubiquitous systems. On top of this infrastructu device connectivity to support communication services. The Cooltown: People, Places, Things Web Presence for the Real World (HP000025 – 34) at 20:

Element	'252 Claim Recitation	Exemplary Citations in Reference
		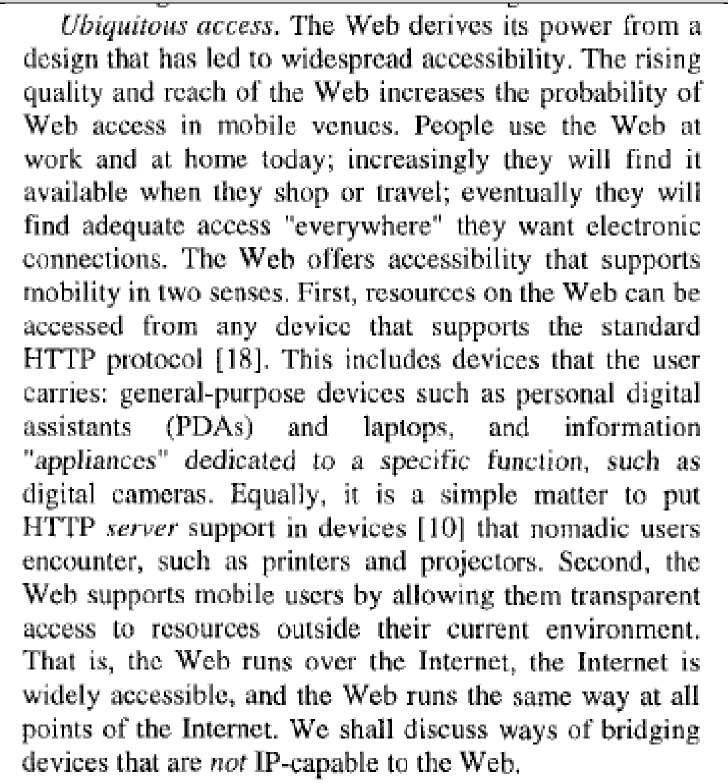 *Ubiquitous access*. The Web derives its power from a design that has led to widespread accessibility. The rising quality and reach of the Web increases the probability of Web access in mobile venues. People use the Web at work and at home today; increasingly they will find it available when they shop or travel; eventually they will find adequate access "everywhere" they want electronic connections. The Web offers accessibility that supports mobility in two senses. First, resources on the Web can be accessed from any device that supports the standard HTTP protocol [18]. This includes devices that the user carries: general-purpose devices such as personal digital assistants (PDAs) and laptops, and information "appliances" dedicated to a specific function, such as digital cameras. Equally, it is a simple matter to put HTTP *server* support in devices [10] that nomadic users encounter, such as printers and projectors. Second, the Web supports mobile users by allowing them transparent access to resources outside their current environment. That is, the Web runs over the Internet, the Internet is widely accessible, and the Web runs the same way at all points of the Internet. We shall discuss ways of bridging devices that are *not* IP-capable to the Web. The Cooltown: People, Places, Things Web Presence for the Real World (HP000025 – 34) at 24:

Element	'252 Claim Recitation	Exemplary Citations in Reference
		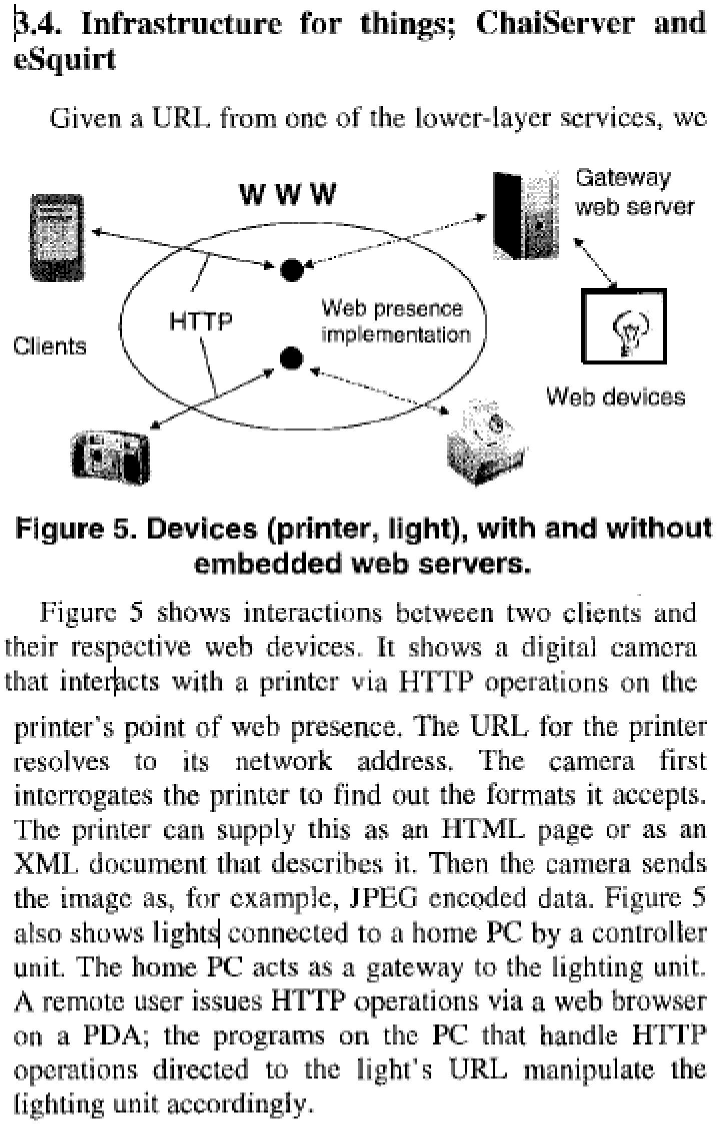 **3.4. Infrastructure for things; ChaiServer and eSquirt** Given a URL from one of the lower-layer services, we **Figure 5. Devices (printer, light), with and without embedded web servers.** Figure 5 shows interactions between two clients and their respective web devices. It shows a digital camera that interacts with a printer via HTTP operations on the printer's point of web presence. The URL for the printer resolves to its network address. The camera first interrogates the printer to find out the formats it accepts. The printer can supply this as an HTML page or as an XML document that describes it. Then the camera sends the image as, for example, JPEG encoded data. Figure 5 also shows lights connected to a home PC by a controller unit. The home PC acts as a gateway to the lighting unit. A remote user issues HTTP operations via a web browser on a PDA; the programs on the PC that handle HTTP operations directed to the light's URL manipulate the lighting unit accordingly. The Cooltown User Experience (HP000088 – 093) at 1:

Element	’252 Claim Recitation	Exemplary Citations in Reference
		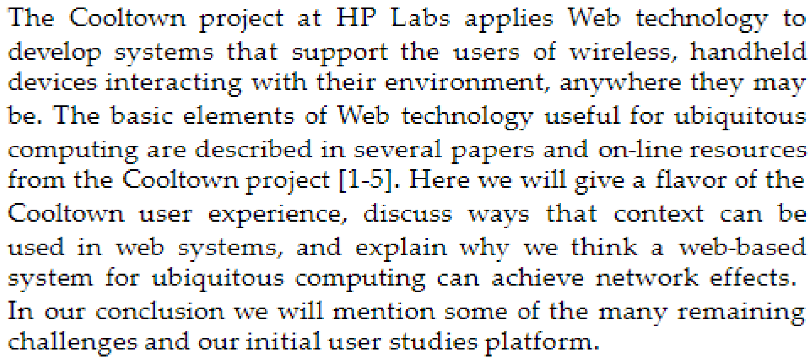 The Cooltown project at HP Labs applies Web technology to develop systems that support the users of wireless, handheld devices interacting with their environment, anywhere they may be. The basic elements of Web technology useful for ubiquitous computing are described in several papers and on-line resources from the Cooltown project [1-5]. Here we will give a flavor of the Cooltown user experience, discuss ways that context can be used in web systems, and explain why we think a web-based system for ubiquitous computing can achieve network effects. In our conclusion we will mention some of the many remaining challenges and our initial user studies platform. A Web-Based Nomadic Computing System (HP000074 – 87) at 2 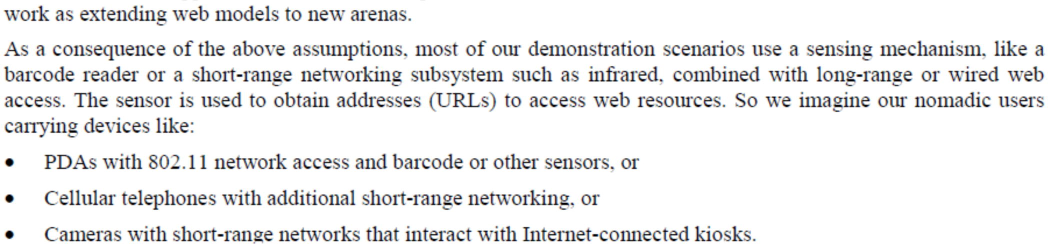 The CoolTown work is not on the devices or wireless communication layers, but rather on the layer of infrastructure that is needed to support nomadic users with portable devices. We work at the "web server" level and we see our work as extending web models to new arenas. As a consequence of the above assumptions, most of our demonstration scenarios use a sensing mechanism, like a barcode reader or a short-range networking subsystem such as infrared, combined with long-range or wired web access. The sensor is used to obtain addresses (URLs) to access web resources. So we imagine our nomadic users carrying devices like: • PDAs with 802.11 network access and barcode or other sensors, or • Cellular telephones with additional short-range networking, or • Cameras with short-range networks that interact with Internet-connected kiosks. A Web-Based Nomadic Computing System (HP000074 – 87) at 3

Element	'252 Claim Recitation	Exemplary Citations in Reference
		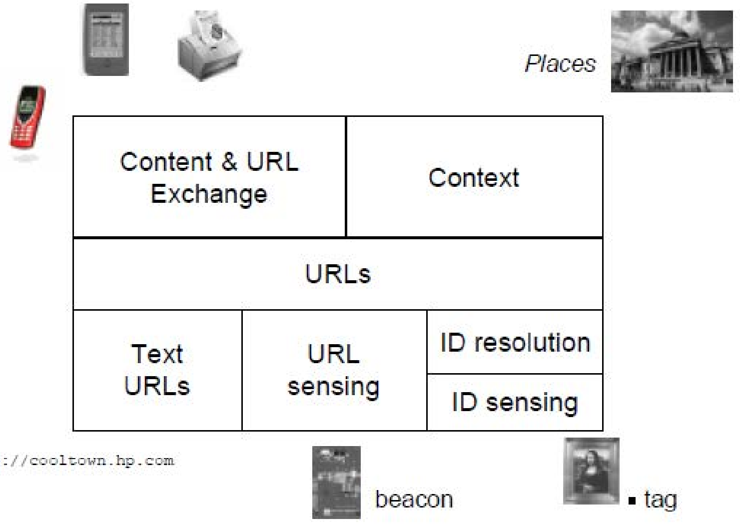 Figure 2. Component relationships. Sources of URLs, including keyboards, beacons, and resolved tags point to context or content. A Web-Based Nomadic Computing System (HP000074 – 87) at 4 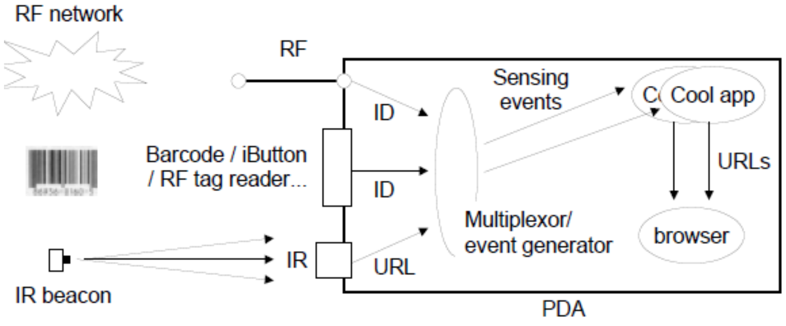 Figure 3. A sketch of a PDA equipped for sensing URLs and identifiers.

Element	'252 Claim Recitation	Exemplary Citations in Reference
		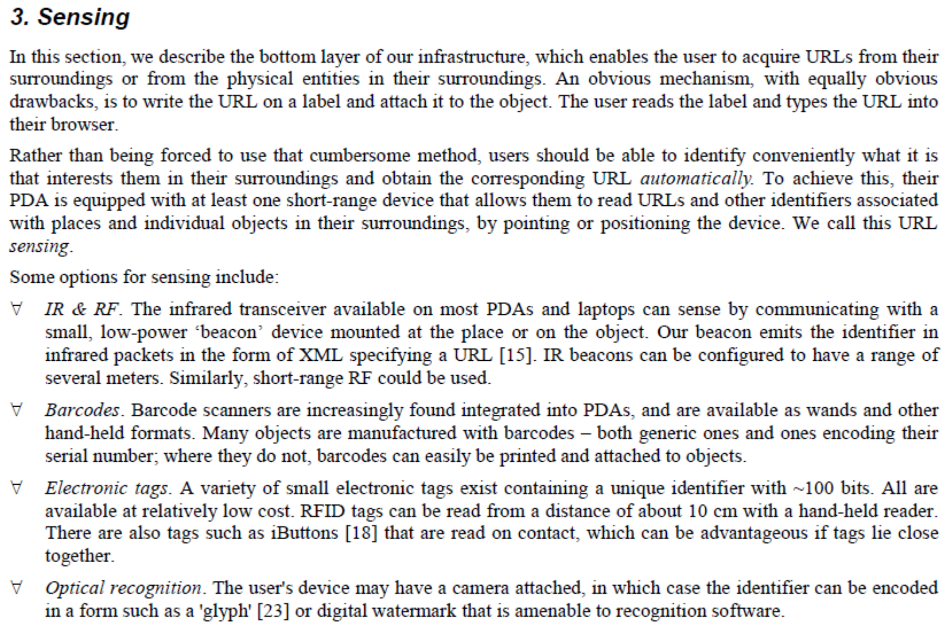 ***3. Sensing*** In this section, we describe the bottom layer of our infrastructure, which enables the user to acquire URLs from their surroundings or from the physical entities in their surroundings. An obvious mechanism, with equally obvious drawbacks, is to write the URL on a label and attach it to the object. The user reads the label and types the URL into their browser. Rather than being forced to use that cumbersome method, users should be able to identify conveniently what it is that interests them in their surroundings and obtain the corresponding URL *automatically*. To achieve this, their PDA is equipped with at least one short-range device that allows them to read URLs and other identifiers associated with places and individual objects in their surroundings, by pointing or positioning the device. We call this URL *sensing*. Some options for sensing include: ∀ *IR & RF*. The infrared transceiver available on most PDAs and laptops can sense by communicating with a small, low-power 'beacon' device mounted at the place or on the object. Our beacon emits the identifier in infrared packets in the form of XML specifying a URL [15]. IR beacons can be configured to have a range of several meters. Similarly, short-range RF could be used. ∀ *Barcodes*. Barcode scanners are increasingly found integrated into PDAs, and are available as wands and other hand-held formats. Many objects are manufactured with barcodes – both generic ones and ones encoding their serial number; where they do not, barcodes can easily be printed and attached to objects. ∀ *Electronic tags*. A variety of small electronic tags exist containing a unique identifier with ~100 bits. All are available at relatively low cost. RFID tags can be read from a distance of about 10 cm with a hand-held reader. There are also tags such as iButtons [18] that are read on contact, which can be advantageous if tags lie close together. ∀ *Optical recognition*. The user's device may have a camera attached, in which case the identifier can be encoded in a form such as a 'glyph' [23] or digital watermark that is amenable to recognition software. To the extent it is found that CoolTown does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that CoolTown does not anticipate this claim, CoolTown renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine CoolTown with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

Element	'252 Claim Recitation	Exemplary Citations in Reference
[18B]	configuring the image processing platform to receive the data relating to the image and to conduct image processing, including:	CoolTown discloses configuring the image processing platform to receive the data relating to the image and to conduct image processing. The Cooltown: People, Places, Things Web Presence for the Real World (HP000025 – 34) at 20: 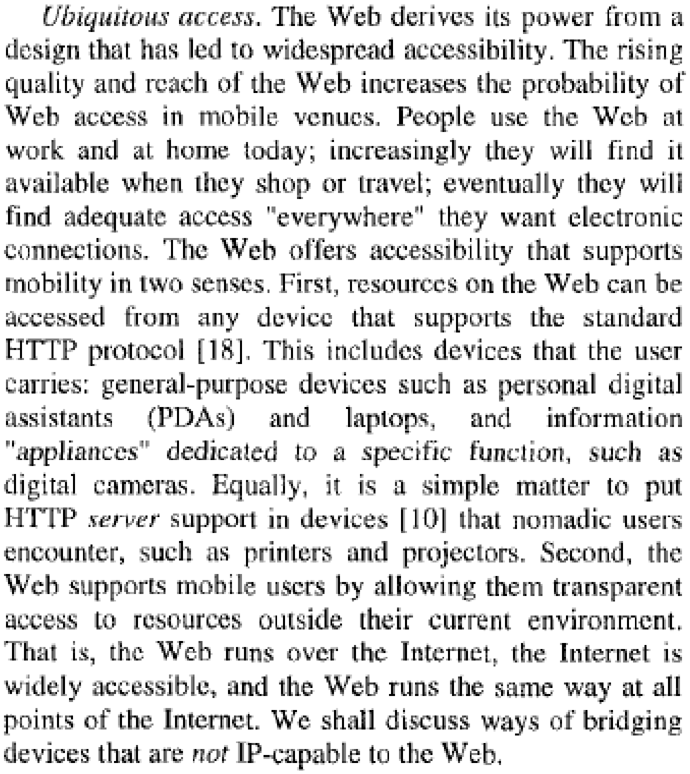 *Ubiquitous access*. The Web derives its power from a design that has led to widespread accessibility. The rising quality and reach of the Web increases the probability of Web access in mobile venues. People use the Web at work and at home today; increasingly they will find it available when they shop or travel; eventually they will find adequate access "everywhere" they want electronic connections. The Web offers accessibility that supports mobility in two senses. First, resources on the Web can be accessed from any device that supports the standard HTTP protocol [18]. This includes devices that the user carries: general-purpose devices such as personal digital assistants (PDAs) and laptops, and information "appliances" dedicated to a specific function, such as digital cameras. Equally, it is a simple matter to put HTTP *server* support in devices [10] that nomadic users encounter, such as printers and projectors. Second, the Web supports mobile users by allowing them transparent access to resources outside their current environment. That is, the Web runs over the Internet, the Internet is widely accessible, and the Web runs the same way at all points of the Internet. We shall discuss ways of bridging devices that are *not* IP-capable to the Web. The Cooltown: People, Places, Things Web Presence for the Real World (HP000025 – 34) at 24:

Element	'252 Claim Recitation	Exemplary Citations in Reference
		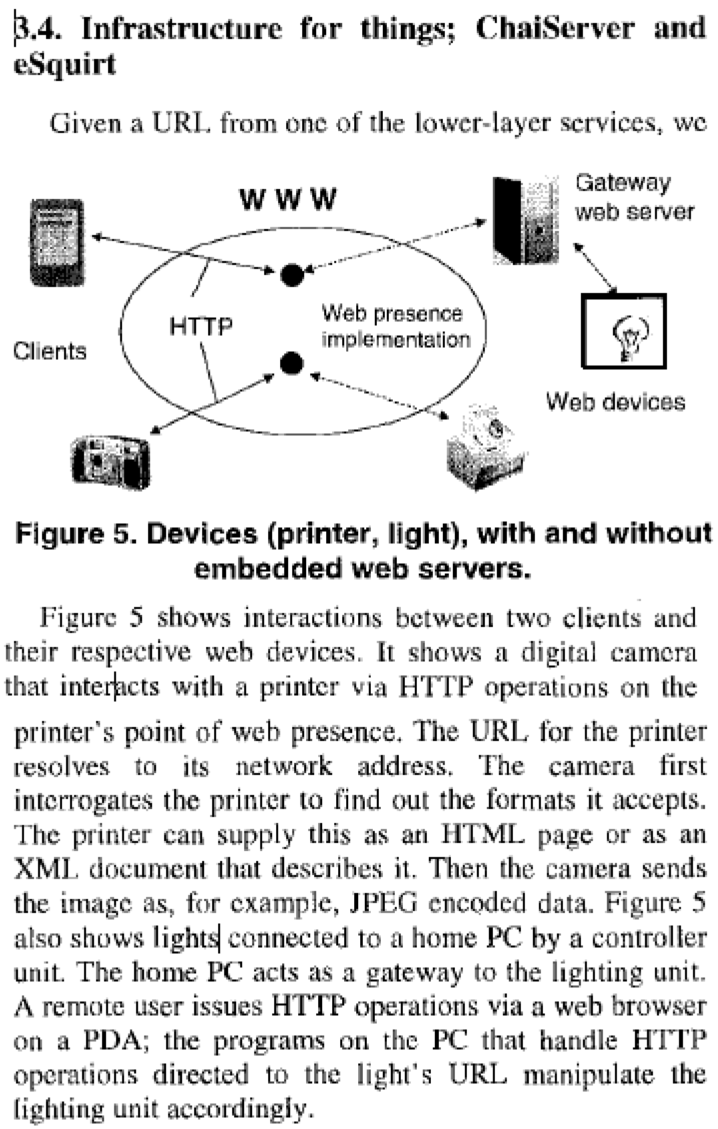 **3.4. Infrastructure for things; ChaiServer and eSquirt** Given a URL from one of the lower-layer services, we **Figure 5. Devices (printer, light), with and without embedded web servers.** Figure 5 shows interactions between two clients and their respective web devices. It shows a digital camera that interacts with a printer via HTTP operations on the printer's point of web presence. The URL for the printer resolves to its network address. The camera first interrogates the printer to find out the formats it accepts. The printer can supply this as an HTML page or as an XML document that describes it. Then the camera sends the image as, for example, JPEG encoded data. Figure 5 also shows lights connected to a home PC by a controller unit. The home PC acts as a gateway to the lighting unit. A remote user issues HTTP operations via a web browser on a PDA; the programs on the PC that handle HTTP operations directed to the light's URL manipulate the lighting unit accordingly. A Web-Based Nomadic Computing System (HP000074 – 87) at 4

Element	'252 Claim Recitation	Exemplary Citations in Reference
		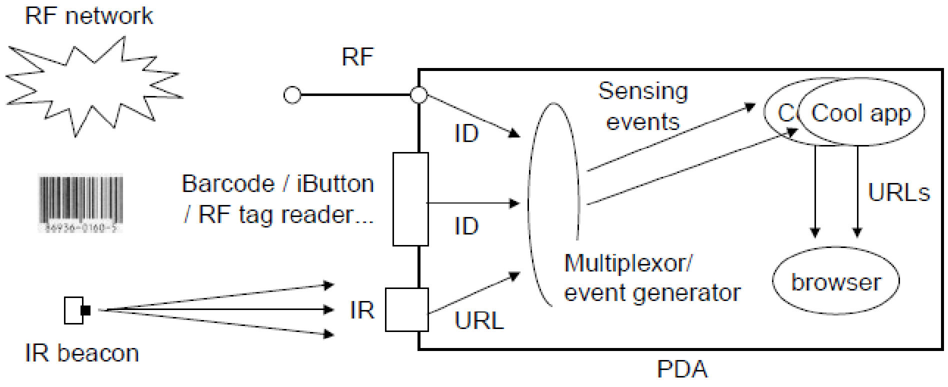 Figure 3. A sketch of a PDA equipped for sensing URLs and identifiers. 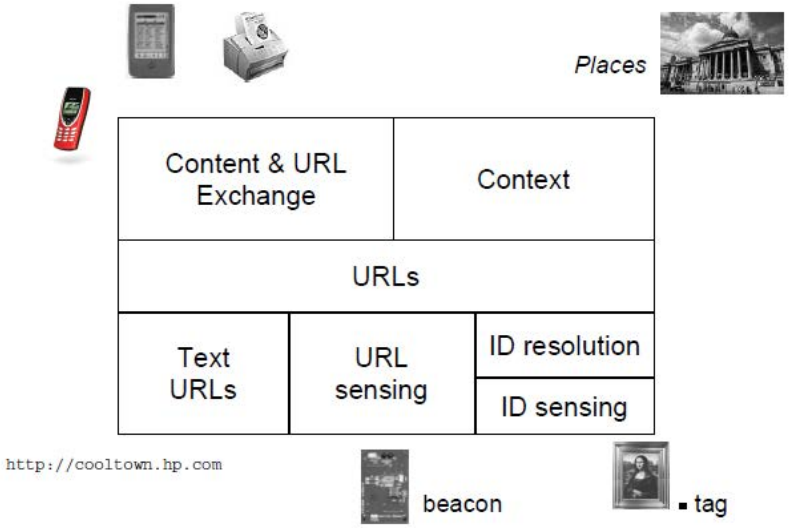 Figure 2. Component relationships. Sources of URLs, including keyboards, beacons, and resolved tags point to context or content. A Web-Based Nomadic Computing System (HP000074 – 87) at 4

Element	'252 Claim Recitation	Exemplary Citations in Reference
		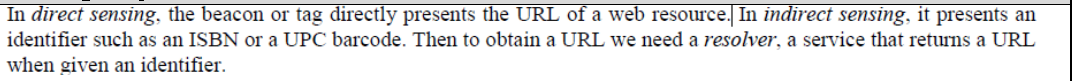In *direct sensing*, the beacon or tag directly presents the URL of a web resource. In *indirect sensing*, it presents an identifier such as an ISBN or a UPC barcode. Then to obtain a URL we need a *resolver*, a service that returns a URL when given an identifier. A Web-Based Nomadic Computing System (HP000074 – 87) at 5 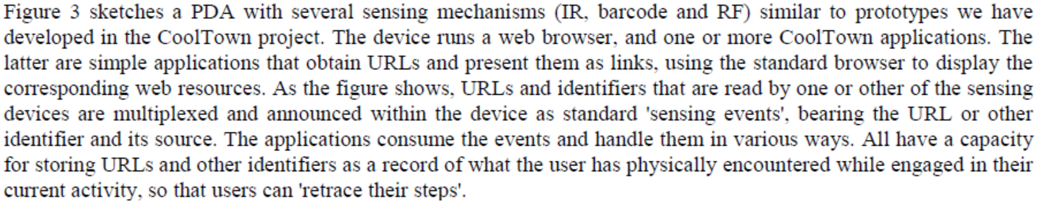Figure 3 sketches a PDA with several sensing mechanisms (IR, barcode and RF) similar to prototypes we have developed in the CoolTown project. The device runs a web browser, and one or more CoolTown applications. The latter are simple applications that obtain URLs and present them as links, using the standard browser to display the corresponding web resources. As the figure shows, URLs and identifiers that are read by one or other of the sensing devices are multiplexed and announced within the device as standard 'sensing events', bearing the URL or other identifier and its source. The applications consume the events and handle them in various ways. All have a capacity for storing URLs and other identifiers as a record of what the user has physically encountered while engaged in their current activity, so that users can 'retrace their steps'. To the extent it is found that CoolTown does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that CoolTown does not anticipate this claim, CoolTown renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine CoolTown with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[18B.i]	operating on the data relating to the image to determine if the image contains one or more recognizable symbols; and	CoolTown discloses operating on the data relating to the image to determine if the image contains one or more recognizable symbols. A Web-Based Nomadic Computing System (HP000074 – 87) at 3

Element	'252 Claim Recitation	Exemplary Citations in Reference
		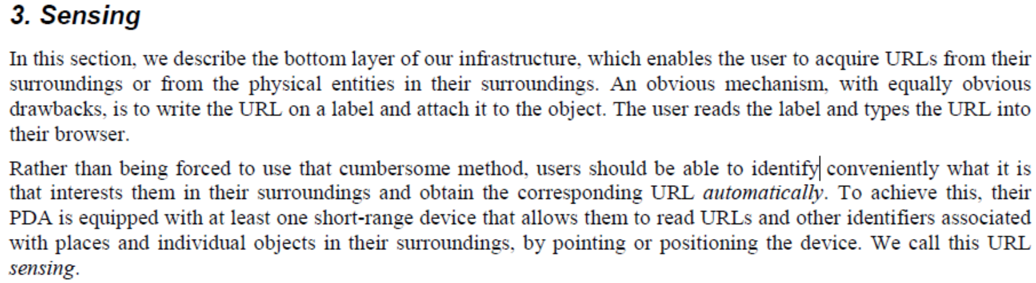 ***3. Sensing*** In this section, we describe the bottom layer of our infrastructure, which enables the user to acquire URLs from their surroundings or from the physical entities in their surroundings. An obvious mechanism, with equally obvious drawbacks, is to write the URL on a label and attach it to the object. The user reads the label and types the URL into their browser. Rather than being forced to use that cumbersome method, users should be able to identify conveniently what it is that interests them in their surroundings and obtain the corresponding URL *automatically*. To achieve this, their PDA is equipped with at least one short-range device that allows them to read URLs and other identifiers associated with places and individual objects in their surroundings, by pointing or positioning the device. We call this URL *sensing*. A Web-Based Nomadic Computing System (HP000074 – 87) at 4 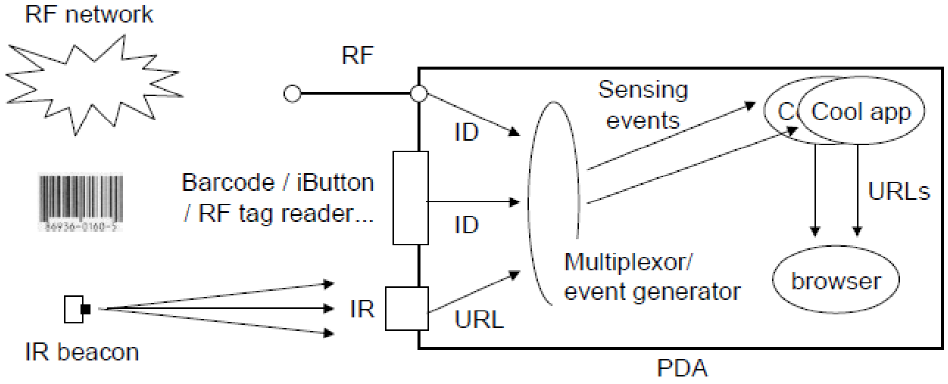 Figure 3. A sketch of a PDA equipped for sensing URLs and identifiers.

Element	'252 Claim Recitation	Exemplary Citations in Reference
		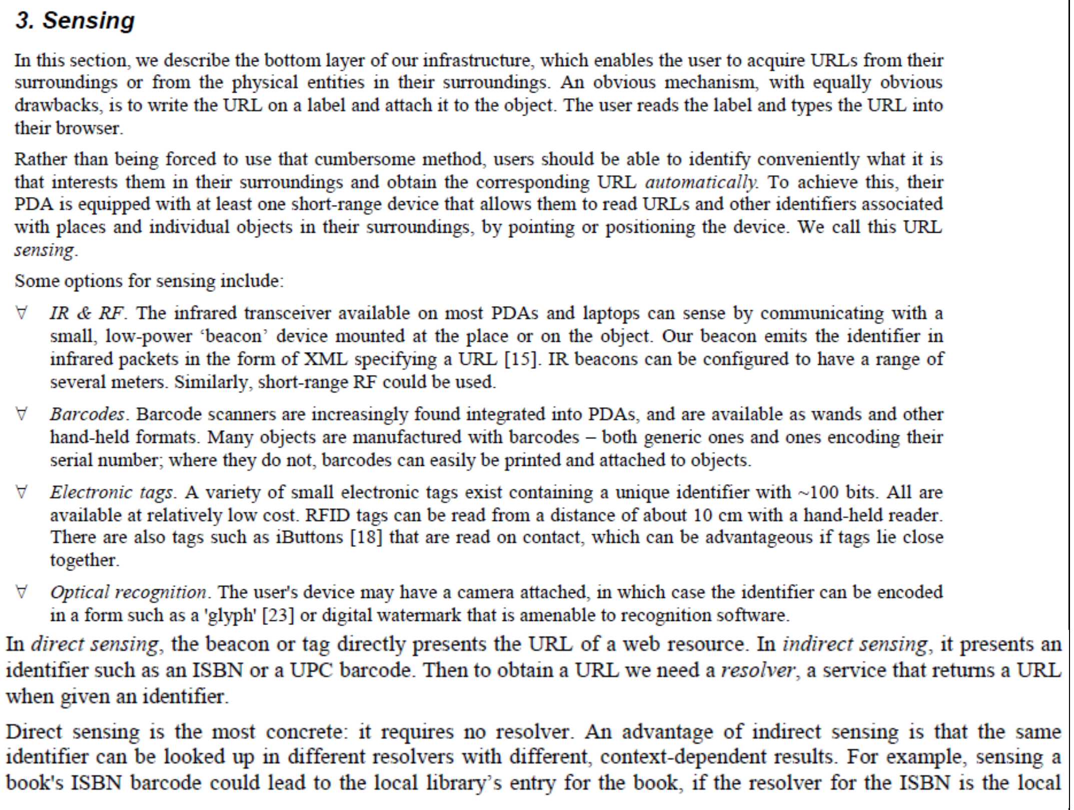 **3. Sensing** In this section, we describe the bottom layer of our infrastructure, which enables the user to acquire URLs from their surroundings or from the physical entities in their surroundings. An obvious mechanism, with equally obvious drawbacks, is to write the URL on a label and attach it to the object. The user reads the label and types the URL into their browser. Rather than being forced to use that cumbersome method, users should be able to identify conveniently what it is that interests them in their surroundings and obtain the corresponding URL *automatically*. To achieve this, their PDA is equipped with at least one short-range device that allows them to read URLs and other identifiers associated with places and individual objects in their surroundings, by pointing or positioning the device. We call this URL *sensing*. Some options for sensing include: ∀ *IR & RF*. The infrared transceiver available on most PDAs and laptops can sense by communicating with a small, low-power 'beacon' device mounted at the place or on the object. Our beacon emits the identifier in infrared packets in the form of XML specifying a URL [15]. IR beacons can be configured to have a range of several meters. Similarly, short-range RF could be used. ∀ *Barcodes*. Barcode scanners are increasingly found integrated into PDAs, and are available as wands and other hand-held formats. Many objects are manufactured with barcodes – both generic ones and ones encoding their serial number; where they do not, barcodes can easily be printed and attached to objects. ∀ *Electronic tags*. A variety of small electronic tags exist containing a unique identifier with ~100 bits. All are available at relatively low cost. RFID tags can be read from a distance of about 10 cm with a hand-held reader. There are also tags such as iButtons [18] that are read on contact, which can be advantageous if tags lie close together. ∀ *Optical recognition*. The user's device may have a camera attached, in which case the identifier can be encoded in a form such as a 'glyph' [23] or digital watermark that is amenable to recognition software. In *direct sensing*, the beacon or tag directly presents the URL of a web resource. In *indirect sensing*, it presents an identifier such as an ISBN or a UPC barcode. Then to obtain a URL we need a *resolver*, a service that returns a URL when given an identifier. Direct sensing is the most concrete: it requires no resolver. An advantage of indirect sensing is that the same identifier can be looked up in different resolvers with different, context-dependent results. For example, sensing a book's ISBN barcode could lead to the local library's entry for the book, if the resolver for the ISBN is the local Cooltown Developer's Network, People, Place, and Things: web presence for the real world (HP000025–34) (HP000025–34) at 20 – 21 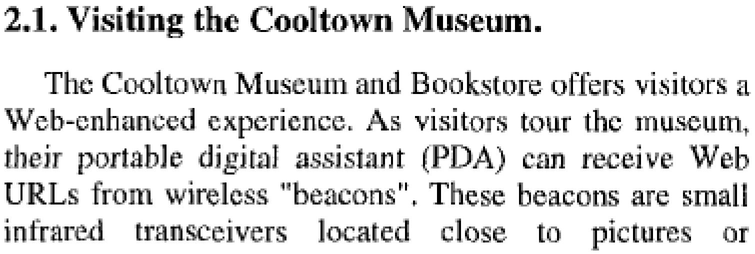 **2.1. Visiting the Cooltown Museum.** The Cooltown Museum and Bookstore offers visitors a Web-enhanced experience. As visitors tour the museum, their portable digital assistant (PDA) can receive Web URLs from wireless "beacons". These beacons are small infrared transceivers located close to pictures or

Element	'252 Claim Recitation	Exemplary Citations in Reference
		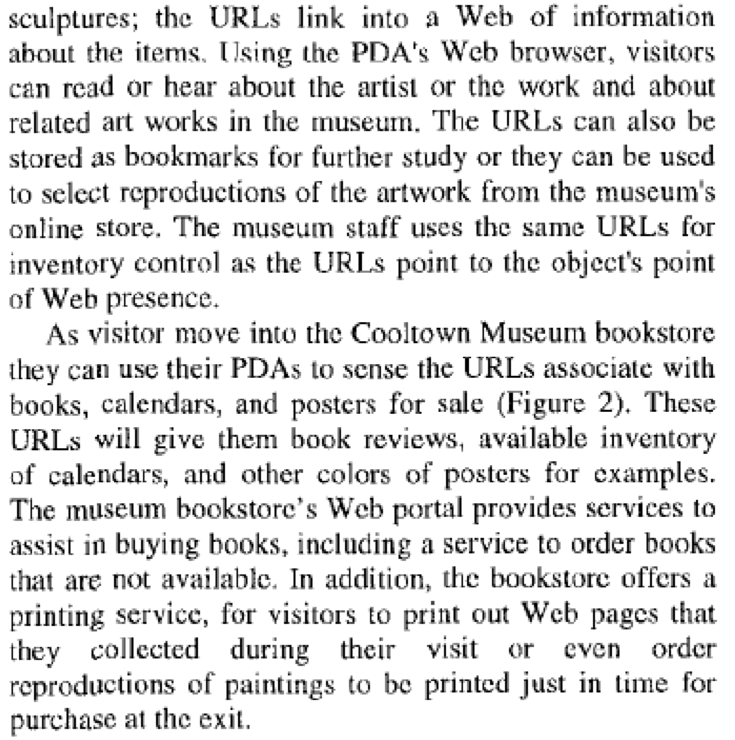 sculptures; the URLs link into a Web of information about the items. Using the PDA's Web browser, visitors can read or hear about the artist or the work and about related art works in the museum. The URLs can also be stored as bookmarks for further study or they can be used to select reproductions of the artwork from the museum's online store. The museum staff uses the same URLs for inventory control as the URLs point to the object's point of Web presence. As visitor move into the Cooltown Museum bookstore they can use their PDAs to sense the URLs associate with books, calendars, and posters for sale (Figure 2). These URLs will give them book reviews, available inventory of calendars, and other colors of posters for examples. The museum bookstore's Web portal provides services to assist in buying books, including a service to order books that are not available. In addition, the bookstore offers a printing service, for visitors to print out Web pages that they collected during their visit or even order reproductions of paintings to be printed just in time for purchase at the exit. Cooltown Developer's Network, People, Place, and Things: web presence for the real world (HP000025–34) (HP000025–34) at 23

Element	'252 Claim Recitation	Exemplary Citations in Reference
		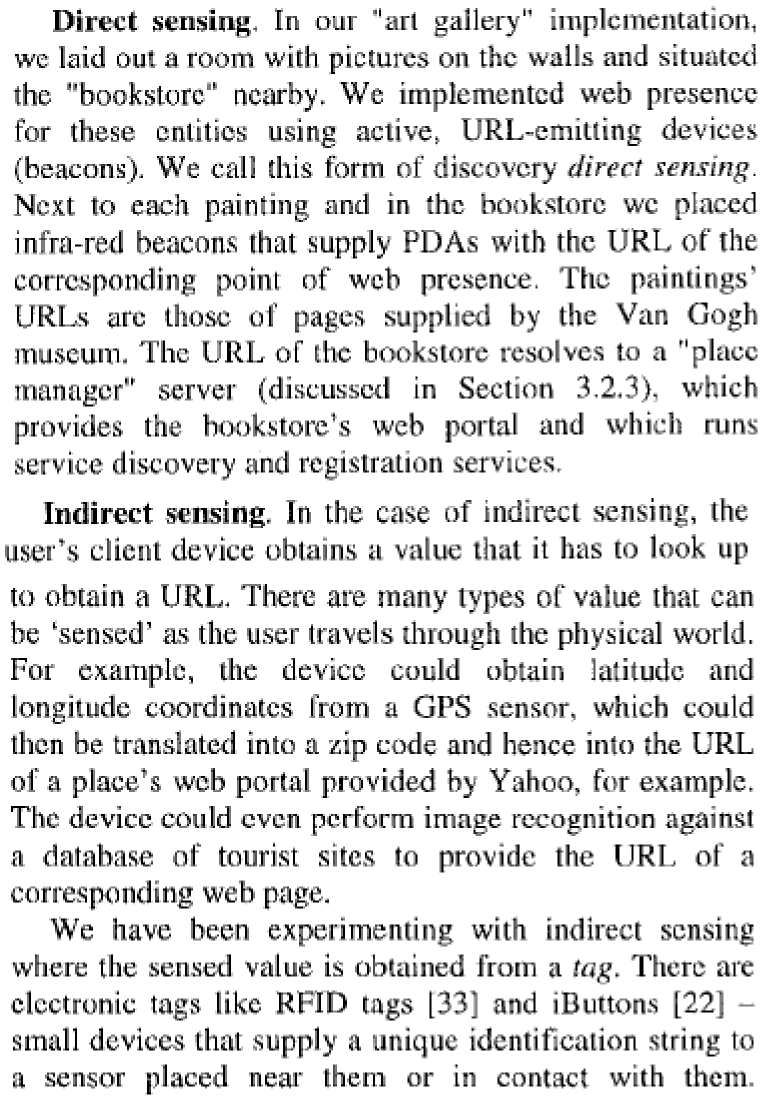 **Direct sensing**. In our "art gallery" implementation, we laid out a room with pictures on the walls and situated the "bookstore" nearby. We implemented web presence for these entities using active, URL-emitting devices (beacons). We call this form of discovery *direct sensing*. Next to each painting and in the bookstore we placed infra-red beacons that supply PDAs with the URL of the corresponding point of web presence. The paintings' URLs are those of pages supplied by the Van Gogh museum. The URL of the bookstore resolves to a "place manager" server (discussed in Section 3.2.3), which provides the bookstore's web portal and which runs service discovery and registration services. **Indirect sensing**. In the case of indirect sensing, the user's client device obtains a value that it has to look up to obtain a URL. There are many types of value that can be 'sensed' as the user travels through the physical world. For example, the device could obtain latitude and longitude coordinates from a GPS sensor, which could then be translated into a zip code and hence into the URL of a place's web portal provided by Yahoo, for example. The device could even perform image recognition against a database of tourist sites to provide the URL of a corresponding web page. We have been experimenting with indirect sensing where the sensed value is obtained from a *tag*. There are electronic tags like RFID tags [33] and iButtons [22] – small devices that supply a unique identification string to a sensor placed near them or in contact with them.

Element	'252 Claim Recitation	Exemplary Citations in Reference
		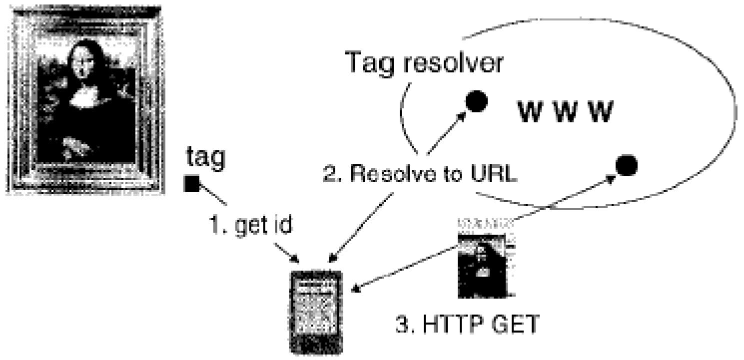 **Figure 4. Indirect sensing: the web presence for a painting.** 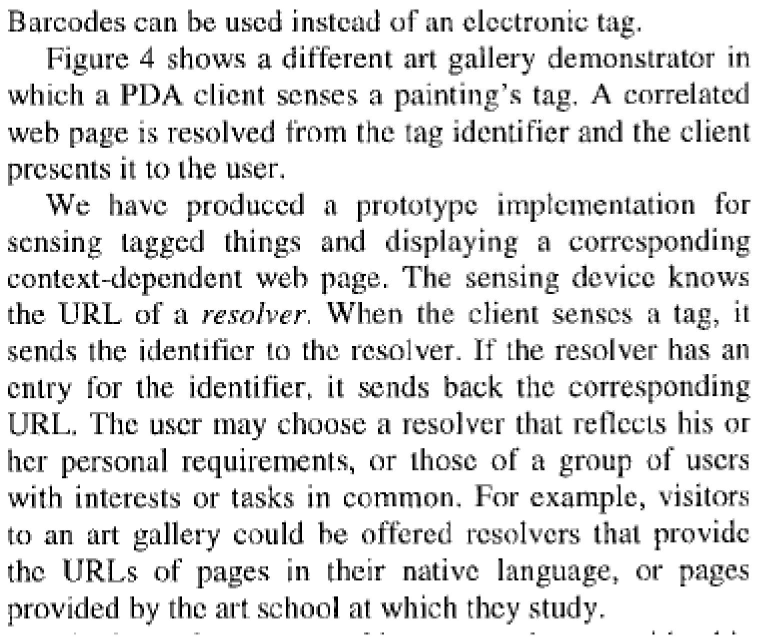 Barcodes can be used instead of an electronic tag. Figure 4 shows a different art gallery demonstrator in which a PDA client senses a painting's tag. A correlated web page is resolved from the tag identifier and the client presents it to the user. We have produced a prototype implementation for sensing tagged things and displaying a corresponding context-dependent web page. The sensing device knows the URL of a *resolver*. When the client senses a tag, it sends the identifier to the resolver. If the resolver has an entry for the identifier, it sends back the corresponding URL. The user may choose a resolver that reflects his or her personal requirements, or those of a group of users with interests or tasks in common. For example, visitors to an art gallery could be offered resolvers that provide the URLs of pages in their native language, or pages provided by the art school at which they study. The Cooltown User Experience (HP000088 – 093) at 4

Element	'252 Claim Recitation	Exemplary Citations in Reference
		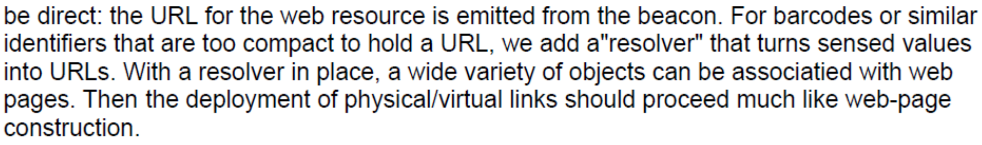 be direct: the URL for the web resource is emitted from the beacon. For barcodes or similar identifiers that are too compact to hold a URL, we add a"resolver" that turns sensed values into URLs. With a resolver in place, a wide variety of objects can be associatied with web pages. Then the deployment of physical/virtual links should proceed much like web-page construction. To the extent it is found that CoolTown does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that CoolTown does not anticipate this claim, CoolTown renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine CoolTown with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[18B.ii]	decoding the recognizable symbols to extract symbol information by analyzing the recognizable symbols according to type;	CoolTown discloses decoding the recognizable symbols to extract symbol information by analyzing the recognizable symbols according to type. A Web-Based Nomadic Computing System (HP000074 – 87) at 3 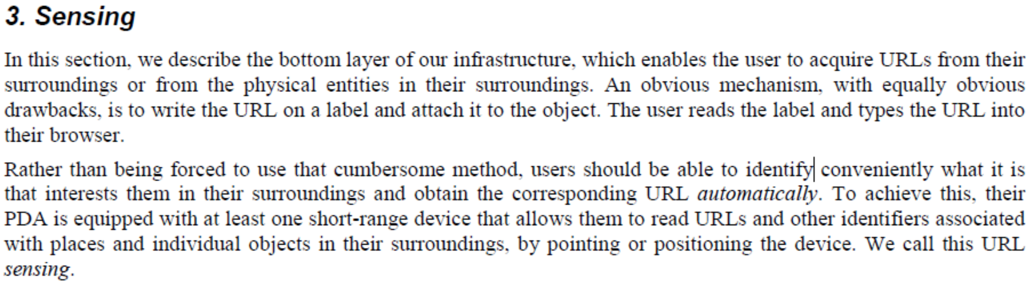 ***3. Sensing*** In this section, we describe the bottom layer of our infrastructure, which enables the user to acquire URLs from their surroundings or from the physical entities in their surroundings. An obvious mechanism, with equally obvious drawbacks, is to write the URL on a label and attach it to the object. The user reads the label and types the URL into their browser. Rather than being forced to use that cumbersome method, users should be able to identify conveniently what it is that interests them in their surroundings and obtain the corresponding URL *automatically*. To achieve this, their PDA is equipped with at least one short-range device that allows them to read URLs and other identifiers associated with places and individual objects in their surroundings, by pointing or positioning the device. We call this URL *sensing*. A Web-Based Nomadic Computing System (HP000074 – 87) at 4

Element	'252 Claim Recitation	Exemplary Citations in Reference
		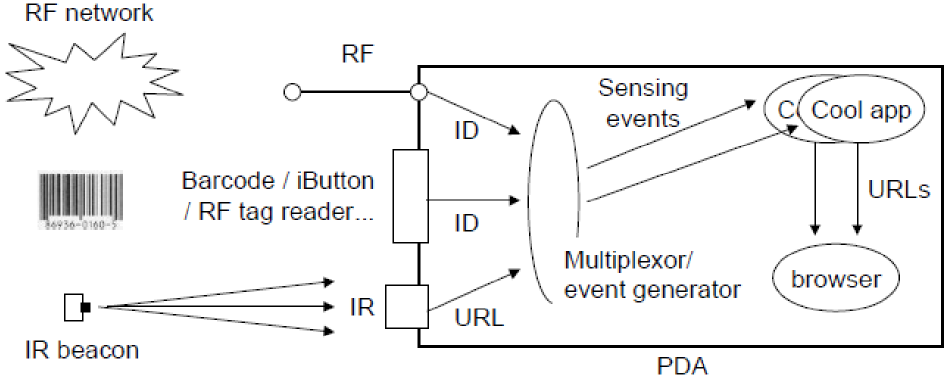 Figure 3. A sketch of a PDA equipped for sensing URLs and identifiers. ***3. Sensing*** In this section, we describe the bottom layer of our infrastructure, which enables the user to acquire URLs from their surroundings or from the physical entities in their surroundings. An obvious mechanism, with equally obvious drawbacks, is to write the URL on a label and attach it to the object. The user reads the label and types the URL into their browser. Rather than being forced to use that cumbersome method, users should be able to identify conveniently what it is that interests them in their surroundings and obtain the corresponding URL *automatically*. To achieve this, their PDA is equipped with at least one short-range device that allows them to read URLs and other identifiers associated with places and individual objects in their surroundings, by pointing or positioning the device. We call this URL *sensing*. Some options for sensing include: ∀ *IR & RF*. The infrared transceiver available on most PDAs and laptops can sense by communicating with a small, low-power 'beacon' device mounted at the place or on the object. Our beacon emits the identifier in infrared packets in the form of XML specifying a URL [15]. IR beacons can be configured to have a range of several meters. Similarly, short-range RF could be used. ∀ *Barcodes*. Barcode scanners are increasingly found integrated into PDAs, and are available as wands and other hand-held formats. Many objects are manufactured with barcodes – both generic ones and ones encoding their serial number; where they do not, barcodes can easily be printed and attached to objects. ∀ *Electronic tags*. A variety of small electronic tags exist containing a unique identifier with ~100 bits. All are available at relatively low cost. RFID tags can be read from a distance of about 10 cm with a hand-held reader. There are also tags such as iButtons [18] that are read on contact, which can be advantageous if tags lie close together. ∀ *Optical recognition*. The user's device may have a camera attached, in which case the identifier can be encoded in a form such as a 'glyph' [23] or digital watermark that is amenable to recognition software.

Element	'252 Claim Recitation	Exemplary Citations in Reference
		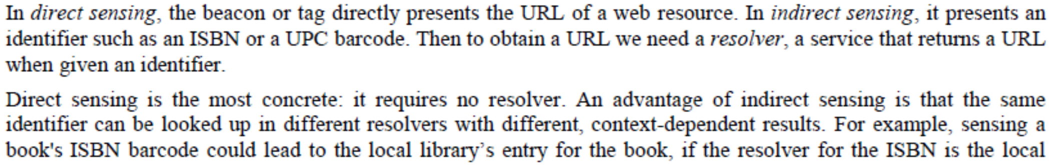 In *direct sensing*, the beacon or tag directly presents the URL of a web resource. In *indirect sensing*, it presents an identifier such as an ISBN or a UPC barcode. Then to obtain a URL we need a *resolver*, a service that returns a URL when given an identifier. Direct sensing is the most concrete: it requires no resolver. An advantage of indirect sensing is that the same identifier can be looked up in different resolvers with different, context-dependent results. For example, sensing a book's ISBN barcode could lead to the local library's entry for the book, if the resolver for the ISBN is the local A Web-Based Nomadic Computing System (HP000074 – 87) at 5 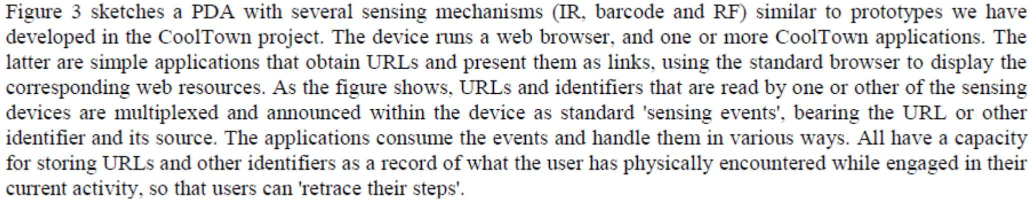 Figure 3 sketches a PDA with several sensing mechanisms (IR, barcode and RF) similar to prototypes we have developed in the CoolTown project. The device runs a web browser, and one or more CoolTown applications. The latter are simple applications that obtain URLs and present them as links, using the standard browser to display the corresponding web resources. As the figure shows, URLs and identifiers that are read by one or other of the sensing devices are multiplexed and announced within the device as standard 'sensing events', bearing the URL or other identifier and its source. The applications consume the events and handle them in various ways. All have a capacity for storing URLs and other identifiers as a record of what the user has physically encountered while engaged in their current activity, so that users can 'retrace their steps'. A Web-Based Nomadic Computing System (HP000074 – 87) at 5 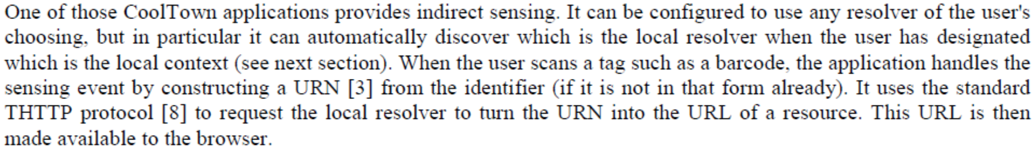 One of those CoolTown applications provides indirect sensing. It can be configured to use any resolver of the user's choosing, but in particular it can automatically discover which is the local resolver when the user has designated which is the local context (see next section). When the user scans a tag such as a barcode, the application handles the sensing event by constructing a URN [3] from the identifier (if it is not in that form already). It uses the standard THTTP protocol [8] to request the local resolver to turn the URN into the URL of a resource. This URL is then made available to the browser. A Web-Based Nomadic Computing System (HP000074 – 87) at 6 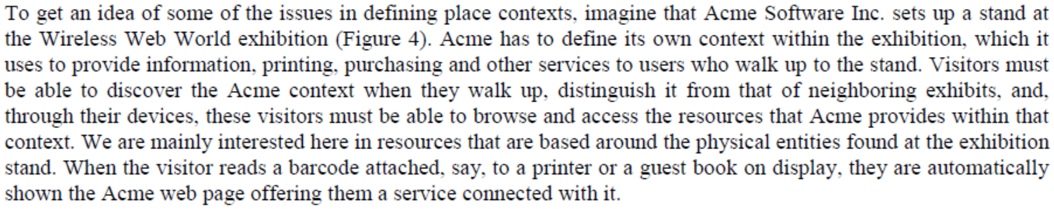 To get an idea of some of the issues in defining place contexts, imagine that Acme Software Inc. sets up a stand at the Wireless Web World exhibition (Figure 4). Acme has to define its own context within the exhibition, which it uses to provide information, printing, purchasing and other services to users who walk up to the stand. Visitors must be able to discover the Acme context when they walk up, distinguish it from that of neighboring exhibits, and, through their devices, these visitors must be able to browse and access the resources that Acme provides within that context. We are mainly interested here in resources that are based around the physical entities found at the exhibition stand. When the visitor reads a barcode attached, say, to a printer or a guest book on display, they are automatically shown the Acme web page offering them a service connected with it. A Web-Based Nomadic Computing System (HP000074 – 87) at 6

Element	'252 Claim Recitation	Exemplary Citations in Reference
		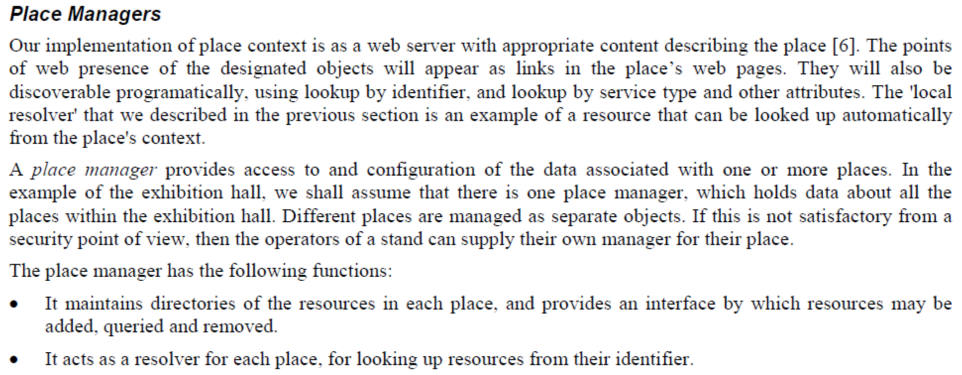 ***Place Managers*** Our implementation of place context is as a web server with appropriate content describing the place [6]. The points of web presence of the designated objects will appear as links in the place's web pages. They will also be discoverable programatically, using lookup by identifier, and lookup by service type and other attributes. The 'local resolver' that we described in the previous section is an example of a resource that can be looked up automatically from the place's context. A *place manager* provides access to and configuration of the data associated with one or more places. In the example of the exhibition hall, we shall assume that there is one place manager, which holds data about all the places within the exhibition hall. Different places are managed as separate objects. If this is not satisfactory from a security point of view, then the operators of a stand can supply their own manager for their place. The place manager has the following functions: • It maintains directories of the resources in each place, and provides an interface by which resources may be added, queried and removed. • It acts as a resolver for each place, for looking up resources from their identifier. Cooltown Developer's Network, People, Place, and Things: web presence for the real world (HP000025–34) (HP000025–34) at 20 – 21 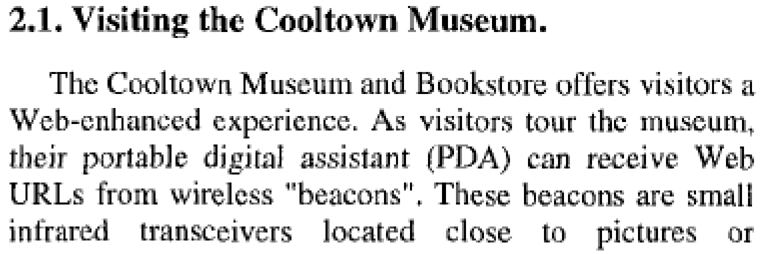 **2.1. Visiting the Cooltown Museum.** The Cooltown Museum and Bookstore offers visitors a Web-enhanced experience. As visitors tour the museum, their portable digital assistant (PDA) can receive Web URLs from wireless "beacons". These beacons are small infrared transceivers located close to pictures or

Element	'252 Claim Recitation	Exemplary Citations in Reference
		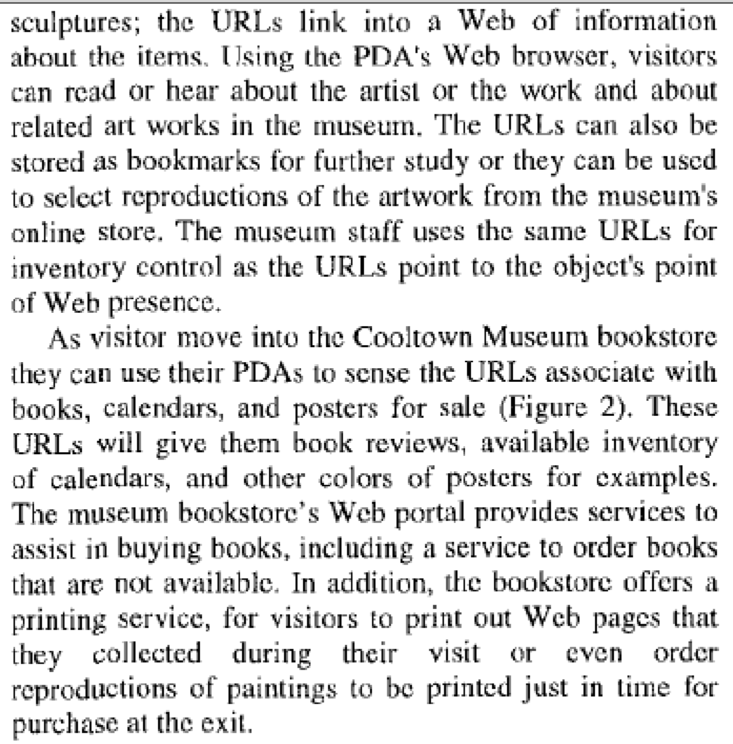 sculptures; the URLs link into a Web of information about the items. Using the PDA's Web browser, visitors can read or hear about the artist or the work and about related art works in the museum. The URLs can also be stored as bookmarks for further study or they can be used to select reproductions of the artwork from the museum's online store. The museum staff uses the same URLs for inventory control as the URLs point to the object's point of Web presence. As visitor move into the Cooltown Museum bookstore they can use their PDAs to sense the URLs associate with books, calendars, and posters for sale (Figure 2). These URLs will give them book reviews, available inventory of calendars, and other colors of posters for examples. The museum bookstore's Web portal provides services to assist in buying books, including a service to order books that are not available. In addition, the bookstore offers a printing service, for visitors to print out Web pages that they collected during their visit or even order reproductions of paintings to be printed just in time for purchase at the exit. Cooltown Developer's Network, People, Place, and Things: web presence for the real world (HP000025–34) (HP000025–34) at 23

Element	'252 Claim Recitation	Exemplary Citations in Reference
		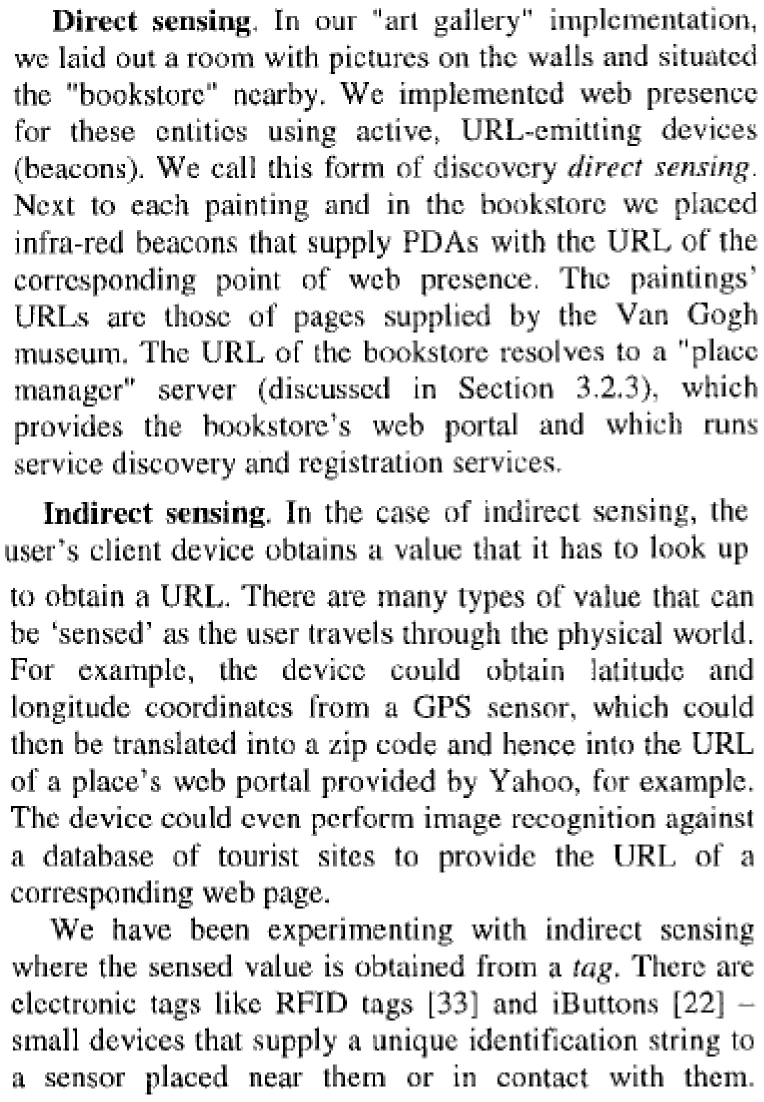 **Direct sensing**. In our "art gallery" implementation, we laid out a room with pictures on the walls and situated the "bookstore" nearby. We implemented web presence for these entities using active, URL-emitting devices (beacons). We call this form of discovery *direct sensing*. Next to each painting and in the bookstore we placed infra-red beacons that supply PDAs with the URL of the corresponding point of web presence. The paintings' URLs are those of pages supplied by the Van Gogh museum. The URL of the bookstore resolves to a "place manager" server (discussed in Section 3.2.3), which provides the bookstore's web portal and which runs service discovery and registration services. **Indirect sensing**. In the case of indirect sensing, the user's client device obtains a value that it has to look up to obtain a URL. There are many types of value that can be 'sensed' as the user travels through the physical world. For example, the device could obtain latitude and longitude coordinates from a GPS sensor, which could then be translated into a zip code and hence into the URL of a place's web portal provided by Yahoo, for example. The device could even perform image recognition against a database of tourist sites to provide the URL of a corresponding web page. We have been experimenting with indirect sensing where the sensed value is obtained from a *tag*. There are electronic tags like RFID tags [33] and iButtons [22] – small devices that supply a unique identification string to a sensor placed near them or in contact with them.

Element	'252 Claim Recitation	Exemplary Citations in Reference
		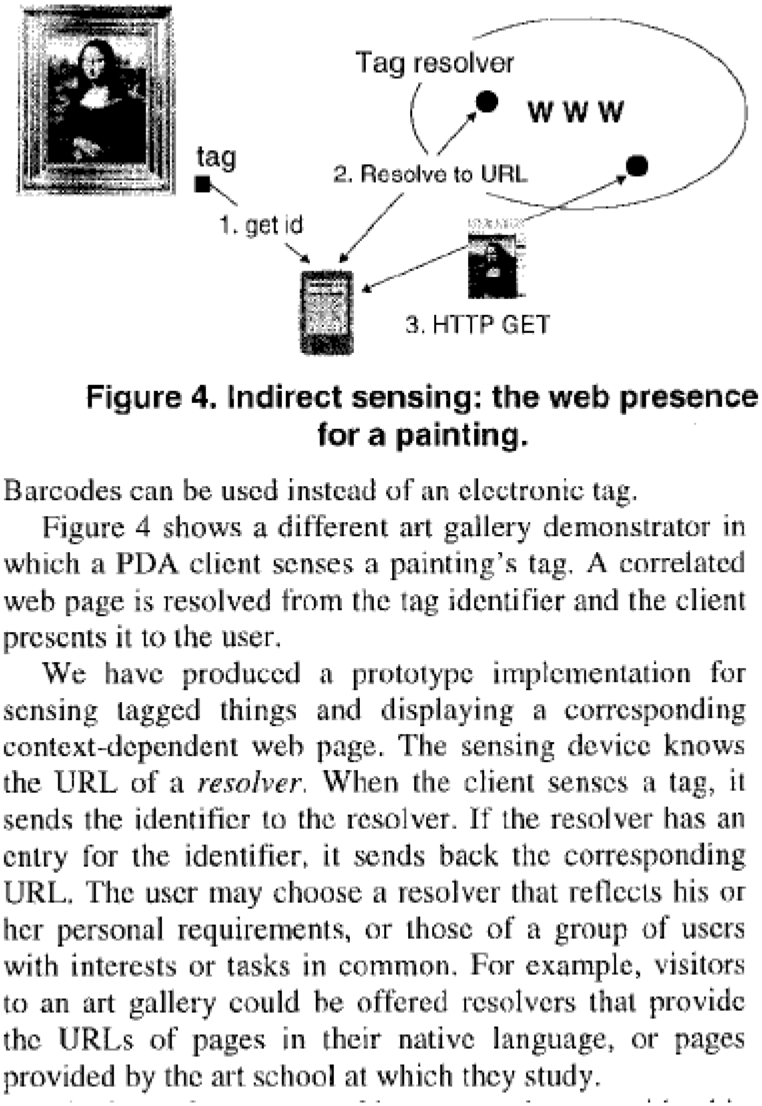 **Figure 4. Indirect sensing: the web presence for a painting.** Barcodes can be used instead of an electronic tag. Figure 4 shows a different art gallery demonstrator in which a PDA client senses a painting's tag. A correlated web page is resolved from the tag identifier and the client presents it to the user. We have produced a prototype implementation for sensing tagged things and displaying a corresponding context-dependent web page. The sensing device knows the URL of a *resolver*. When the client senses a tag, it sends the identifier to the resolver. If the resolver has an entry for the identifier, it sends back the corresponding URL. The user may choose a resolver that reflects his or her personal requirements, or those of a group of users with interests or tasks in common. For example, visitors to an art gallery could be offered resolvers that provide the URLs of pages in their native language, or pages provided by the art school at which they study. The Cooltown User Experience (HP000088 – 093) at 4

Element	'252 Claim Recitation	Exemplary Citations in Reference
		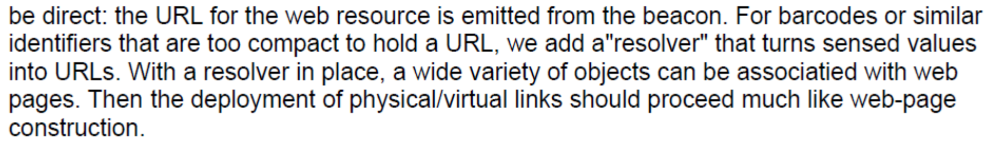 be direct: the URL for the web resource is emitted from the beacon. For barcodes or similar identifiers that are too compact to hold a URL, we add a"resolver" that turns sensed values into URLs. With a resolver in place, a wide variety of objects can be associatied with web pages. Then the deployment of physical/virtual links should proceed much like web-page construction. To the extent it is found that CoolTown does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that CoolTown does not anticipate this claim, CoolTown renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine CoolTown with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[18C]	providing access to a distal server configured to use a database to identify pertinent information associated with the recognizable symbols based on the symbol information; and	CoolTown discloses providing access to a distal server configured to use a database to identify pertinent information associated with the recognizable symbols based on the symbol information. A Web-Based Nomadic Computing System (HP000074 – 87) at Fig. 2

Element	'252 Claim Recitation	Exemplary Citations in Reference
		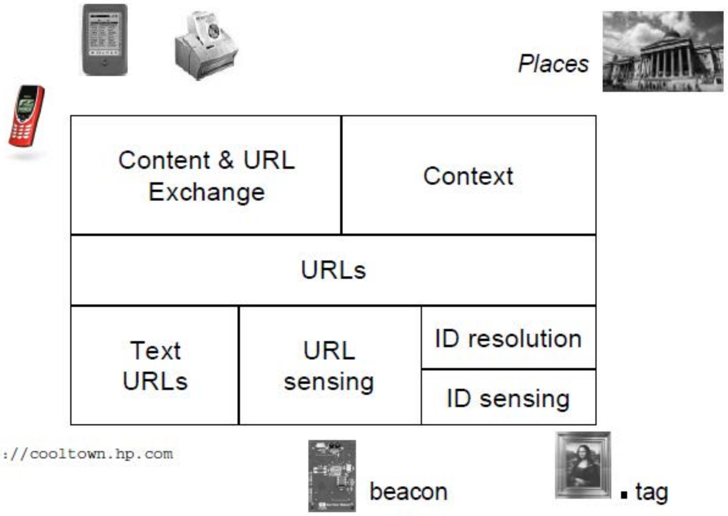 Figure 2. Component relationships. Sources of URLs, including keyboards, beacons, and resolved tags point to context or content. A Web-Based Nomadic Computing System (HP000074 – 87) at 3 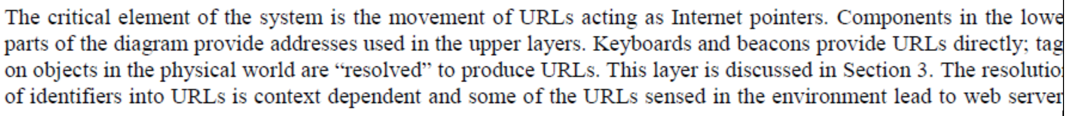 The critical element of the system is the movement of URLs acting as Internet pointers. Components in the lowe parts of the diagram provide addresses used in the upper layers. Keyboards and beacons provide URLs directly; tag on objects in the physical world are "resolved" to produce URLs. This layer is discussed in Section 3. The resolutio of identifiers into URLs is context dependent and some of the URLs sensed in the environment lead to web server A Web-Based Nomadic Computing System (HP000074 – 87) at 4

Element	'252 Claim Recitation	Exemplary Citations in Reference
		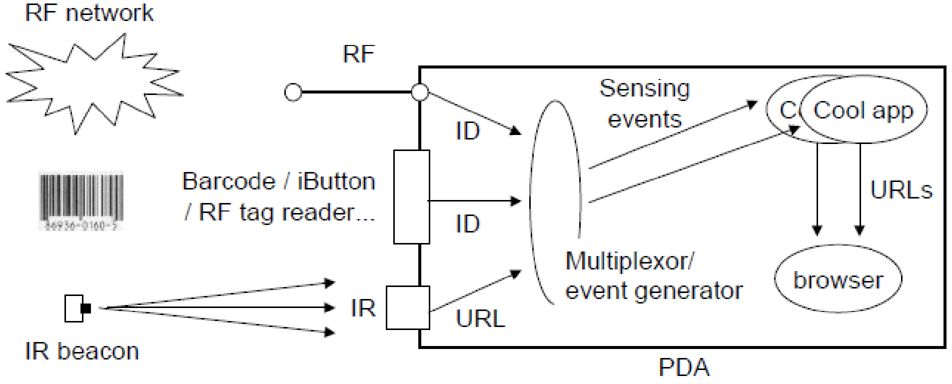 Figure 3. A sketch of a PDA equipped for sensing URLs and identifiers. A Web-Based Nomadic Computing System (HP000074 – 87) at 4–5 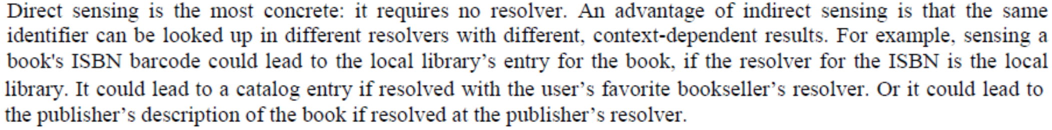 Direct sensing is the most concrete: it requires no resolver. An advantage of indirect sensing is that the same identifier can be looked up in different resolvers with different, context-dependent results. For example, sensing a book's ISBN barcode could lead to the local library's entry for the book, if the resolver for the ISBN is the local library. It could lead to a catalog entry if resolved with the user's favorite bookseller's resolver. Or it could lead to the publisher's description of the book if resolved at the publisher's resolver. A Web-Based Nomadic Computing System (HP000074 – 87) at 5 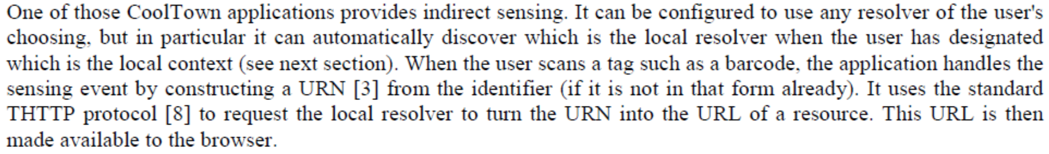 One of those CoolTown applications provides indirect sensing. It can be configured to use any resolver of the user's choosing, but in particular it can automatically discover which is the local resolver when the user has designated which is the local context (see next section). When the user scans a tag such as a barcode, the application handles the sensing event by constructing a URN [3] from the identifier (if it is not in that form already). It uses the standard THTTP protocol [8] to request the local resolver to turn the URN into the URL of a resource. This URL is then made available to the browser. A Web-Based Nomadic Computing System (HP000074 – 87) at 6 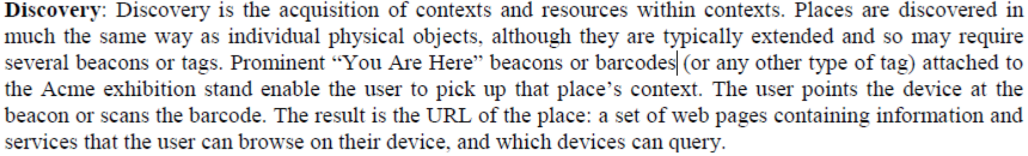 **Discovery**: Discovery is the acquisition of contexts and resources within contexts. Places are discovered in much the same way as individual physical objects, although they are typically extended and so may require several beacons or tags. Prominent "You Are Here" beacons or barcodes (or any other type of tag) attached to the Acme exhibition stand enable the user to pick up that place's context. The user points the device at the beacon or scans the barcode. The result is the URL of the place: a set of web pages containing information and services that the user can browse on their device, and which devices can query. A Web-Based Nomadic Computing System (HP000074 – 87) at 6

Element	'252 Claim Recitation	Exemplary Citations in Reference
		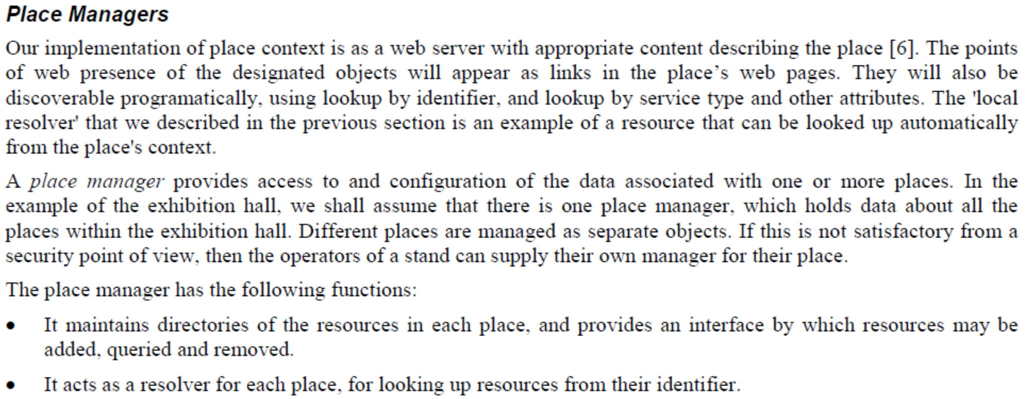 ***Place Managers*** Our implementation of place context is as a web server with appropriate content describing the place [6]. The points of web presence of the designated objects will appear as links in the place's web pages. They will also be discoverable programatically, using lookup by identifier, and lookup by service type and other attributes. The 'local resolver' that we described in the previous section is an example of a resource that can be looked up automatically from the place's context. A *place manager* provides access to and configuration of the data associated with one or more places. In the example of the exhibition hall, we shall assume that there is one place manager, which holds data about all the places within the exhibition hall. Different places are managed as separate objects. If this is not satisfactory from a security point of view, then the operators of a stand can supply their own manager for their place. The place manager has the following functions: • It maintains directories of the resources in each place, and provides an interface by which resources may be added, queried and removed. • It acts as a resolver for each place, for looking up resources from their identifier. Cooltown Developer's Network, People, Place, and Things: web presence for the real world (HP000025–34) (HP000025–34) at 20 – 21 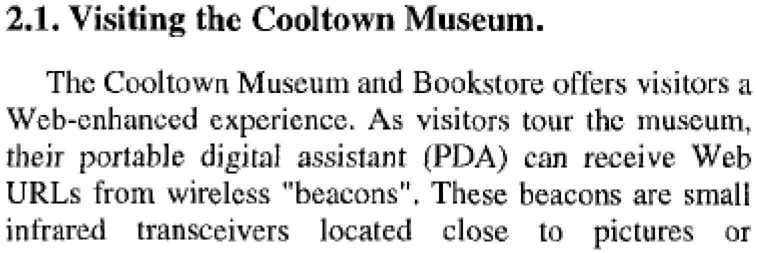 **2.1. Visiting the Cooltown Museum.** The Cooltown Museum and Bookstore offers visitors a Web-enhanced experience. As visitors tour the museum, their portable digital assistant (PDA) can receive Web URLs from wireless "beacons". These beacons are small infrared transceivers located close to pictures or

Element	'252 Claim Recitation	Exemplary Citations in Reference
		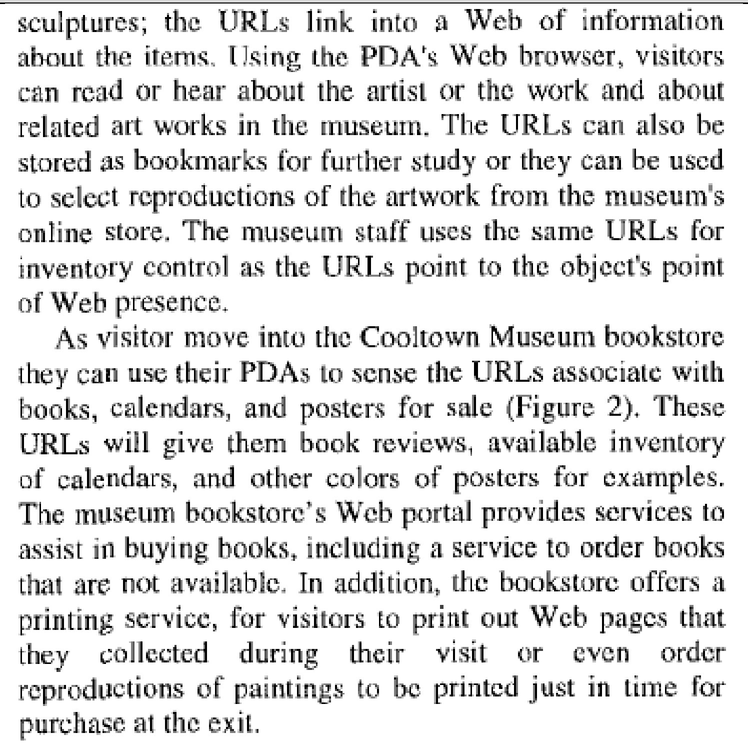 sculptures; the URLs link into a Web of information about the items. Using the PDA's Web browser, visitors can read or hear about the artist or the work and about related art works in the museum. The URLs can also be stored as bookmarks for further study or they can be used to select reproductions of the artwork from the museum's online store. The museum staff uses the same URLs for inventory control as the URLs point to the object's point of Web presence. As visitor move into the Cooltown Museum bookstore they can use their PDAs to sense the URLs associate with books, calendars, and posters for sale (Figure 2). These URLs will give them book reviews, available inventory of calendars, and other colors of posters for examples. The museum bookstore's Web portal provides services to assist in buying books, including a service to order books that are not available. In addition, the bookstore offers a printing service, for visitors to print out Web pages that they collected during their visit or even order reproductions of paintings to be printed just in time for purchase at the exit. Cooltown Developer's Network, People, Place, and Things: web presence for the real world (HP000025–34) (HP000025–34) at 23

Element	'252 Claim Recitation	Exemplary Citations in Reference
		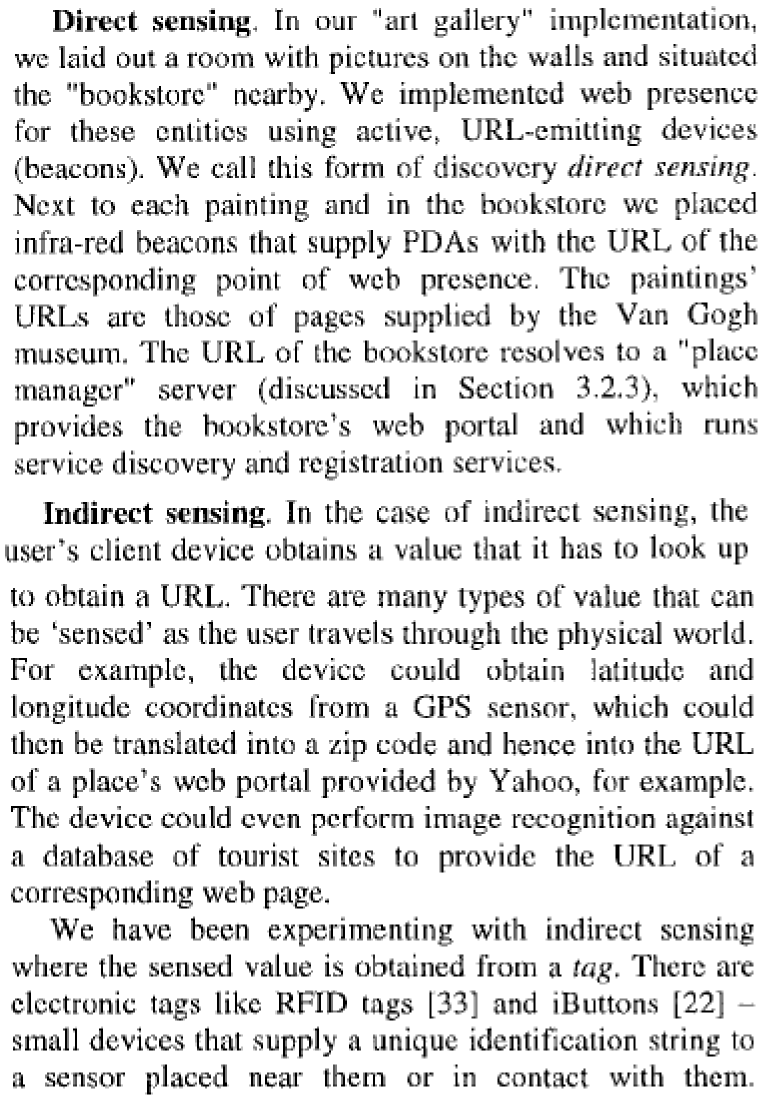 **Direct sensing**. In our "art gallery" implementation, we laid out a room with pictures on the walls and situated the "bookstore" nearby. We implemented web presence for these entities using active, URL-emitting devices (beacons). We call this form of discovery *direct sensing*. Next to each painting and in the bookstore we placed infra-red beacons that supply PDAs with the URL of the corresponding point of web presence. The paintings' URLs are those of pages supplied by the Van Gogh museum. The URL of the bookstore resolves to a "place manager" server (discussed in Section 3.2.3), which provides the bookstore's web portal and which runs service discovery and registration services. **Indirect sensing**. In the case of indirect sensing, the user's client device obtains a value that it has to look up to obtain a URL. There are many types of value that can be 'sensed' as the user travels through the physical world. For example, the device could obtain latitude and longitude coordinates from a GPS sensor, which could then be translated into a zip code and hence into the URL of a place's web portal provided by Yahoo, for example. The device could even perform image recognition against a database of tourist sites to provide the URL of a corresponding web page. We have been experimenting with indirect sensing where the sensed value is obtained from a *tag*. There are electronic tags like RFID tags [33] and iButtons [22] – small devices that supply a unique identification string to a sensor placed near them or in contact with them.

Element	'252 Claim Recitation	Exemplary Citations in Reference
		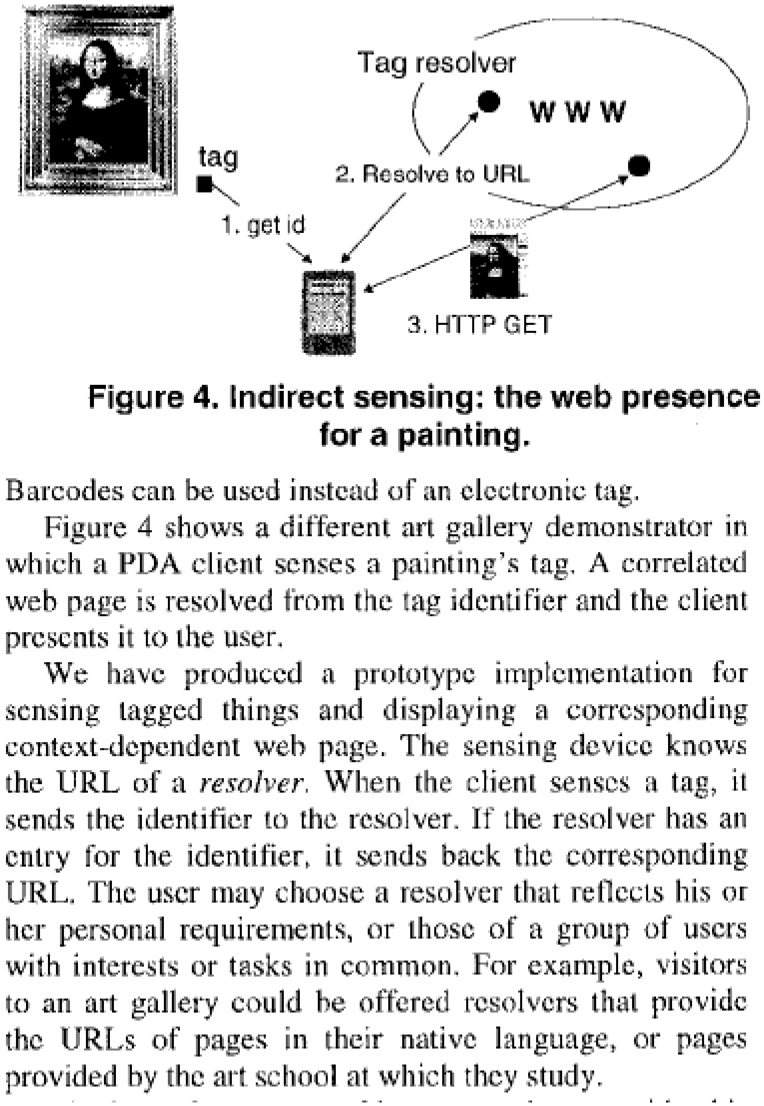 **Figure 4. Indirect sensing: the web presence for a painting.** Barcodes can be used instead of an electronic tag. Figure 4 shows a different art gallery demonstrator in which a PDA client senses a painting's tag. A correlated web page is resolved from the tag identifier and the client presents it to the user. We have produced a prototype implementation for sensing tagged things and displaying a corresponding context-dependent web page. The sensing device knows the URL of a *resolver*. When the client senses a tag, it sends the identifier to the resolver. If the resolver has an entry for the identifier, it sends back the corresponding URL. The user may choose a resolver that reflects his or her personal requirements, or those of a group of users with interests or tasks in common. For example, visitors to an art gallery could be offered resolvers that provide the URLs of pages in their native language, or pages provided by the art school at which they study. A Web-Based Nomadic Computing System (HP000074 – 87) at 6

Element	'252 Claim Recitation	Exemplary Citations in Reference
		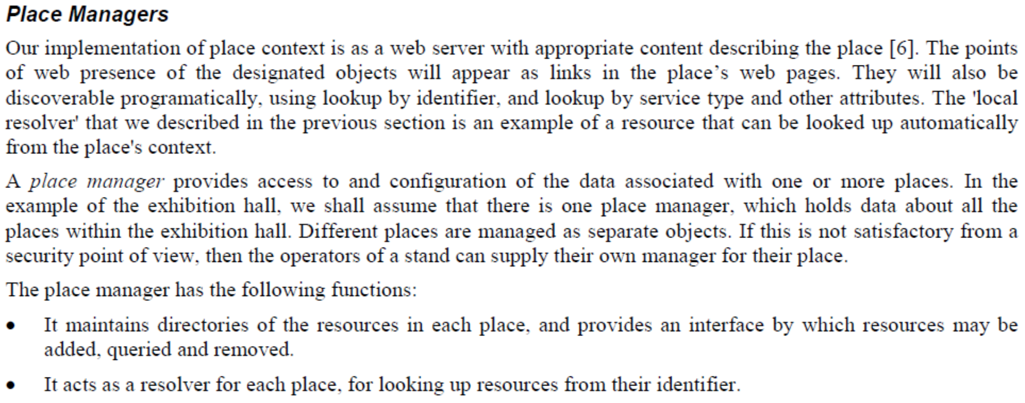 ***Place Managers*** Our implementation of place context is as a web server with appropriate content describing the place [6]. The points of web presence of the designated objects will appear as links in the place's web pages. They will also be discoverable programatically, using lookup by identifier, and lookup by service type and other attributes. The 'local resolver' that we described in the previous section is an example of a resource that can be looked up automatically from the place's context. A *place manager* provides access to and configuration of the data associated with one or more places. In the example of the exhibition hall, we shall assume that there is one place manager, which holds data about all the places within the exhibition hall. Different places are managed as separate objects. If this is not satisfactory from a security point of view, then the operators of a stand can supply their own manager for their place. The place manager has the following functions: • It maintains directories of the resources in each place, and provides an interface by which resources may be added, queried and removed. • It acts as a resolver for each place, for looking up resources from their identifier. The Cooltown User Experience (HP000088 – 093) at 4 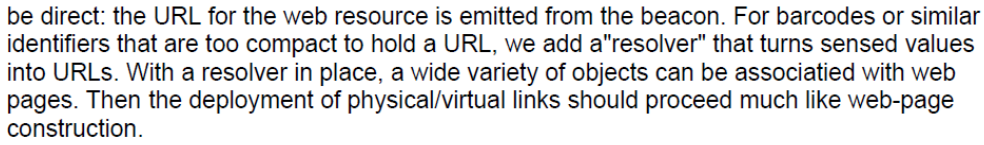 be direct: the URL for the web resource is emitted from the beacon. For barcodes or similar identifiers that are too compact to hold a URL, we add a"resolver" that turns sensed values into URLs. With a resolver in place, a wide variety of objects can be associatied with web pages. Then the deployment of physical/virtual links should proceed much like web-page construction. To the extent it is found that CoolTown does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that CoolTown does not anticipate this claim, CoolTown renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine CoolTown with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

Element	'252 Claim Recitation	Exemplary Citations in Reference
[18D]	allowing the mobile device to receive the pertinent information over a network.	CoolTown discloses allowing the mobile device to receive the pertinent information over a network. A Web-Based Nomadic Computing System (HP000074 – 87) at Fig. 2 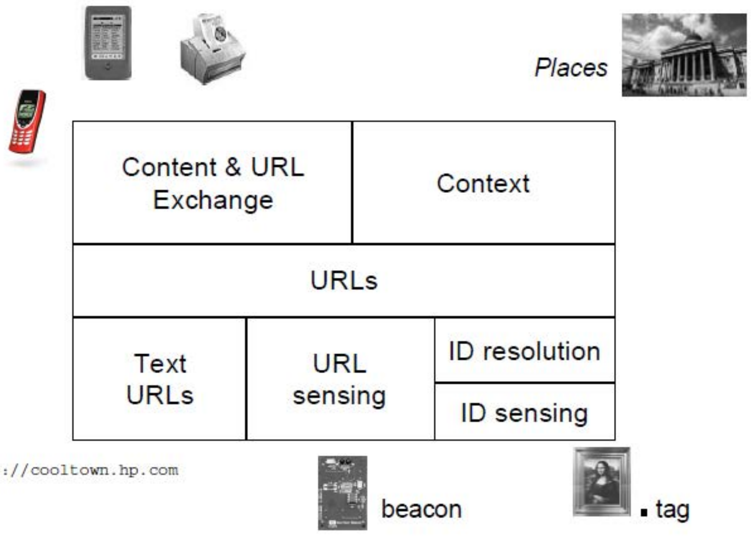 Figure 2. Component relationships. Sources of URLs, including keyboards, beacons, and resolved tags point to context or content. A Web-Based Nomadic Computing System (HP000074 – 87) at 4

Element	'252 Claim Recitation	Exemplary Citations in Reference
		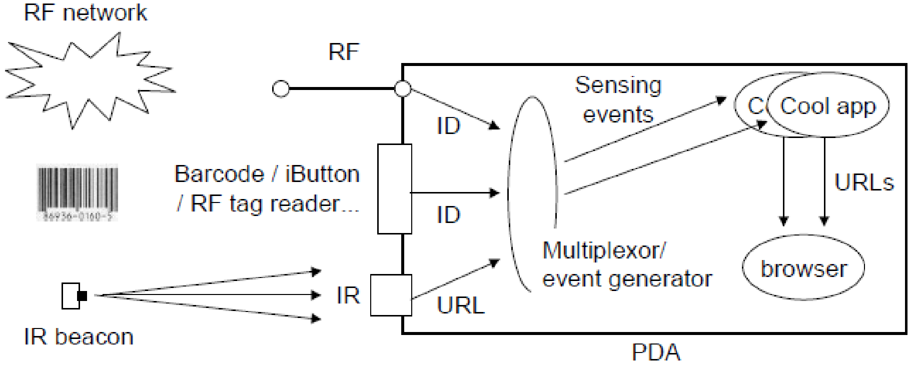 Figure 3. A sketch of a PDA equipped for sensing URLs and identifiers. A Web-Based Nomadic Computing System (HP000074 – 87) at 5 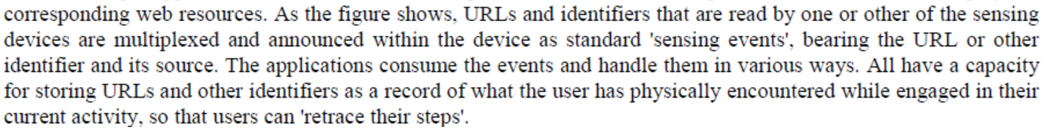Figure 3 sketches a PDA with several sensing mechanisms (IR, barcode and RF) similar to prototypes we have developed in the CoolTown project. The device runs a web browser, and one or more CoolTown applications. The latter are simple applications that obtain URLs and present them as links, using the standard browser to display the corresponding web resources. As the figure shows, URLs and identifiers that are read by one or other of the sensing devices are multiplexed and announced within the device as standard 'sensing events', bearing the URL or other identifier and its source. The applications consume the events and handle them in various ways. All have a capacity for storing URLs and other identifiers as a record of what the user has physically encountered while engaged in their current activity, so that users can 'retrace their steps'. A Web-Based Nomadic Computing System (HP000074 – 87) at 5 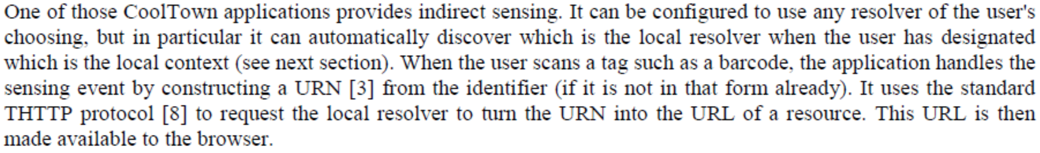One of those CoolTown applications provides indirect sensing. It can be configured to use any resolver of the user's choosing, but in particular it can automatically discover which is the local resolver when the user has designated which is the local context (see next section). When the user scans a tag such as a barcode, the application handles the sensing event by constructing a URN [3] from the identifier (if it is not in that form already). It uses the standard THTTP protocol [8] to request the local resolver to turn the URN into the URL of a resource. This URL is then made available to the browser. A Web-Based Nomadic Computing System (HP000074 – 87) at 6 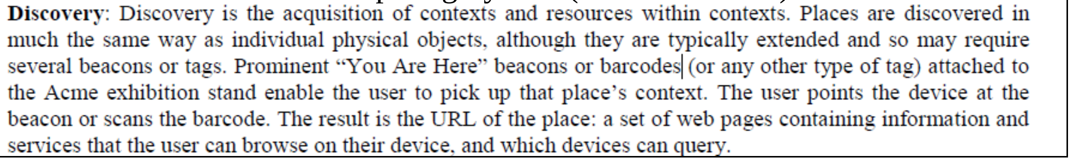**Discovery**: Discovery is the acquisition of contexts and resources within contexts. Places are discovered in much the same way as individual physical objects, although they are typically extended and so may require several beacons or tags. Prominent "You Are Here" beacons or barcodes (or any other type of tag) attached to the Acme exhibition stand enable the user to pick up that place's context. The user points the device at the beacon or scans the barcode. The result is the URL of the place: a set of web pages containing information and services that the user can browse on their device, and which devices can query.

Element	'252 Claim Recitation	Exemplary Citations in Reference
		Cooltown Developer's Network, People, Place, and Things: web presence for the real world (HP000025–34) at 5 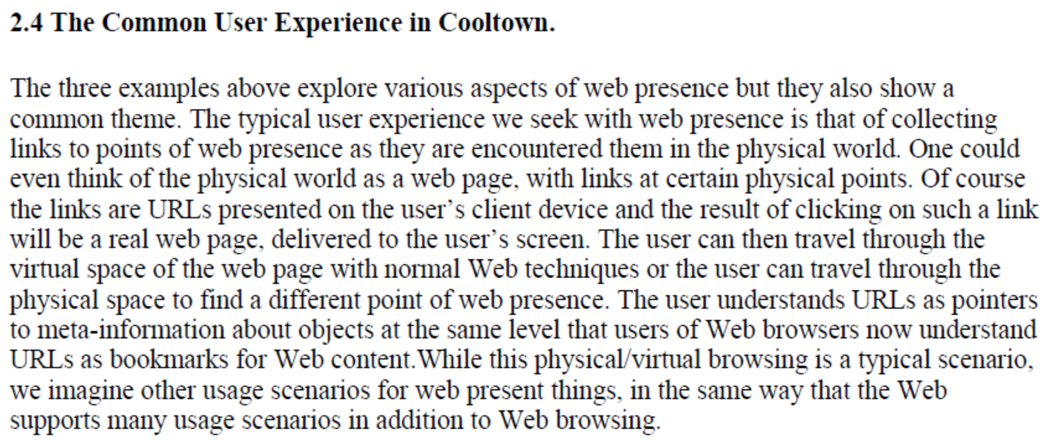 **2.4 The Common User Experience in Cooltown.** The three examples above explore various aspects of web presence but they also show a common theme. The typical user experience we seek with web presence is that of collecting links to points of web presence as they are encountered them in the physical world. One could even think of the physical world as a web page, with links at certain physical points. Of course the links are URLs presented on the user's client device and the result of clicking on such a link will be a real web page, delivered to the user's screen. The user can then travel through the virtual space of the web page with normal Web techniques or the user can travel through the physical space to find a different point of web presence. The user understands URLs as pointers to meta-information about objects at the same level that users of Web browsers now understand URLs as bookmarks for Web content.While this physical/virtual browsing is a typical scenario, we imagine other usage scenarios for web present things, in the same way that the Web supports many usage scenarios in addition to Web browsing. Cooltown Developer's Network, People, Place, and Things: web presence for the real world (HP000025–34) at 9

Element	'252 Claim Recitation	Exemplary Citations in Reference
		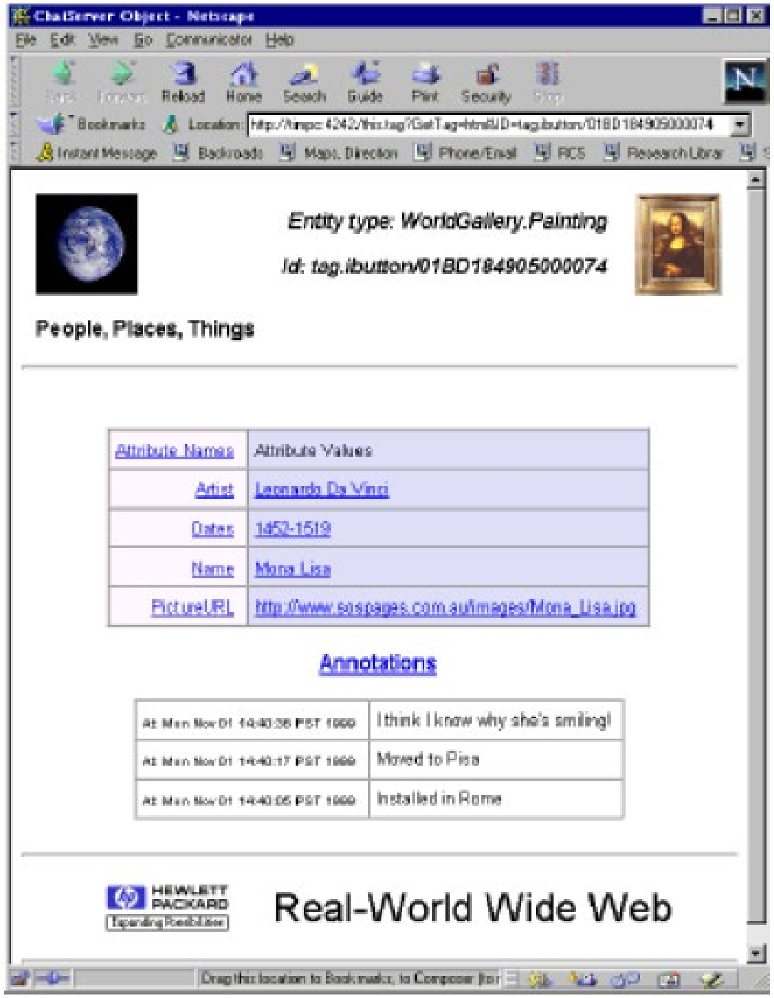 *Entity type: WorldGallery.Painting* *Id: tag.ibutton/01BD184905000074* **People, Places, Things** Attribute Names — Attribute Values Artist — Leonardo Da Vinci Dates — 1452-1519 Name — Mona Lisa PictureURL — http://www.sospages.com.au/images/Mona_Lisa.jpg Annotations At Mon Nov 01 14:40:36 PST 1999 — I think I know why she's smiling! At Mon Nov 01 14:40:17 PST 1999 — Moved to Pisa At Mon Nov 01 14:40:05 PST 1999 — Installed in Rome **Figure 5. A Web presence for *Mona Lisa*.** Cooltown Developer's Network, People, Place, and Things: web presence for the real world (HP000025–34) (HP000025–34) at 23

Element	'252 Claim Recitation	Exemplary Citations in Reference
		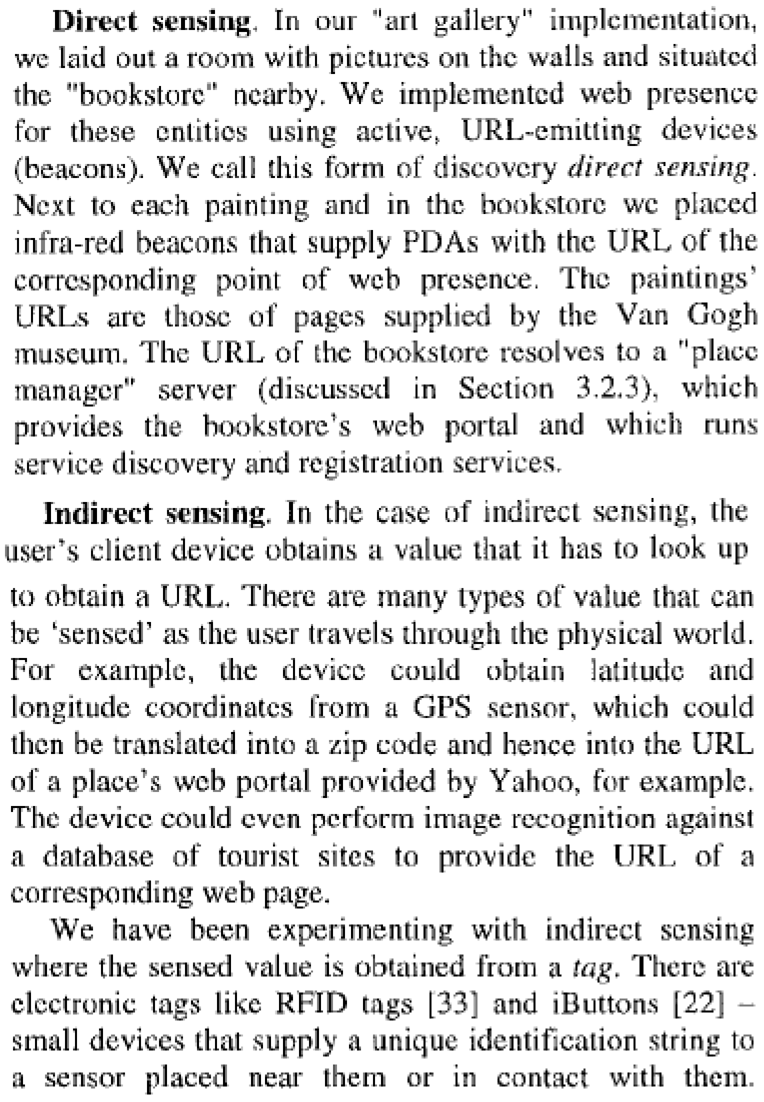 **Direct sensing**. In our "art gallery" implementation, we laid out a room with pictures on the walls and situated the "bookstore" nearby. We implemented web presence for these entities using active, URL-emitting devices (beacons). We call this form of discovery *direct sensing*. Next to each painting and in the bookstore we placed infra-red beacons that supply PDAs with the URL of the corresponding point of web presence. The paintings' URLs are those of pages supplied by the Van Gogh museum. The URL of the bookstore resolves to a "place manager" server (discussed in Section 3.2.3), which provides the bookstore's web portal and which runs service discovery and registration services. **Indirect sensing**. In the case of indirect sensing, the user's client device obtains a value that it has to look up to obtain a URL. There are many types of value that can be 'sensed' as the user travels through the physical world. For example, the device could obtain latitude and longitude coordinates from a GPS sensor, which could then be translated into a zip code and hence into the URL of a place's web portal provided by Yahoo, for example. The device could even perform image recognition against a database of tourist sites to provide the URL of a corresponding web page. We have been experimenting with indirect sensing where the sensed value is obtained from a *tag*. There are electronic tags like RFID tags [33] and iButtons [22] – small devices that supply a unique identification string to a sensor placed near them or in contact with them.

Element	'252 Claim Recitation	Exemplary Citations in Reference
		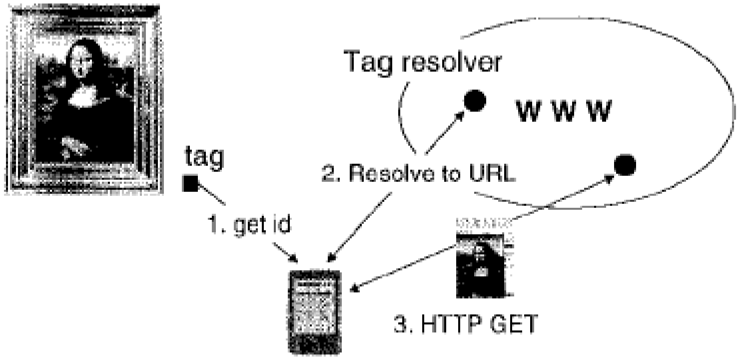 Figure 4. Indirect sensing: the web presence for a painting. Barcodes can be used instead of an electronic tag. Figure 4 shows a different art gallery demonstrator in which a PDA client senses a painting's tag. A correlated web page is resolved from the tag identifier and the client presents it to the user. We have produced a prototype implementation for sensing tagged things and displaying a corresponding context-dependent web page. The sensing device knows the URL of a *resolver*. When the client senses a tag, it sends the identifier to the resolver. If the resolver has an entry for the identifier, it sends back the corresponding URL. The user may choose a resolver that reflects his or her personal requirements, or those of a group of users with interests or tasks in common. For example, visitors to an art gallery could be offered resolvers that provide the URLs of pages in their native language, or pages provided by the art school at which they study.

Element	'252 Claim Recitation	Exemplary Citations in Reference
		To the extent it is found that CoolTown does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that CoolTown does not anticipate this claim, CoolTown renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine CoolTown with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

2. Claim 26

Claim	'252 Claim Recitation	Exemplary Citations in Reference
26	The method of claim 18, wherein the data relating to the image comprises the image.	CoolTown discloses the method of claim 18, wherein the data relating to the image comprises the image. A Web-Based Nomadic Computing System (HP000074 – 87) at Fig. 3

		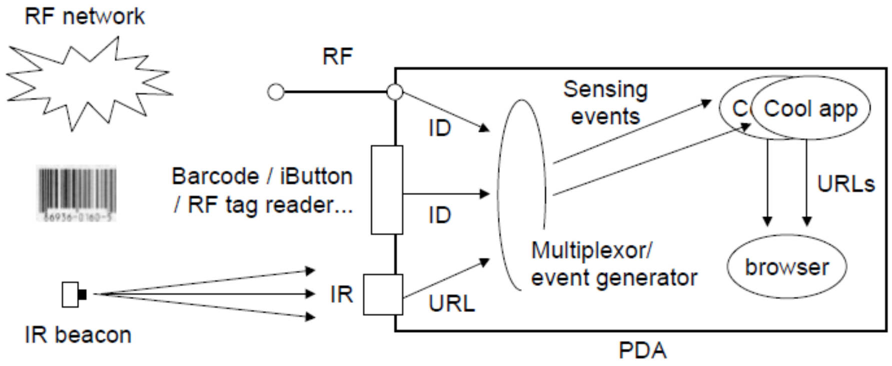 Figure 3. A sketch of a PDA equipped for sensing URLs and identifiers. A Web-Based Nomadic Computing System (HP000074 – 87) at 4 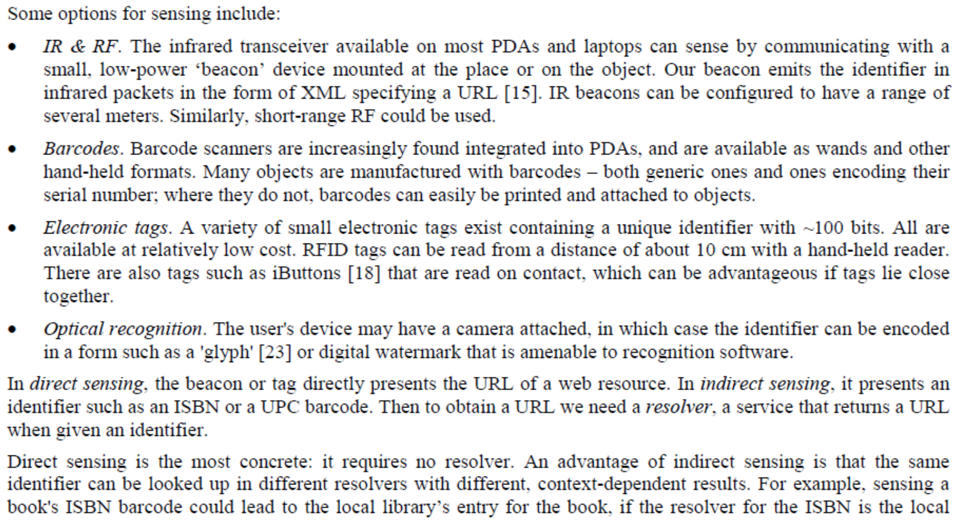 Some options for sensing include: - *IR & RF*. The infrared transceiver available on most PDAs and laptops can sense by communicating with a small, low-power 'beacon' device mounted at the place or on the object. Our beacon emits the identifier in infrared packets in the form of XML specifying a URL [15]. IR beacons can be configured to have a range of several meters. Similarly, short-range RF could be used. - *Barcodes*. Barcode scanners are increasingly found integrated into PDAs, and are available as wands and other hand-held formats. Many objects are manufactured with barcodes – both generic ones and ones encoding their serial number; where they do not, barcodes can easily be printed and attached to objects. - *Electronic tags*. A variety of small electronic tags exist containing a unique identifier with ~100 bits. All are available at relatively low cost. RFID tags can be read from a distance of about 10 cm with a hand-held reader. There are also tags such as iButtons [18] that are read on contact, which can be advantageous if tags lie close together. - *Optical recognition*. The user's device may have a camera attached, in which case the identifier can be encoded in a form such as a 'glyph' [23] or digital watermark that is amenable to recognition software. In *direct sensing*, the beacon or tag directly presents the URL of a web resource. In *indirect sensing*, it presents an identifier such as an ISBN or a UPC barcode. Then to obtain a URL we need a *resolver*, a service that returns a URL when given an identifier. Direct sensing is the most concrete: it requires no resolver. An advantage of indirect sensing is that the same identifier can be looked up in different resolvers with different, context-dependent results. For example, sensing a book's ISBN barcode could lead to the local library's entry for the book, if the resolver for the ISBN is the local A Web-Based Nomadic Computing System (HP000074 – 87) at 8

		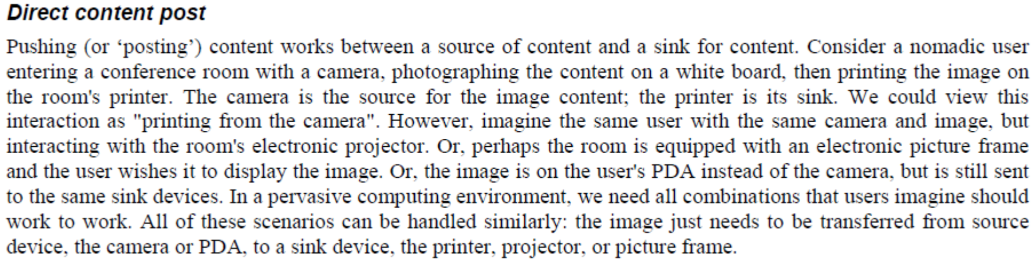 ***Direct content post*** Pushing (or 'posting') content works between a source of content and a sink for content. Consider a nomadic user entering a conference room with a camera, photographing the content on a white board, then printing the image on the room's printer. The camera is the source for the image content; the printer is its sink. We could view this interaction as "printing from the camera". However, imagine the same user with the same camera and image, but interacting with the room's electronic projector. Or, perhaps the room is equipped with an electronic picture frame and the user wishes it to display the image. Or, the image is on the user's PDA instead of the camera, but is still sent to the same sink devices. In a pervasive computing environment, we need all combinations that users imagine should work to work. All of these scenarios can be handled similarly: the image just needs to be transferred from source device, the camera or PDA, to a sink device, the printer, projector, or picture frame. Cooltown Developer's Network, People, Place, and Things: web presence for the real world (HP000025–34) (HP000025–34) at 20 – 21 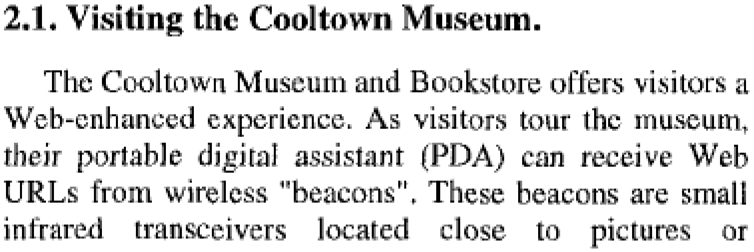 **2.1. Visiting the Cooltown Museum.** The Cooltown Museum and Bookstore offers visitors a Web-enhanced experience. As visitors tour the museum, their portable digital assistant (PDA) can receive Web URLs from wireless "beacons". These beacons are small infrared transceivers located close to pictures or

		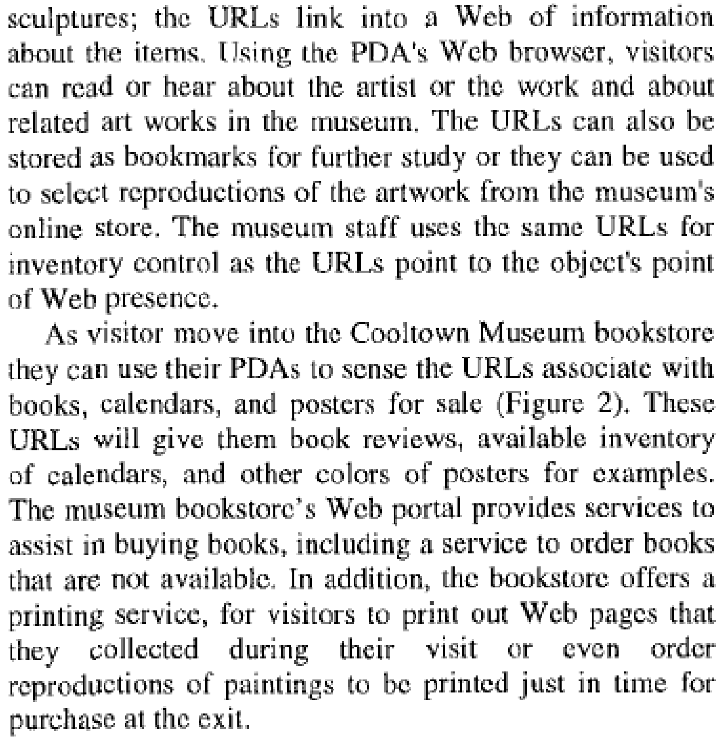 sculptures; the URLs link into a Web of information about the items. Using the PDA's Web browser, visitors can read or hear about the artist or the work and about related art works in the museum. The URLs can also be stored as bookmarks for further study or they can be used to select reproductions of the artwork from the museum's online store. The museum staff uses the same URLs for inventory control as the URLs point to the object's point of Web presence. As visitor move into the Cooltown Museum bookstore they can use their PDAs to sense the URLs associate with books, calendars, and posters for sale (Figure 2). These URLs will give them book reviews, available inventory of calendars, and other colors of posters for examples. The museum bookstore's Web portal provides services to assist in buying books, including a service to order books that are not available. In addition, the bookstore offers a printing service, for visitors to print out Web pages that they collected during their visit or even order reproductions of paintings to be printed just in time for purchase at the exit. Cooltown Developer's Network, People, Place, and Things: web presence for the real world (HP000025–34) (HP000025–34) at 23

		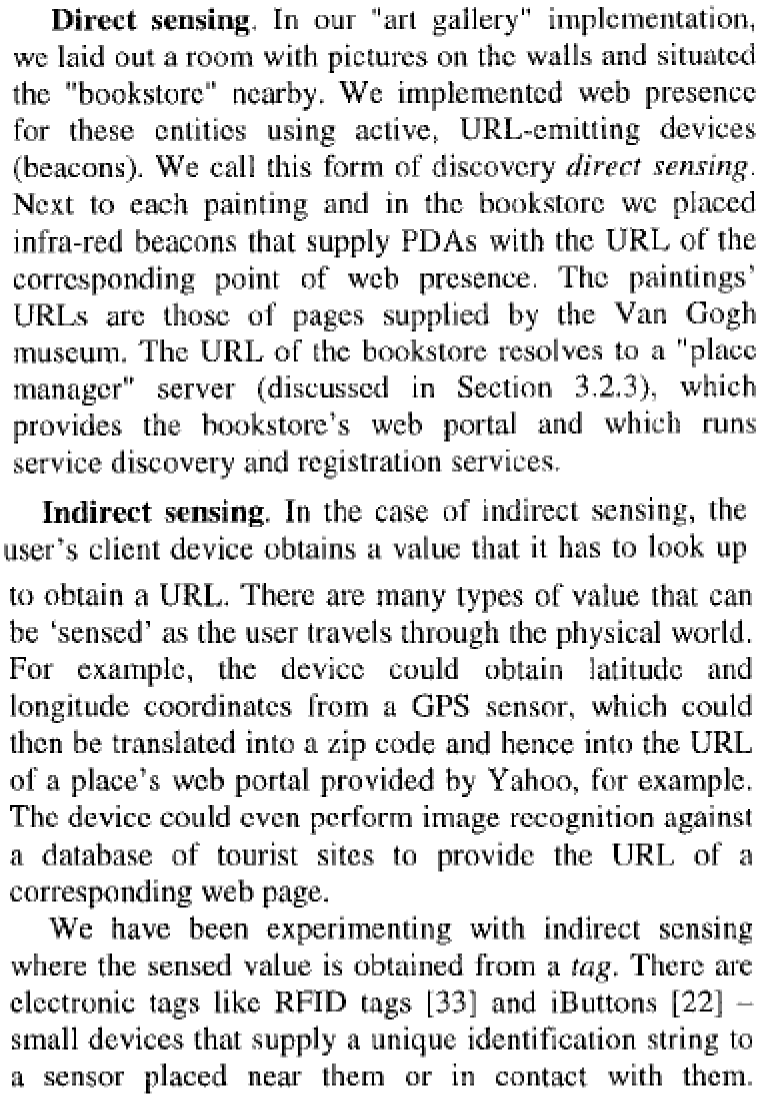 **Direct sensing**. In our "art gallery" implementation, we laid out a room with pictures on the walls and situated the "bookstore" nearby. We implemented web presence for these entities using active, URL-emitting devices (beacons). We call this form of discovery *direct sensing*. Next to each painting and in the bookstore we placed infra-red beacons that supply PDAs with the URL of the corresponding point of web presence. The paintings' URLs are those of pages supplied by the Van Gogh museum. The URL of the bookstore resolves to a "place manager" server (discussed in Section 3.2.3), which provides the bookstore's web portal and which runs service discovery and registration services. **Indirect sensing**. In the case of indirect sensing, the user's client device obtains a value that it has to look up to obtain a URL. There are many types of value that can be 'sensed' as the user travels through the physical world. For example, the device could obtain latitude and longitude coordinates from a GPS sensor, which could then be translated into a zip code and hence into the URL of a place's web portal provided by Yahoo, for example. The device could even perform image recognition against a database of tourist sites to provide the URL of a corresponding web page. We have been experimenting with indirect sensing where the sensed value is obtained from a *tag*. There are electronic tags like RFID tags [33] and iButtons [22] – small devices that supply a unique identification string to a sensor placed near them or in contact with them.

		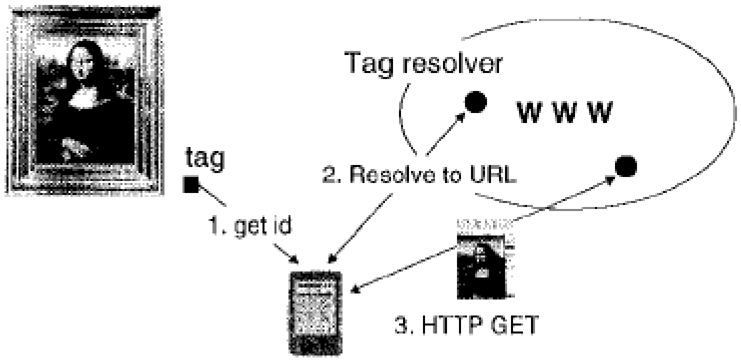 **Figure 4. Indirect sensing: the web presence for a painting.** Barcodes can be used instead of an electronic tag. Figure 4 shows a different art gallery demonstrator in which a PDA client senses a painting's tag. A correlated web page is resolved from the tag identifier and the client presents it to the user. We have produced a prototype implementation for sensing tagged things and displaying a corresponding context-dependent web page. The sensing device knows the URL of a *resolver*. When the client senses a tag, it sends the identifier to the resolver. If the resolver has an entry for the identifier, it sends back the corresponding URL. The user may choose a resolver that reflects his or her personal requirements, or those of a group of users with interests or tasks in common. For example, visitors to an art gallery could be offered resolvers that provide the URLs of pages in their native language, or pages provided by the art school at which they study. To the extent it is found that CoolTown does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that CoolTown does not

		anticipate this claim, CoolTown renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine CoolTown with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

3. Claim 27

Claim	'252 Claim Recitation	Exemplary Citations in Reference
27	The method of claim 26, wherein image processing further comprises compressing the image to form the data relating to the image.	CoolTown discloses the method of claim 26, wherein image processing further comprises compressing the image to form the data relating to the image. A Web-Based Nomadic Computing System (HP000074 – 87) at Fig. 3 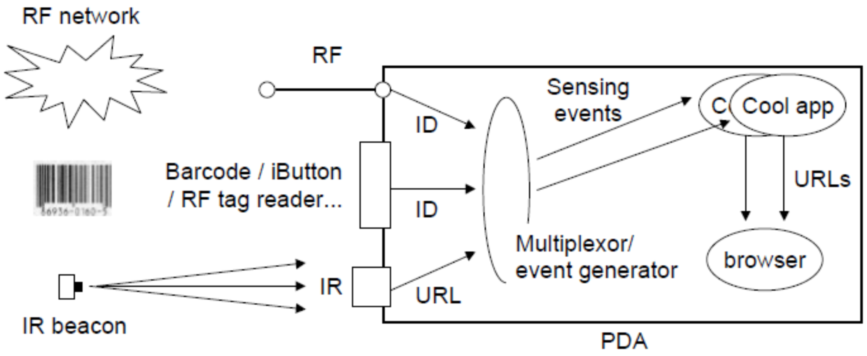 Figure 3. A sketch of a PDA equipped for sensing URLs and identifiers. A Web-Based Nomadic Computing System at 4

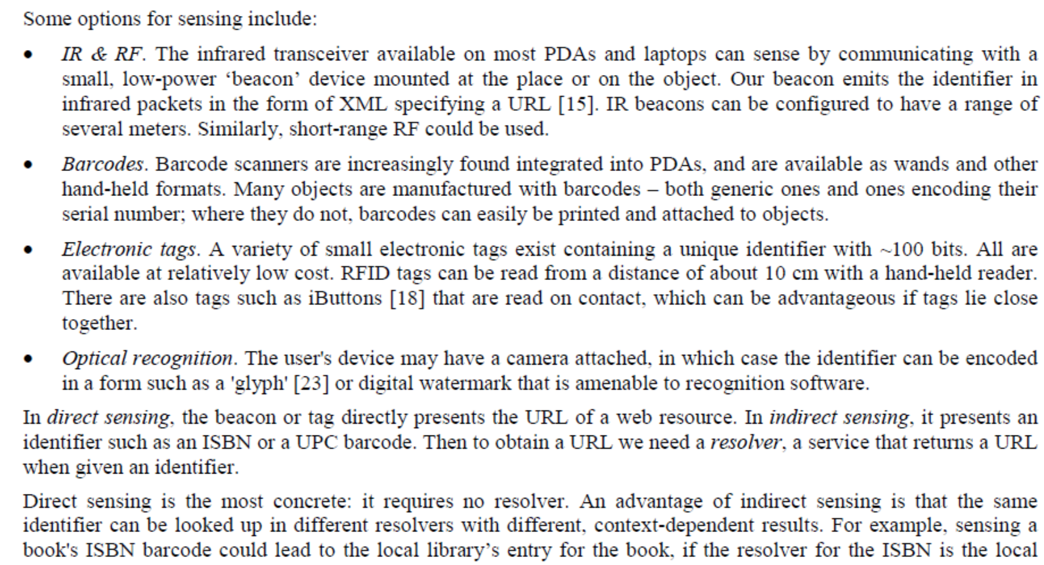

Some options for sensing include:

- *IR & RF*. The infrared transceiver available on most PDAs and laptops can sense by communicating with a small, low-power 'beacon' device mounted at the place or on the object. Our beacon emits the identifier in infrared packets in the form of XML specifying a URL [15]. IR beacons can be configured to have a range of several meters. Similarly, short-range RF could be used.
- *Barcodes*. Barcode scanners are increasingly found integrated into PDAs, and are available as wands and other hand-held formats. Many objects are manufactured with barcodes – both generic ones and ones encoding their serial number; where they do not, barcodes can easily be printed and attached to objects.
- *Electronic tags*. A variety of small electronic tags exist containing a unique identifier with ~100 bits. All are available at relatively low cost. RFID tags can be read from a distance of about 10 cm with a hand-held reader. There are also tags such as iButtons [18] that are read on contact, which can be advantageous if tags lie close together.
- *Optical recognition*. The user's device may have a camera attached, in which case the identifier can be encoded in a form such as a 'glyph' [23] or digital watermark that is amenable to recognition software.

In *direct sensing*, the beacon or tag directly presents the URL of a web resource. In *indirect sensing*, it presents an identifier such as an ISBN or a UPC barcode. Then to obtain a URL we need a *resolver*, a service that returns a URL when given an identifier.

Direct sensing is the most concrete: it requires no resolver. An advantage of indirect sensing is that the same identifier can be looked up in different resolvers with different, context-dependent results. For example, sensing a book's ISBN barcode could lead to the local library's entry for the book, if the resolver for the ISBN is the local

A Web-Based Nomadic Computing System (HP000074 – 87) at 8

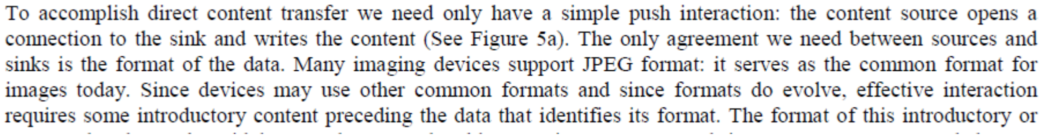

To accomplish direct content transfer we need only have a simple push interaction: the content source opens a connection to the sink and writes the content (See Figure 5a). The only agreement we need between sources and sinks is the format of the data. Many imaging devices support JPEG format: it serves as the common format for images today. Since devices may use other common formats and since formats do evolve, effective interaction requires some introductory content preceding the data that identifies its format. The format of this introductory or

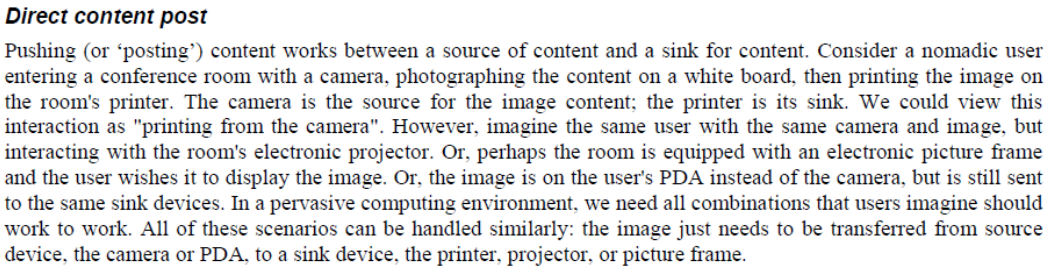

Direct content post

Pushing (or 'posting') content works between a source of content and a sink for content. Consider a nomadic user entering a conference room with a camera, photographing the content on a white board, then printing the image on the room's printer. The camera is the source for the image content; the printer is its sink. We could view this interaction as "printing from the camera". However, imagine the same user with the same camera and image, but interacting with the room's electronic projector. Or, perhaps the room is equipped with an electronic picture frame and the user wishes it to display the image. Or, the image is on the user's PDA instead of the camera, but is still sent to the same sink devices. In a pervasive computing environment, we need all combinations that users imagine should work to work. All of these scenarios can be handled similarly: the image just needs to be transferred from source device, the camera or PDA, to a sink device, the printer, projector, or picture frame.

To the extent it is found that CoolTown does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that CoolTown does not

		anticipate this claim, CoolTown renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine CoolTown with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

4. Claim 29

Claim	'252 Claim Recitation	Exemplary Citations in Reference
29	The method of claim 18, wherein image processing further includes converting at least a portion of the data relating to the image into a grayscale image.	To the extent it is found that CoolTown does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that CoolTown does not anticipate this claim, CoolTown renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine CoolTown with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

5. Claim 31

Claim	'252 Claim Recitation	Exemplary Citations in Reference
31	The method of claim 18, further comprising the distal server	CoolTown discloses the method of claim 18, further comprising the distal server using database to identify the pertinent information based on image characteristics

	using database to identify the pertinent information based on image characteristics of the image derived by the image processing platform based on the data relating to the image.	of the image derived by the image processing platform based on the data relating to the image. A Web-Based Nomadic Computing System (HP000074 – 87) at Fig. 2 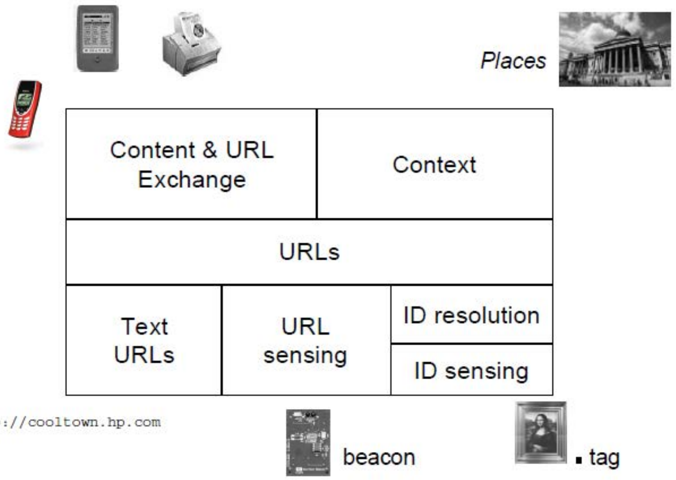 Figure 2. Component relationships. Sources of URLs, including keyboards, beacons, and resolved tags point to context or content. A Web-Based Nomadic Computing System (HP000074 – 87) at 4

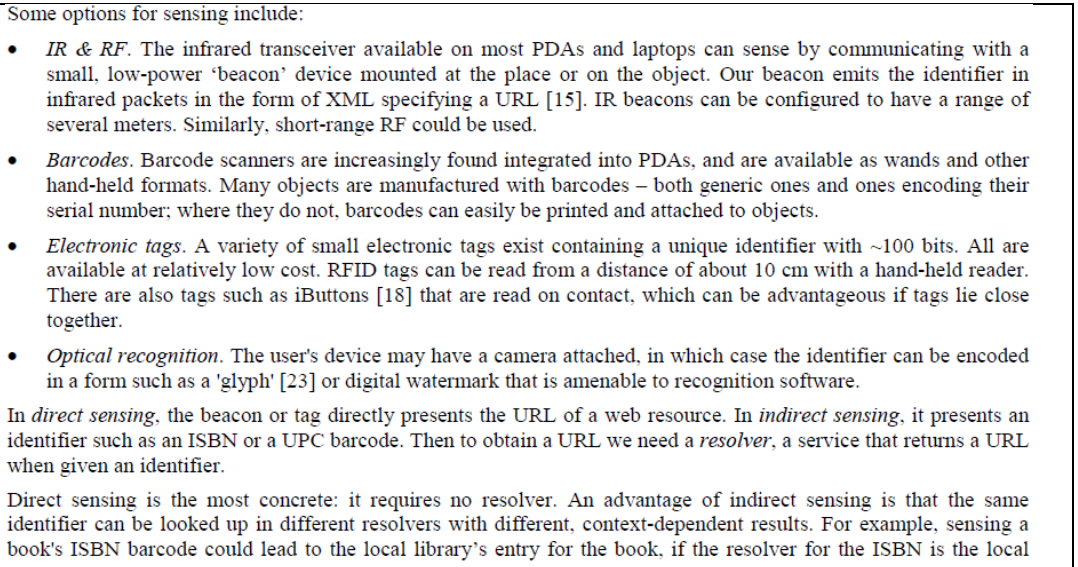

Some options for sensing include:

- *IR & RF*. The infrared transceiver available on most PDAs and laptops can sense by communicating with a small, low-power 'beacon' device mounted at the place or on the object. Our beacon emits the identifier in infrared packets in the form of XML specifying a URL [15]. IR beacons can be configured to have a range of several meters. Similarly, short-range RF could be used.
- *Barcodes*. Barcode scanners are increasingly found integrated into PDAs, and are available as wands and other hand-held formats. Many objects are manufactured with barcodes – both generic ones and ones encoding their serial number; where they do not, barcodes can easily be printed and attached to objects.
- *Electronic tags*. A variety of small electronic tags exist containing a unique identifier with ~100 bits. All are available at relatively low cost. RFID tags can be read from a distance of about 10 cm with a hand-held reader. There are also tags such as iButtons [18] that are read on contact, which can be advantageous if tags lie close together.
- *Optical recognition*. The user's device may have a camera attached, in which case the identifier can be encoded in a form such as a 'glyph' [23] or digital watermark that is amenable to recognition software.

In *direct sensing*, the beacon or tag directly presents the URL of a web resource. In *indirect sensing*, it presents an identifier such as an ISBN or a UPC barcode. Then to obtain a URL we need a *resolver*, a service that returns a URL when given an identifier.

Direct sensing is the most concrete: it requires no resolver. An advantage of indirect sensing is that the same identifier can be looked up in different resolvers with different, context-dependent results. For example, sensing a book's ISBN barcode could lead to the local library's entry for the book, if the resolver for the ISBN is the local

A Web-Based Nomadic Computing System (HP000074 – 87) at 5

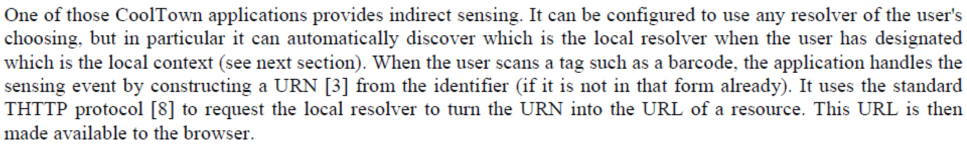

One of those CoolTown applications provides indirect sensing. It can be configured to use any resolver of the user's choosing, but in particular it can automatically discover which is the local resolver when the user has designated which is the local context (see next section). When the user scans a tag such as a barcode, the application handles the sensing event by constructing a URN [3] from the identifier (if it is not in that form already). It uses the standard THTTP protocol [8] to request the local resolver to turn the URN into the URL of a resource. This URL is then made available to the browser.

A Web-Based Nomadic Computing System (HP000074 – 87) at 6

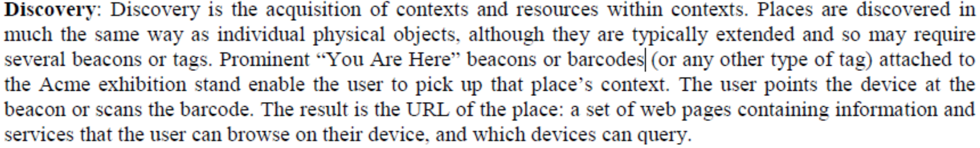

Discovery: Discovery is the acquisition of contexts and resources within contexts. Places are discovered in much the same way as individual physical objects, although they are typically extended and so may require several beacons or tags. Prominent "You Are Here" beacons or barcodes (or any other type of tag) attached to the Acme exhibition stand enable the user to pick up that place's context. The user points the device at the beacon or scans the barcode. The result is the URL of the place: a set of web pages containing information and services that the user can browse on their device, and which devices can query.

A Web-Based Nomadic Computing System (HP000074–87) at 6

		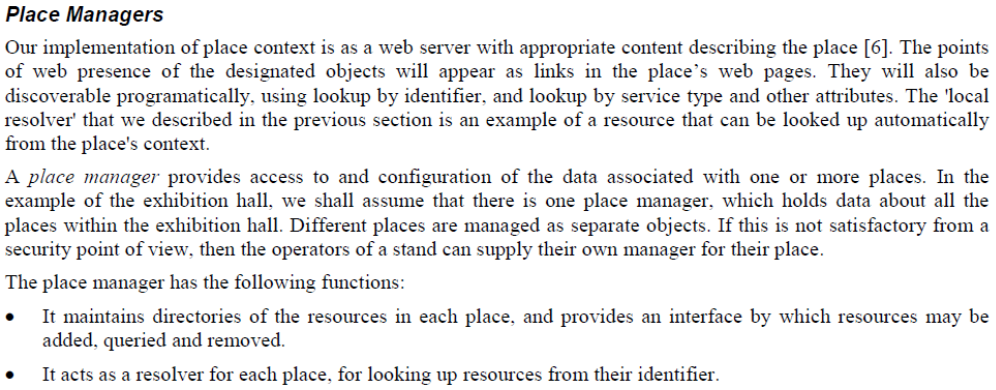 ***Place Managers*** Our implementation of place context is as a web server with appropriate content describing the place [6]. The points of web presence of the designated objects will appear as links in the place's web pages. They will also be discoverable programatically, using lookup by identifier, and lookup by service type and other attributes. The 'local resolver' that we described in the previous section is an example of a resource that can be looked up automatically from the place's context. A *place manager* provides access to and configuration of the data associated with one or more places. In the example of the exhibition hall, we shall assume that there is one place manager, which holds data about all the places within the exhibition hall. Different places are managed as separate objects. If this is not satisfactory from a security point of view, then the operators of a stand can supply their own manager for their place. The place manager has the following functions: • It maintains directories of the resources in each place, and provides an interface by which resources may be added, queried and removed. • It acts as a resolver for each place, for looking up resources from their identifier. Cooltown Developer's Network, People, Place, and Things: web presence for the real world (HP000025–34) at 8 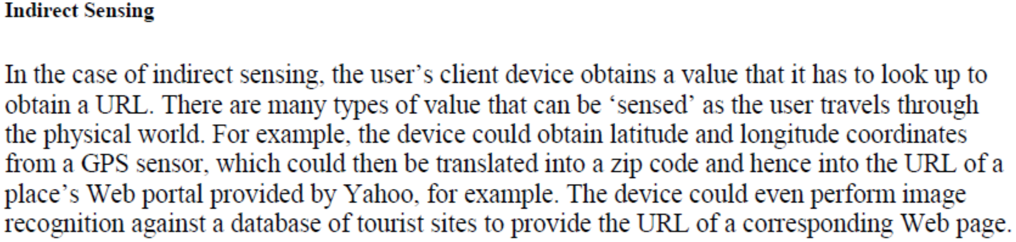 **Indirect Sensing** In the case of indirect sensing, the user's client device obtains a value that it has to look up to obtain a URL. There are many types of value that can be 'sensed' as the user travels through the physical world. For example, the device could obtain latitude and longitude coordinates from a GPS sensor, which could then be translated into a zip code and hence into the URL of a place's Web portal provided by Yahoo, for example. The device could even perform image recognition against a database of tourist sites to provide the URL of a corresponding Web page. Cooltown Developer's Network, People, Place, and Things: web presence for the real world (HP000025–34) (HP000025–34) at 23

		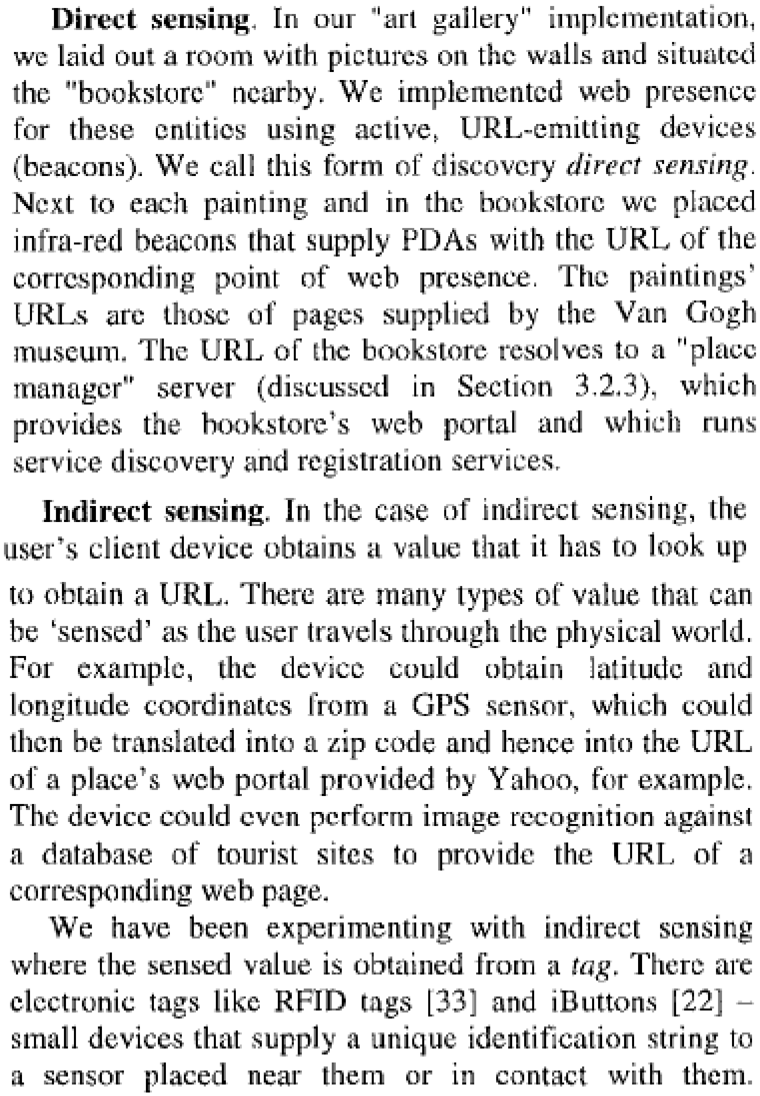 **Direct sensing.** In our "art gallery" implementation, we laid out a room with pictures on the walls and situated the "bookstore" nearby. We implemented web presence for these entities using active, URL-emitting devices (beacons). We call this form of discovery *direct sensing*. Next to each painting and in the bookstore we placed infra-red beacons that supply PDAs with the URL of the corresponding point of web presence. The paintings' URLs are those of pages supplied by the Van Gogh museum. The URL of the bookstore resolves to a "place manager" server (discussed in Section 3.2.3), which provides the bookstore's web portal and which runs service discovery and registration services. **Indirect sensing.** In the case of indirect sensing, the user's client device obtains a value that it has to look up to obtain a URL. There are many types of value that can be 'sensed' as the user travels through the physical world. For example, the device could obtain latitude and longitude coordinates from a GPS sensor, which could then be translated into a zip code and hence into the URL of a place's web portal provided by Yahoo, for example. The device could even perform image recognition against a database of tourist sites to provide the URL of a corresponding web page. We have been experimenting with indirect sensing where the sensed value is obtained from a *tag*. There are electronic tags like RFID tags [33] and iButtons [22] – small devices that supply a unique identification string to a sensor placed near them or in contact with them.

		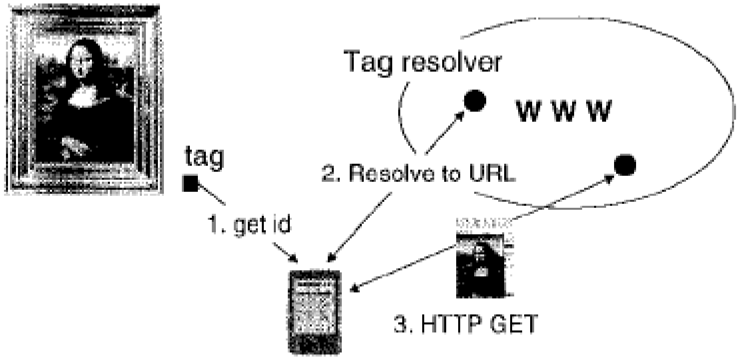 **Figure 4. Indirect sensing: the web presence for a painting.** Barcodes can be used instead of an electronic tag. Figure 4 shows a different art gallery demonstrator in which a PDA client senses a painting's tag. A correlated web page is resolved from the tag identifier and the client presents it to the user. We have produced a prototype implementation for sensing tagged things and displaying a corresponding context-dependent web page. The sensing device knows the URL of a *resolver*. When the client senses a tag, it sends the identifier to the resolver. If the resolver has an entry for the identifier, it sends back the corresponding URL. The user may choose a resolver that reflects his or her personal requirements, or those of a group of users with interests or tasks in common. For example, visitors to an art gallery could be offered resolvers that provide the URLs of pages in their native language, or pages provided by the art school at which they study. To the extent it is found that CoolTown does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that CoolTown does not

		anticipate this claim, CoolTown renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine CoolTown with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

Nantworks, LLC v. Bank of America Corporation
Central District of California, No. 2:20-cv-7872
Invalidity Chart for U.S. Patent 9,031,278 based on
U.S. Patent No. 6,512,919 to Ogasawara ("Ogasawara")

Exhibit F-1

Bank of America Corporation ("BofA" or "Defendant") provides this chart subject to all reservations, objections, statements, and disclaimers set forth herein and in BofA's Final Invalidity Contentions Cover Pleading, as well as any amendment, supplement, or modification thereof, which are incorporated herein by reference in their entirety.

On information and belief, the application that matured into Ogasawara was filed on March 30, 1999, claiming the priority date of December 14, 1998. Ogasawara published at least as of May 30, 2002 and is available as prior art at least under 35 U.S.C. § 102(e).

As shown in the chart below, CoolTown anticipates asserted claims 1, 4, and 5 of the '278 Patent. To the extent CoolTown is found not to expressly disclose certain limitations in the asserted claims, such limitations are inherent. To the extent CoolTown is found not to anticipate any asserted claims or claim elements of the '278 Patent, the reference nevertheless renders those claims or claim elements obvious under 35 U.S.C. § 103, either alone or in combination with other art identified herein or in the cover pleading of BofA's Invalidity Contentions, for the reasons identified in the cover pleading. These Final Invalidity Contentions are not an admission by BofA that the accused products, including any current or past versions of the accused products, are covered by, or infringe the asserted claims, but are based instead on the recognition that if the claims are interpreted to be broad enough to encompass the accused products, the claims must also be construed to have that same scope when considering whether they are invalid.

The following chart is based on the Court's claim construction rulings (Dkt. Nos. 145, 153, 236), and to the extent not contradictory, Plaintiffs' apparent interpretations of the claims in Plaintiffs' "Preliminary" Final Infringement Contentions. BofA notes that in many instances, Plaintiffs' "Preliminary" Final Infringement Contentions ignore and/or contradict the Court's claim construction rulings. BofA does not accept the assumptions concerning the scope and meaning implicit in Plaintiffs' Preliminary Final Infringement Contentions, to the extent those assumptions are discernible, and reserves the right to challenge Plaintiffs' proposed (or implied) constructions. Furthermore, Plaintiffs' "Preliminary" Final Contentions fail to comply with the level of specificity required by SPR 2.1.2 and 2.1.3, and include improper reservation of rights to supplement with "additional ways in which BoA Accused Products infringe." Any alleged deficient discovery is of Plaintiffs own making. Most if not all of the deficiencies in Plaintiffs' Preliminary Final Contentions have been documented by Defendants through multiple letters since at least May 18, 2021 and up to Bank of America's most recent letter of August 19, 2023. Plaintiffs' shifting infringement contentions, deficiencies in its infringement contentions, and changing claim scope have prejudiced BofA from being able to reasonably prepare its defenses, including its invalidity contentions. BofA also reserves the right to revise and supplement these charts if and when Plaintiffs are permitted to provide revised or supplemental Infringement Contentions. Citations given in the chart below are merely representative of the respective elements to provide notice to Plaintiffs of BofA's invalidity theories and are not meant to be exhaustive.

1. Claim 1

Element	'278 Claim Recitation	Exemplary Citations in Reference
[1Pre]	A computer-assisted method, comprising:	Ogasawara discloses a computer-assisted method. Ogasawara at Abstract • "An electronic shopping system facilitates purchase transactions via a wireless videophone. A purchase transaction program is downloaded from the seller's server to a purchaser's wireless videophone via a program loader contained within the purchaser's wireless videophone. The purchase transaction program is stored in a program memory and is used by the purchaser to facilitate the selection of items to be purchased, as well as payment therefor." Ogasawara at 1:16–21 • "The present invention relates generally to electronic shopping systems and, more particularly to an electronic shopping system which utilizes a program downloadable wireless video phone as a purpose-type dedicated terminal which enables a shopper to capture, recognize and decode captured images."
[1A]	receiving, via a mobile device, an image comprising a representation of at least a portion of a document;	Ogasawara discloses receiving, via a mobile device, an image comprising a representation of at least a portion of a document. *See* Ogasawara at Figs. 1 and 11

Element	'278 Claim Recitation	Exemplary Citations in Reference
		FIG. 1 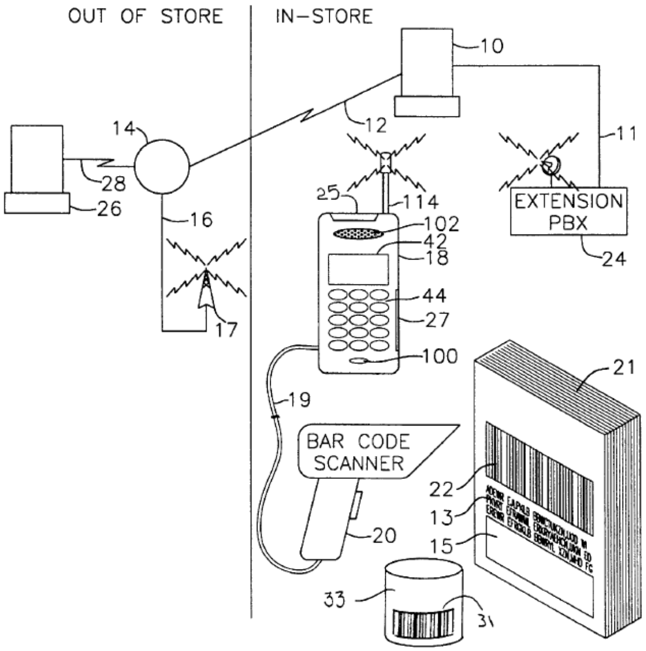

Element	'278 Claim Recitation	Exemplary Citations in Reference
		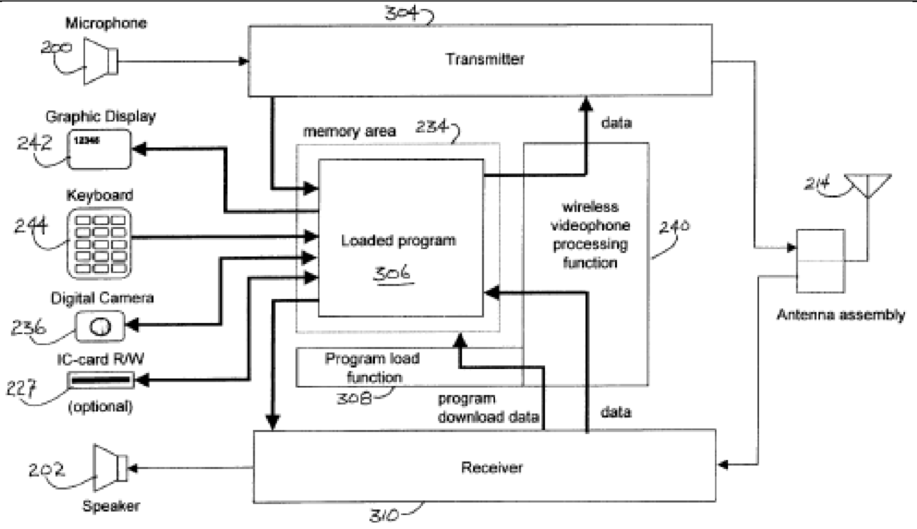 *FIG. 11* Ogasawara at 5:40–47 • "A catalog 21 of the items which can be purchased contains a bar code 22 for each such item, and preferably also contains descriptive text 13 and a picture 15 of each item. The use of such a catalog 21 or the like facilitates the purchasing of products via the electronic shopping system of the present invention when the purchaser is not in the store where the items are sold. Typically, each item 33 also has a bar code 31 applied thereto." Ogasawara at 7:26–33 • "When the wireless telephone 18 is used to make purchases within a store, bar codes on merchandise or bar codes on the store shelf where the product is displayed, upon the product to be purchased or within a catalog, may be scanned to facilitate selection of desired items to be purchased. When the

Element	'278 Claim Recitation	Exemplary Citations in Reference
		wireless telephone 18 is used to make purchases while away from the store, then a catalog 21 or any other source of bar codes may be utilized." Ogasawara at 16:7–13 • "In addition, the digital camera **236** includes processing circuitry which translates the visual image acquired by the camera's lens into digital signals suitable for processing by the microprocessor **238** into a form which can be broadcast transmitted (i.e., JPEG encoded, gif encoded, and the like) by the functional electronic section **240**." Ogasawara at 19:46–65 • "In accordance with practice of principles of the invention, videographic image data is received by the system's digital camera **236** and provided to the system's microprocessor **238** where the received videographic information is processed by an application program previously downloaded to the wireless videophone and stored as a loaded program 306. Following processing, the videographic image data is encoded by the channel encoded circuitry 324 for modulation and transmission to a store server. Optionally, if character or pattern recognition software has not been downloaded into the system, videographic image data is received by the microprocessor **238** and passed directly to the channel encoder circuitry **324** for modulation and broadcast by he [sic] system. During receipt of videographic image data, once the digital videographic data has been demodulated and equalized by the demodulator and equalizer circuitry **348** and 346 respectively, the videographic image information is decoded by channel decoder circuitry **344** and directed to the microprocessor **238** for display on the graphic display screen **242**." Ogasawara at 21:54 – 65

Element	'278 Claim Recitation	Exemplary Citations in Reference
		• "In addition to its use as a simple, efficient electronic shopping system, the wireless videophone of the present invention is not limited to use in retail transactions. Because its program download capability is not application specific, it will be understood that any application program of an appropriate size is able to be downloaded and executed by the wireless videophone of the invention. The program loader capability in combination with the system's digital video camera and graphics display screen allows the system to be eminently suitable for catalog shopping application, in which items are selected for purchase by merely scanning a printed bar code in proximity to an item on the catalog page." Ogasawara at claim 24 • "24. The program downloadable wireless videophone according to claim 23, wherein a bar code identifies particular ones of a plurality of items, the items defining an application selected from the group consisting of an electronic shopping system merchandise ordering display, a bar code telephone directory, a printed bar code identified merchandise catalog, and a product affixed identification label." To the extent it is found that Ogasawara does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Ogasawara does not anticipate this claim, Ogasawara renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Ogasawara with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to

Element	'278 Claim Recitation	Exemplary Citations in Reference
		address similar known problems, according to known methods which would yield predictable results.
[1B]	determining that symbolic content is on the at least the portion of the document based on the image;	Ogasawara discloses determining that symbolic content is on the at least the portion of the document based on the image. Ogasawara at 3:13–17 • "An integral digital camera is used to scan the images of bar codes of purchased items, and pattern recognition software resident either in the videophone or in the server, translates the bar code image data into an alpha-numeric product identification." Ogasawara at 18:11–22 • "For those purchasers who are equipped with a wireless videophone in accordance with the present invention, a tailored purchase transaction program is provided which accommodates the enhanced functional capabilities of the device. In particular, since the wireless videophone **218** is provided with a digital camera **236** in place of a bar code scanner, the tailored purchase transaction program might additionally include character recognition and/or pattern recognition, as well as bar code decode, software which would allow the wireless videophone to function in a manner similar to the wireless telephone and bar code scanner embodiment described above." Ogasawara at 6:46–50 • "When the store server **10** or remote server **26** receives bar code data from the customer's wireless telephone **18**, then the store server **10** or remote server **26** searches a database and obtains a description and price for the item scanned." Ogasawara at 18:23–58

Element	'278 Claim Recitation	Exemplary Citations in Reference
		• "It should also be understood that the character recognition, pattern recognition and/or bar code decode routines need not be directly downloaded by the program loader firmware memory of the wireless videophone in order that the system have this particular capability. The wireless videophone is perfectly capable of capturing digital videographic information, such as a bar code pattern or a graphics image pattern, and transferring this information directly to the store's server through the functional electronics section. Once this information is received by the server, a character recognition and/or pattern recognition application routine can be invoked within the server in order to either decode the bar code videographic data or to perform pattern recognition functions on a icon-like patter captured by the digital video camera. A server personal shopping application 254 is also resident in the server system 210 and might be characterized as a program for facilitating operation of the server in order to perform electronic shopping. The server personal shopping application 254 facilitates the downloading of the purchase transaction program and/or any appropriate character or pattern recognition routines to the wireless videophone 218 after the wireless videophone has established a connection with the server as discussed above. Server personal shopping application 254 also facilitates the reception and processing-of product selections made by a purchaser utilizing the wireless videophone as discussed previously as well as optionally facilitating the reception and decoding of bar codes, character or pattern information by an optional character and/or pattern recognition software routine. The server personal shopping application 254 is also adapted to receive and store payment information, such as credit card account numbers, expiration dates, and the like and is also capable of facilitating reading and updating of information on a purchaser's IC card via the IC card reader/writer 227, if utilized."

Element	'278 Claim Recitation	Exemplary Citations in Reference
		Ogasawara at 20:41–67 • "With particular reference to FIG. 13, once the purchase transaction program has been downloaded, a purchaser may use the wireless videophone, in particular the digital camera, to select various items which the consumer desires to purchase. Such selection is made by capturing an image of an item's UPC bar code data, or the like, within the image field of the digital camera. While the consumer is scanning a particular bar code with the digital camera, the image seen by the camera is displayed on the system's graphic display screen so that the consumer can visually verify that the entire image has been captured. If character and/or pattern recognition software was downloaded to the system as part of the program, the customer need only inform the system that the image is ready for processing by pressing a pre-defined key on the system's keyboard. Those skilled in the art will appreciate that various other codes, indicia, text, icons, and the like may be scanned and recognized by modern character and pattern recognition application routines. Thus, product selection need not be made in accordance with a UPC bar code but might be made with reference to a selection of product specific icons, each of which represent a particular product or service offered for purchase. For example, when grocery shopping, product specific icons might represent the stylistic outlines of a loaf of bread, carton of milk, a bunch of broccoli, and the like, allowing a consumer to quickly and efficiently select items for purchase without recourse to the complications of reading bar codes." Ogasawara at 21:15–24. • "In the case of FIG. 13, if the bar code read key has been depressed the program captures the bar code image taken by the digital camera. Once the bar code image has been captured, the program decodes the bar code image data to its corresponding numeric bar code data, by operating on the bar code image with pattern recognition software. Once the bar code has been

Element	'278 Claim Recitation	Exemplary Citations in Reference
		decoded to its numeric values, the bar code data is transmitted to the store's server through the wireless telephone function of the system." Ogasawara at 21:54 – 65 • "In addition to its use as a simple, efficient electronic shopping system, the wireless videophone of the present invention is not limited to use in retail transactions. Because its program download capability is not application specific, it will be understood that any application program of an appropriate size is able to be downloaded and executed by the wireless videophone of the invention. The program loader capability in combination with the system's digital video camera and graphics display screen allows the system to be eminently suitable for catalog shopping application, in which items are selected for purchase by merely scanning a printed bar code in proximity to an item on the catalog page." Ogasawara at 22:41 – 23:10 • "Notwithstanding the foregoing, it will also be understood that the merchandise item information need not be displayed in bar code form for the invention to perform as intended. Indeed, merchandise item information may be presented in alpha-numeric form (i.e., an item name) or as a numeric code. In either case, character recognition software allows the wireless videophone to identify the item in a seamless fashion. Character recognition software supports identifying a product based on a numeric or alpha-numeric code, while pattern recognition software supports recognition of product characteristics based on bar code data. As was described above, pattern and/or character recognition can be accomplished by software downloaded to the wireless videophone, or, alternatively, by transmitting the videographic image of the numeric, alpha-numeric or bar code identifier to a server for such recognition processing.

Element	'278 Claim Recitation	Exemplary Citations in Reference
		Processing results are then transmitted back to the wireless videophone for display to the customer. It will be understood that the system of the present invention allows bar code and/or alpha-numeric information to be captured from almost any form of display means by merely placing the information within the visual image field of the system's digital camera. If the system's downloaded program includes pattern and/or character recognition routines, recognition data processing is performed within the video phone. Where the downloaded program does not include recognition software, the program transmits the videographic image to a server and a recognition processing software resident in the server processes the data and returns recognized merchandise information to the graphics display screen of the videophone. In its simplest form, the downloaded program might transmit videographic image data to a store server which, in turn, displays the information to a store clerk. The store clerk would then manually identify the merchandise associated with the transmitted data (bar code and/or alpha-numeric) and return identified merchandise information to the consumer's videophone graphics display screen." To the extent it is found that Ogasawara does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Ogasawara does not anticipate this claim, Ogasawara renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Ogasawara with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

Element	**'278 Claim Recitation**	**Exemplary Citations in Reference**
[1C]	extracting symbol information based on the symbolic content according to symbol type;	Ogasawara discloses extracting symbol information based on the symbolic content according to symbol type. *See* [1B] above. *See also:* Ogasawara at Abstract • "An electronic shopping system facilitates purchase transactions via a wireless videophone. A purchase transaction program is downloaded from the seller's server to a purchaser's wireless videophone via a program loader contained within the purchaser's wireless videophone. The purchase transaction program is stored in a program memory and is used by the purchaser to facilitate the selection of items to be purchased, as well as payment therefor. An integral digital camera is attached to the wireless telephone to facilitate the selection of items to be purchased and is controlled via the downloaded purchase transaction program to function as a bar code or product icon image capture device. Character or pattern recognition software translates the bar or icon code image into an appropriate item identifier." Ogasawara at 22:46–51. • "In either case, character recognition software allows the wireless videophone to identify the item in a seamless fashion. Character recognition software supports identifying a product based on a numeric or alpha-numeric code, while pattern recognition software supports recognition of product characteristics based on bar code data." Ogasawara at claims 13 – 16 • "13. The electronic shopping system according to claim 12, wherein the item indicia is a bar code, the pattern recognition software translating the bar code data into a numerical product identification.

Element	**'278 Claim Recitation**	**Exemplary Citations in Reference**
		14. The electronic shopping system according to claim 12, wherein the item indicia is a product specific pattern shape, the pattern recognition software translating the pattern shape data into a numerical product identification. 15. The electronic shopping system according to claim 5, the server configured to host a pattern recognition software program receiving the visual image data associated with the item indicia, the pattern recognition software program further recognizing the item indicia and translating the item indicia into a corresponding desired product identification. 16. The electronic shopping system according to claim 15, wherein the item indicia is a bar code, the pattern recognition program translating the bar code data into a numerical product identification. To the extent it is found that Ogasawara does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Ogasawara does not anticipate this claim, Ogasawara renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Ogasawara with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1D]	determining a validity of the document based at least in part on the image and the symbol information; and	Ogasawara discloses determining a validity of the document based at least in part on the image and the symbol information. Ogasawara at 3:13–17

Element	'278 Claim Recitation	Exemplary Citations in Reference
		• "An integral digital camera is used to scan the images of bar codes of purchased items, and pattern recognition software resident either in the videophone or in the server, translates the bar code image data into an alpha-numeric product identification." Ogasawara at 5:40–47 • "A catalog 21 of the items which can be purchased contains a bar code 22 for each such item, and preferably also contains descriptive text 13 and a picture 15 of each item. The use of such a catalog 21 or the like facilitates the purchasing of products via the electronic shopping system of the present invention when the purchaser is not in the store where the items are sold. Typically, each item 33 also has a bar code 31 applied thereto." Ogasawara at 6:46–50 • "When the store server **10** or remote server **26** receives bar code data from the customer's wireless telephone **18**, then the store server **10** or remote server **26** searches a database and obtains a description and price for the item scanned." Ogasawara at 7:26–33 • "When the wireless telephone 18 is used to make purchases within a store, bar codes on merchandise or bar codes on the store shelf where the product is displayed, upon the product to be purchased or within a catalog, may be scanned to facilitate selection of desired items to be purchased. When the wireless telephone 18 is used to make purchases while away from the store, then a catalog 21 or any other source of bar codes may be utilized." Ogasawara at 18:11–22 • "For those purchasers who are equipped with a wireless videophone in accordance with the present invention, a tailored purchase transaction

Element	'278 Claim Recitation	Exemplary Citations in Reference
		program is provided which accommodates the enhanced functional capabilities of the device. In particular, since the wireless videophone **218** is provided with a digital camera **236** in place of a bar code scanner, the tailored purchase transaction program might additionally include character recognition and/or pattern recognition, as well as bar code decode, software which would allow the wireless videophone to function in a manner similar to the wireless telephone and bar code scanner embodiment described above." Ogasawara at 18:23–58 • "It should also be understood that the character recognition, pattern recognition and/or bar code decode routines need not be directly downloaded by the program loader firmware memory of the wireless videophone in order that the system have this particular capability. The wireless videophone is perfectly capable of capturing digital videographic information, such as a bar code pattern or a graphics image pattern, and transferring this information directly to the store's server through the functional electronics section. Once this information is received by the server, a character recognition and/or pattern recognition application routine can be invoked within the server in order to either decode the bar code videographic data or to perform pattern recognition functions on a icon-like patter captured by the digital video camera. A server personal shopping application 254 is also resident in the server system 210 and might be characterized as a program for facilitating operation of the server in order to perform electronic shopping. The server personal shopping application 254 facilitates the downloading of the purchase transaction program and/or any appropriate character or pattern recognition routines to the wireless videophone 218 after the wireless videophone has established a connection with the server as discussed above. Server personal shopping application 254 also facilitates the reception and

Element	'278 Claim Recitation	Exemplary Citations in Reference
		processing-of product selections made by a purchaser utilizing the wireless videophone as discussed previously as well as optionally facilitating the reception and decoding of bar codes, character or pattern information by an optional character and/or pattern recognition software routine. The server personal shopping application 254 is also adapted to receive and store payment information, such as credit card account numbers, expiration dates, and the like and is also capable of facilitating reading and updating of information on a purchaser's IC card via the IC card reader/writer 227, if utilized." Ogasawara at 20:41–21:8 • "With particular reference to FIG. 13, once the purchase transaction program has been downloaded, a purchaser may use the wireless videophone, in particular the digital camera, to select various items which the consumer desires to purchase. Such selection is made by capturing an image of an item's UPC bar code data, or the like, within the image field of the digital camera. . . . If character and/or pattern recognition software was downloaded to the system as part of the program, the customer need only inform the system that the image is ready for processing by pressing a pre-defined key on the system's keyboard. Those skilled in the art will appreciate that various other codes, indicia, text, icons, and the like may be scanned and recognized by modern character and pattern recognition application routines. Thus, product selection need not be made in accordance with a UPC bar code but might be made with reference to a selection of product specific icons, each of which represent a particular product or service offered for purchase. For example, when grocery shopping, product specific icons might represent the stylistic outlines of a loaf of bread, carton of milk, a bunch of broccoli, and the like, allowing a consumer to quickly and efficiently select items for purchase without recourse to the complications of reading bar codes.

Element	'278 Claim Recitation	Exemplary Citations in Reference
		However, no matter the form of the indicia captured by the digital video camera, the system operates in identical manner for each indicia and is governed solely by operation of the downloaded program which is coded to recognize either characters or patterns as appropriate to the application. However, in the following, the types of videographic images captured by the video camera will all be subsumed under the generic term 'bar code' image data." Ogasawara at 21:24–35 • "Once in the server, the numeric bar code data is associated to a corresponding product which can be retrieved from a Price Look Up (PLU) table or file. An "item selection complete" response is then issued by the server and received by the customer's wireless videophone through its telephone function and the response is displayed to the customer on the system's integral graphics display screen. The response might include the numeric bar code data of the items selected for purchase as well as the item name, item price, and the like, in order that the consumer may confirm that this is, indeed, the item that they selected." Ogasawara at 21:54 – 65 • "In addition to its use as a simple, efficient electronic shopping system, the wireless videophone of the present invention is not limited to use in retail transactions. Because its program download capability is not application specific, it will be understood that any application program of an appropriate size is able to be downloaded and executed by the wireless videophone of the invention. The program loader capability in combination with the system's digital video camera and graphics display screen allows the system to be eminently suitable for catalog shopping application, in which items are selected for purchase by merely scanning a printed bar code in proximity to an item on the catalog page."

Element	'278 Claim Recitation	Exemplary Citations in Reference
		Ogasawara at claim 24 • "24. The program downloadable wireless videophone according to claim 23, wherein a bar code identifies particular ones of a plurality of items, the items defining an application selected from the group consisting of an electronic shopping system merchandise ordering display, a bar code telephone directory, a printed bar code identified merchandise catalog, and a product affixed identification label." To the extent it is found that Ogasawara does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Ogasawara does not anticipate this claim, Ogasawara renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Ogasawara with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1E]	recognizing the document as a first target object based at least in part on	Ogasawara discloses this limitation. Ogasawara at 6:46–50 • "When the store server **10** or remote server **26** receives bar code data from the customer's wireless telephone **18**, then the store server **10** or remote server **26** searches a database and obtains a description and price for the item scanned." Ogasawara at 18:11–22

Element	'278 Claim Recitation	Exemplary Citations in Reference
		- "For those purchasers who are equipped with a wireless videophone in accordance with the present invention, a tailored purchase transaction program is provided which accommodates the enhanced functional capabilities of the device. In particular, since the wireless videophone **218** is provided with a digital camera **236** in place of a bar code scanner, the tailored purchase transaction program might additionally include character recognition and/or pattern recognition, as well as bar code decode, software which would allow the wireless videophone to function in a manner similar to the wireless telephone and bar code scanner embodiment described above." Ogasawara at 18:23–58 - "It should also be understood that the character recognition, pattern recognition and/or bar code decode routines need not be directly downloaded by the program loader firmware memory of the wireless videophone in order that the system have this particular capability. The wireless videophone is perfectly capable of capturing digital videographic information, such as a bar code pattern or a graphics image pattern, and transferring this information directly to the store's server through the functional electronics section. Once this information is received by the server, a character recognition and/or pattern recognition application routine can be invoked within the server in order to either decode the bar code videographic data or to perform pattern recognition functions on a icon-like patter captured by the digital video camera. A server personal shopping application 254 is also resident in the server system 210 and might be characterized as a program for facilitating operation of the server in order to perform electronic shopping. The server personal shopping application 254 facilitates the downloading of the purchase transaction program and/or any appropriate character or pattern

Element	'278 Claim Recitation	Exemplary Citations in Reference
		recognition routines to the wireless videophone 218 after the wireless videophone has established a connection with the server as discussed above. Server personal shopping application 254 also facilitates the reception and processing-of product selections made by a purchaser utilizing the wireless videophone as discussed previously as well as optionally facilitating the reception and decoding of bar codes, character or pattern information by an optional character and/or pattern recognition software routine. The server personal shopping application 254 is also adapted to receive and store payment information, such as credit card account numbers, expiration dates, and the like and is also capable of facilitating reading and updating of information on a purchaser's IC card via the IC card reader/writer 227, if utilized." Ogasawara at 20:41–21:8 • "With particular reference to FIG. 13, once the purchase transaction program has been downloaded, a purchaser may use the wireless videophone, in particular the digital camera, to select various items which the consumer desires to purchase. Such selection is made by capturing an image of an item's UPC bar code data, or the like, within the image field of the digital camera. . . . If character and/or pattern recognition software was downloaded to the system as part of the program, the customer need only inform the system that the image is ready for processing by pressing a pre-defined key on the system's keyboard. Those skilled in the art will appreciate that various other codes, indicia, text, icons, and the like may be scanned and recognized by modern character and pattern recognition application routines. Thus, product selection need not be made in accordance with a UPC bar code but might be made with reference to a selection of product specific icons, each of which represent a particular product or service offered for purchase. For example, when grocery shopping, product specific icons might represent the stylistic outlines of a loaf of bread, carton of milk, a bunch of broccoli, and the like, allowing a

Element	'278 Claim Recitation	Exemplary Citations in Reference
		consumer to quickly and efficiently select items for purchase without recourse to the complications of reading bar codes. However, no matter the form of the indicia captured by the digital video camera, the system operates in identical manner for each indicia and is governed solely by operation of the downloaded program which is coded to recognize either characters or patterns as appropriate to the application. However, in the following, the types of videographic images captured by the video camera will all be subsumed under the generic term 'bar code' image data." Ogasawara at 22:40–51 • "Notwithstanding the foregoing, it will also be understood that the merchandise item information need not be displayed in bar code form for the invention to perform as intended. Indeed, merchandise item information may be presented in alpha-numeric form (i.e., an item name) or as a numeric code. In either case, character recognition software allows the wireless videophone to identify the item in a seamless fashion. Character recognition software supports identifying a product based on a numeric or alpha-numeric code, while pattern recognition software supports recognition of product characteristics based on bar code data." Ogasawara at 23:11–31 • "Advanced pattern recognition software is able to enhance the performance of the wireless videophone, in accord with the invention, by offering the capability to capture merchandise information from items that are not identified by either a bar code or an alpha-numeric label. Advanced pattern recognition software allows a consumer to capture a videographic image of an apple, for example, and to have the apple be recognized as such by the pattern recognition software. This capability is useful for any merchandise

Element	'278 Claim Recitation	Exemplary Citations in Reference
		item having a distinct or identifiable shape or other visually identifiable characteristic. Such pattern recognition is accomplished within the wireless videophone, by specialized downloaded software. Alternatively, a visual image of the merchandise item is transmitted to a server, wherein pattern recognition software makes the requisite merchandise item identification." To the extent it is found that Ogasawara does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Ogasawara does not anticipate this claim, Ogasawara renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Ogasawara with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1E.i]	the image,	Ogasawara discloses this limitation. *See* limitation [1E] above. To the extent it is found that Ogasawara does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Ogasawara does not anticipate this claim, Ogasawara renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Ogasawara with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these

Element	'278 Claim Recitation	Exemplary Citations in Reference
		references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1E.ii]	the symbol information,	Ogasawara discloses this limitation. *See* limitation [1E] above. To the extent it is found that Ogasawara does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Ogasawara does not anticipate this claim, Ogasawara renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Ogasawara with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1E.iii]	and a query of a database storing target object information associated with a plurality of target objects including the first target object;	Ogasawara discloses this limitation. *See* limitation [1E] above. To the extent it is found that Ogasawara does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Ogasawara does not anticipate this claim, Ogasawara renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Ogasawara with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success

Element	'278 Claim Recitation	Exemplary Citations in Reference
		for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1F]	receiving, via an address, first target object information associated with the first target object,	Ogasawara discloses receiving, via an address, first target object information associated with the first target object. *See* limitation [1A] above. Ogasawara at 6:42–50 • "The store server **10** or remote server **26** personal shopping application facilitates purchase transactions. Each message coming from a wireless telephone **18** is associated with the customer's telephone number, the customer ID, or some other unique identification. When the store server **10** or remote server **26** receives bar code data from the customer's wireless telephone **18**, then the store server **10** or remote server **26** searches a database and obtains a description and price for the item scanned. The item description and price is then transmitted to the customer's wireless telephone 18 and is preferably displayed upon the display 42 thereo" Ogasawara at 7:35 – 42 • "As each bar code is read, the purchase transaction program sends bar code data, such as SKU (Stock Keeping Unit) code or the Universal Product Code represented thereby, to the server and the server then preferably responds by sending a description and price for the product back to the wireless telephone 18, where the information is preferably shown upon the display 42 thereof. Also, the total price of items selected for purchase is preferably displayed." Ogasawara at 14:23 – 43 • "According to the preferred embodiment of the present invention, after each

Element	'278 Claim Recitation	Exemplary Citations in Reference
		product is selected, a description of the product and the price thereof is shown in the display 42 of the wireless telephone 18. This information may comprise part of the purchase transaction program, or alternatively may be communicated from the server 10, 26. According to the preferred embodiment of the present invention, the purchaser is given an opportunity to either confirm a purchase or to delete the item from the purchase list after each selection is made. The purchaser is preferably also given a choice to confirm or delete each purchase selection once all purchase selections have been made, prior to paying for the purchases. According to the preferred embodiment of the present invention, the purchaser indicates that all desired purchases have been made by pressing a predetermined key of the keypad 44. The wireless telephone 18 then responds by displaying the total price of all purchases and also preferably provides an opportunity to delete purchases from the list as discussed above." Ogasawara at 21:24–35 • "Once in the server, the numeric bar code data is associated to a corresponding product which can be retrieved from a Price Look Up (PLU) table or file. An "item selection complete" response is then issued by the server and received by the customer's wireless videophone through its telephone function and the response is displayed to the customer on the system's integral graphics display screen. The response might include the numeric bar code data of the items selected for purchase as well as the item name, item price, and the like, in order that the consumer may confirm that this is, indeed, the item that they selected." To the extent it is found that Ogasawara does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Ogasawara does not anticipate this claim, Ogasawara renders it obvious, either alone, in combination

Element	'278 Claim Recitation	Exemplary Citations in Reference
		with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Ogasawara with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1F.i]	wherein the first target object information comprises a response regarding the validity of the document.	Ogasawara discloses the first target object information comprises a response regarding the validity of the document *See* limitations [1E] and [1F] above. To the extent it is found that Ogasawara does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Ogasawara does not anticipate this claim, Ogasawara renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Ogasawara with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

2. Claim 3 (Non-asserted)

Claim	'004 Claim Recitation	Exemplary Citations in Reference

3	The method of claim 1, wherein the symbolic content comprises a unique code.	Ogasawara discloses the method of claim 1, wherein the symbolic content comprises a unique code. Ogasawara at 21:24–35 • "Once in the server, the numeric bar code data is associated to a corresponding product which can be retrieved from a Price Look Up (PLU) table or file. An "item selection complete" response is then issued by the server and received by the customer's wireless videophone through its telephone function and the response is displayed to the customer on the system's integral graphics display screen. The response might include the numeric bar code data of the items selected for purchase as well as the item name, item price, and the like, in order that the consumer may confirm that this is, indeed, the item that they selected." Ogasawara at 22:44–51 • "Indeed, merchandise item information may be presented in alpha-numeric form (i.e., an item name) or as a numeric code. In either case, character recognition software allows the wireless videophone to identify the item in a seamless fashion. Character recognition software supports identifying a product based on a numeric or alpha-numeric code, while pattern recognition software supports recognition of product characteristics based on bar code data." Ogasawara at 22:59–63 • "It will be understood that the system of the present invention allows bar code and/or alpha-numeric information to be captured from almost any form of display means by merely placing the information within the visual image field of the system's digital camera." Ogasawara at 20:41–58

		• "With particular reference to FIG. 13, once the purchase transaction program has been downloaded, a purchaser may use the wireless videophone, in particular the digital camera, to select various items which the consumer desires to purchase. Such selection is made by capturing an image of an item's UPC bar code data, or the like, within the image field of the digital camera. While the consumer is scanning a particular bar code with the digital camera, the image seen by the camera is displayed on the system's graphic display screen so that the consumer can visually verify that the entire image has been captured. If character and/or pattern recognition software was downloaded to the system as part of the program, the customer need only inform the system that the image is ready for processing by pressing a pre-defined key on the system's keyboard. Those skilled in the art will appreciate that various other codes, indicia, text, icons, and the like may be scanned and recognized by modern character and pattern recognition application routines." To the extent it is found that Ogasawara does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Ogasawara does not anticipate this claim, Ogasawara renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Preliminary Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Ogasawara with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

3. Claim 4

Claim	'004 Claim Recitation	Exemplary Citations in Reference

4	The method of claim 3, wherein the image includes a representation of the code.	Ogasawara discloses the method of claim 3, wherein the image includes a representation of the code. Ogasawara at 20:41–21:8 • "With particular reference to FIG. 13, once the purchase transaction program has been downloaded, a purchaser may use the wireless videophone, in particular the digital camera, to select various items which the consumer desires to purchase. Such selection is made by capturing an image of an item's UPC bar code data, or the like, within the image field of the digital camera. . . . If character and/or pattern recognition software was downloaded to the system as part of the program, the customer need only inform the system that the image is ready for processing by pressing a pre-defined key on the system's keyboard. Those skilled in the art will appreciate that various other codes, indicia, text, icons, and the like may be scanned and recognized by modern character and pattern recognition application routines. Thus, product selection need not be made in accordance with a UPC bar code but might be made with reference to a selection of product specific icons, each of which represent a particular product or service offered for purchase. For example, when grocery shopping, product specific icons might represent the stylistic outlines of a loaf of bread, carton of milk, a bunch of broccoli, and the like, allowing a consumer to quickly and efficiently select items for purchase without recourse to the complications of reading bar codes. However, no matter the form of the indicia captured by the digital video camera, the system operates in identical manner for each indicia and is governed solely by operation of the downloaded program which is coded to recognize either characters or patterns as appropriate to the application. However, in the following, the types of videographic images captured by the video camera will all be subsumed under the generic term 'bar code' image data."

Ogasawara at 21:24–35

- "Once in the server, the numeric bar code data is associated to a corresponding product which can be retrieved from a Price Look Up (PLU) table or file. An "item selection complete" response is then issued by the server and received by the customer's wireless videophone through its telephone function and the response is displayed to the customer on the system's integral graphics display screen. The response might include the numeric bar code data of the items selected for purchase as well as the item name, item price, and the like, in order that the consumer may confirm that this is, indeed, the item that they selected."

Ogasawara at 22:44–51

- "Indeed, merchandise item information may be presented in alpha-numeric form (i.e., an item name) or as a numeric code. In either case, character recognition software allows the wireless videophone to identify the item in a seamless fashion. Character recognition software supports identifying a product based on a numeric or alpha-numeric code, while pattern recognition software supports recognition of product characteristics based on bar code data."

Ogasawara at 22:59–63

- "It will be understood that the system of the present invention allows bar code and/or alpha-numeric information to be captured from almost any form of display means by merely placing the information within the visual image field of the system's digital camera."

To the extent it is found that Ogasawara does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Ogasawara does not anticipate this claim, Ogasawara renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover

		pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Ogasawara with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

4. Claim 5

Claim	'004 Claim Recitation	Exemplary Citations in Reference
5	The method of claim 3, further comprising detecting whether there is a match within a specified tolerance and returning a message indicating the absence of a match within a specified tolerance.	Ogasawara discloses the method of claim 3, further comprising detecting whether there is a match within a specified tolerance and returning a message indicating the absence of a match within a specified tolerance. Ogasawara at 7:34–41 • "As each bar code is read, the purchase transaction program sends bar code data, such as SKU (Stock Keeping Unit) code or the Universal Product Code represented thereby, to the server and the server then preferably responds by sending a description and price for the product back to the wireless telephone **18**, where the information is preferably shown upon the display **42** thereof. Also, the total price of items selected for purchase is preferably displayed." Ogasawara at 21:24–35 • "Once in the server, the numeric bar code data is associated to a corresponding product which can be retrieved from a Price Look Up (PLU) table or file. An "item selection complete" response is then issued by the server and received by the customer's wireless videophone through its telephone function and the response is displayed to the customer on the system's integral graphics display screen. The response might include

the numeric bar code data of the items selected for purchase as well as the item name, item price, and the like, in order that the consumer may confirm that this is, indeed, the item that they selected."

Ogasawara at 23:12–31

- "Advanced pattern recognition software is able to enhance the performance of the wireless videophone, in accord with the invention, by offering the capability to capture merchandise information from items that are not identified by either a bar code or an alpha-numeric label. Advanced pattern recognition software allows a consumer to capture a videographic image of an apple, for example, and to have the apple be recognized as such by the pattern recognition software. This capability is useful for any merchandise item having a distinct or identifiable shape or other visually identifiable characteristic. Such pattern recognition is accomplished within the wireless videophone, by specialized downloaded software. Alternatively, a visual image of the merchandise item is transmitted to a server, wherein pattern recognition software makes the requisite merchandise item identification. Here too, the system allows for the videographic image of an item to be transmitted to a server, which displays the item to a store clerk for manual recognition and identification. Merchandise information is then returned to the customer for display by the wireless videophone."

To the extent it is found that Ogasawara does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Ogasawara does not anticipate this claim, Ogasawara renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Ogasawara with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including

		because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

Nantworks, LLC v. Bank of America Corporation
Central District of California, No. 2:20-cv-7872
Invalidity Chart for U.S. Patent 9,031,278 based on
U.S. Patent No.6,036,086 to Sizer ("Sizer")

Exhibit F-5

Bank of America Corporation ("BofA" or "Defendant") provides this chart subject to all reservations, objections, statements, and disclaimers set forth herein and in BofA's Final Invalidity Contentions Cover Pleading, as well as any amendment, supplement, or modification thereof, which are incorporated herein by reference in their entirety.

On information and belief, the application that matured into Sizer was filed on March 28, 1997. Sizer published at least as of March 14, 2000, and is available as prior art at least under 35 U.S.C. § 102 (a), (b), (e).

As shown in the chart below, Sizer anticipates asserted claims 1, 4, 5 of the '278 Patent. To the extent Sizer is found not to expressly disclose certain limitations in the asserted claims, such limitations are inherent. To the extent Sizer is found not to anticipate any asserted claims or claim elements of the '278 Patent, the reference nevertheless renders those claims or claim elements obvious under 35 U.S.C. § 103, either alone or in combination with other art identified herein or in the cover pleading of BofA's Final Invalidity Contentions, for the reasons identified in the cover pleading. These Final Invalidity Contentions are not an admission by BofA that the accused products, including any current or past versions of the accused products, are covered by, or infringe the asserted claims, but are based instead on the recognition that if the claims are interpreted to be broad enough to encompass the accused products, the claims must also be construed to have that same scope when considering whether they are invalid.

The following chart is based on the Court's claim construction rulings (Dkt. Nos. 145, 153, 236), and to the extent not contradictory, Plaintiffs' apparent interpretations of the claims in Plaintiffs' "Preliminary" Final Infringement Contentions. BofA notes that in many instances, Plaintiffs' "Preliminary" Final Infringement Contentions ignore and/or contradict the Court's claim construction rulings. BofA does not accept the assumptions concerning the scope and meaning implicit in Plaintiffs' Preliminary Final Infringement Contentions, to the extent those assumptions are discernible, and reserves the right to challenge Plaintiffs' proposed (or implied) constructions. Furthermore, Plaintiffs' "Preliminary" Final Contentions fail to comply with the level of specificity required by SPR 2.1.2 and 2.1.3, and include improper reservation of rights to supplement with "additional ways in which BoA Accused Products infringe." Any alleged deficient discovery is of Plaintiffs own making. Most if not all of the deficiencies in Plaintiffs' Preliminary Final Contentions have been documented by Defendants through multiple letters since at least May 18, 2021 and up to Bank of America's most recent letter of August 19, 2023. Plaintiffs' shifting infringement contentions, deficiencies in its infringement contentions, and changing claim scope have prejudiced BofA from being able to reasonably prepare its defenses, including its invalidity contentions. BofA also reserves the right to revise and supplement these charts if and when Plaintiffs are permitted to provide revised or supplemental Infringement Contentions. Citations given in the chart below are merely representative of the respective elements to provide notice to Plaintiffs of BofA's invalidity theories and are not meant to be exhaustive.

1. Claim 1

Element	'278 Claim Recitation	Exemplary Citations in Reference
[1Pre]	A computer-assisted method, comprising:	Sizer discloses a computer assisted method. Sizer at FIG. 1 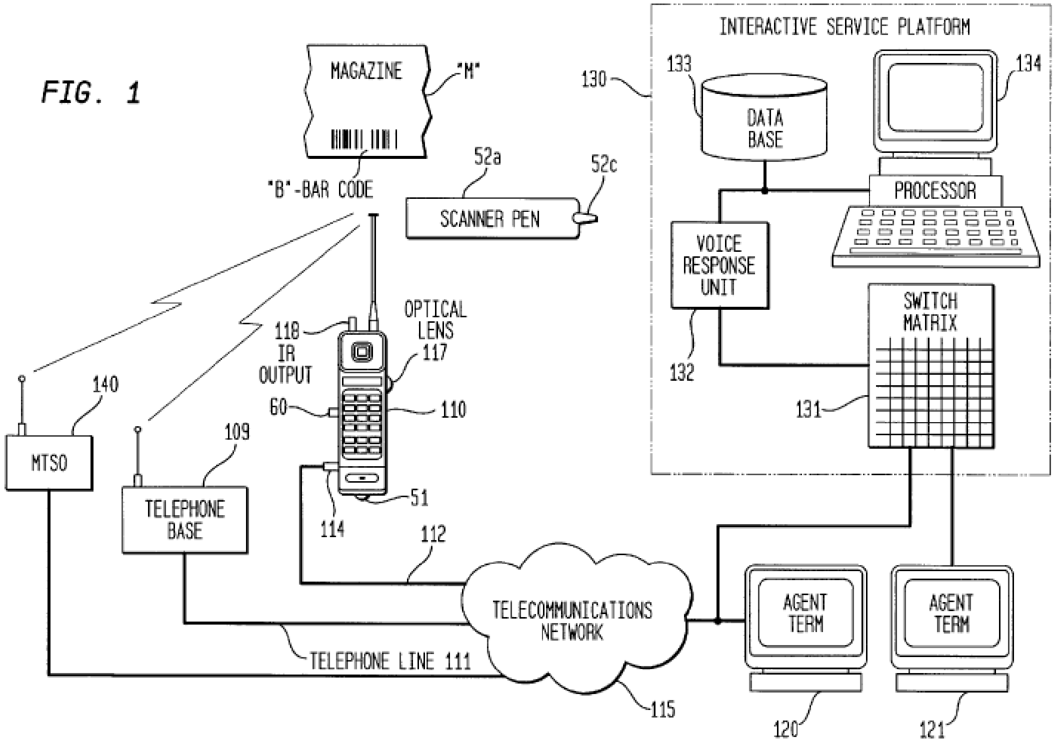 Sizer 3:31-33 • "Referring now to FIGS. 1 and 5, there is illustrated at 50 a scanner unit (FIG. 5) that is part of portable capture device 110 (FIG. 1)."

Element	'278 Claim Recitation	Exemplary Citations in Reference
[1A]	receiving, via a mobile device, an image comprising a representation of at least a portion of a document;	Sizer discloses receiving, via a mobile device, an image comprising a representation of at least a portion of a document. Sizer at FIG. 1 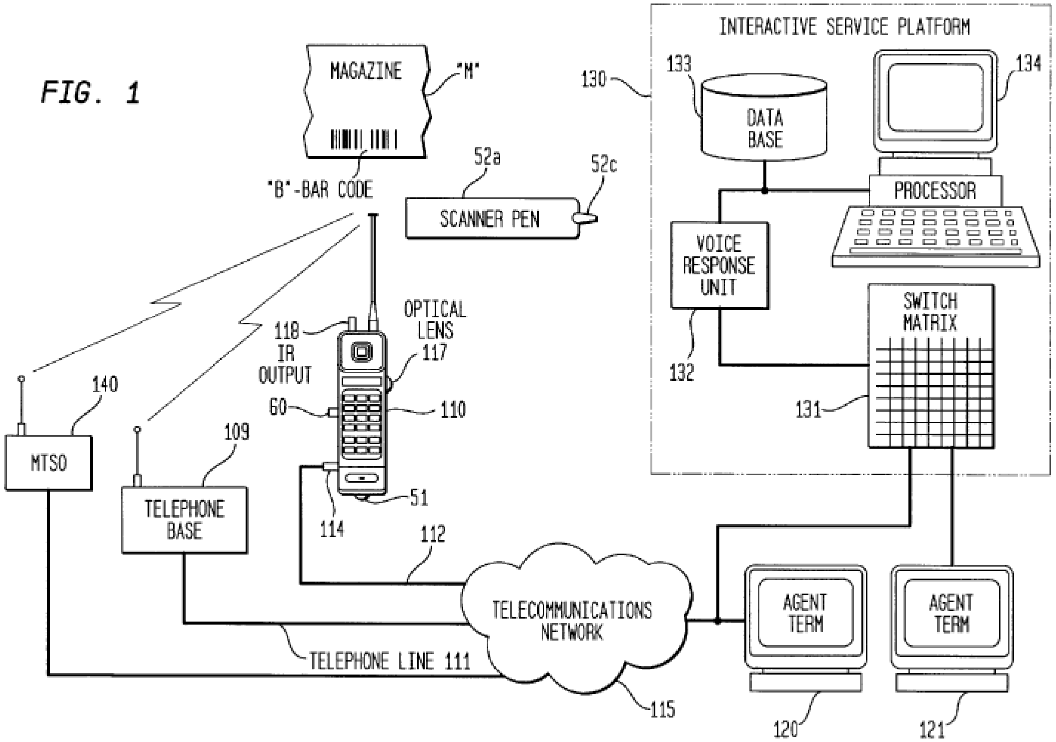 FIG. 1 Sizer 3:31-33

Element	'278 Claim Recitation	Exemplary Citations in Reference
		• "Referring now to FIGS. 1 and 5, there is illustrated at 50 a scanner unit (FIG. 5) that is part of portable capture device 110 (FIG. 1)." Sizer at 4:11-23 • "The scanner unit 50 of the present invention can be used in combination with marks that represents a unique code for transaction data. A standard bar code format could be used, which has encoded therein an 800 number to be called and other necessary information for completing a transaction. Xerox glyphs also could be appropriate because the glyphs can encode a large amount of data in a short space. FIG. 1 illustrates an object "M" in the form of a Magazine page having a bar code "B" adjacent an advertisement so that information can be gathered concerning the advertisement for initiating a transaction. Additionally, alphanumeric characters could be scanned, which specify a unique identifier for various advertisements." Sizer at 6:41-62 • "If the encoded non-perceptible data is inserted in the video portion of a television signal, the images displayed on a visual display area a monitor or television are captured by optical lens 117 and an associated photo-detector 203, which is arranged to supply an electrical signal to a filter/receiver 215 representing the image. On the other hand, if the encoded non-perceptible data is inserted in the audio portion of a television signal, e.g. the output from a speaker is captured by a microphone 211 and an associated amplifier 213, which is arranged to supply an electrical signal to filter/receiver 215 representing the sound energy. In either event, the encoded non-perceptible data is decoded in the filter/receiver 215 in a manner consistent with the manner in which the original data was encoded. Thus, filter/receiver 215 can be arranged to perform the same functionality as the elements illustrated in

Element	'278 Claim Recitation	Exemplary Citations in Reference
		FIG. 4 of the above referenced Broughton et al. patent, or, alternatively, the operation of filter/receiver 215 can be as described in conjunction with FIGS. 4 or 7 in the copending application of T. Sizer." Sizer 12:34-35 "It is further to be noted that the present invention can be used in the context of automobiles and mobile telephones." To the extent it is found that Sizer does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Sizer does not anticipate this claim, Sizer renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Sizer with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1B]	determining that symbolic content is on the at least the portion of the document based on the image;	Sizer discloses determining that symbolic content is on the at least the portion of the document based on the image. Sizer at 4:11-23 • "The scanner unit 50 of the present invention can be used in combination with marks that represents a unique code for transaction data. A standard bar code format could be used, which has encoded therein an 800 number to be called and other necessary information for completing a transaction. Xerox glyphs also could be appropriate because the glyphs can encode a large

Element	'278 Claim Recitation	Exemplary Citations in Reference
		amount of data in a short space. FIG. 1 illustrates an object "M" in the form of a Magazine page having a bar code "B" adjacent an advertisement so that information can be gathered concerning the advertisement for initiating a transaction. Additionally, alphanumeric characters could be scanned, which specify a unique identifier for various advertisements." Sizer at 9:33-53 • "For background purposes, a televised advertisement for a product could include additional ordering or "coupon" information embedded in a subliminal visual pattern within the video image displayed on a television. The coupon information describes the item(s) offered for sale, the regular price, the coupon value, and contains additional information needed for ordering. This embedded information could be transmitted between a television and capture device at a rate of about 60 bits per second, that is sufficient such that the necessary information can be transmitted in a very brief period of time. In the application just described, the embedded product information, dialing number, coupon information and necessary checkbits together comprise approximately 2,000 bits of information, so that between three and four seconds would be required for transmission. Once the coupon information is stored in capture device 110, the user of the capture device can supplement the coupon information with user entered data which can be entered using keypad 247. This supplemental information may include credit card or other billing information pertaining to the user, as an example." To the extent it is found that Sizer does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Sizer does not anticipate this claim, Sizer renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final

Element	'278 Claim Recitation	Exemplary Citations in Reference
		Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Sizer with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1C]	extracting symbol information based on the symbolic content according to symbol type;	Sizer discloses extracting symbol information based on the symbolic content according to symbol type. Sizer at 4:11-23 • "The scanner unit 50 of the present invention can be used in combination with marks that represents a unique code for transaction data. A standard bar code format could be used, which has encoded therein an 800 number to be called and other necessary information for completing a transaction. Xerox glyphs also could be appropriate because the glyphs can encode a large amount of data in a short space. FIG. 1 illustrates an object "M" in the form of a Magazine page having a bar code "B" adjacent an advertisement so that information can be gathered concerning the advertisement for initiating a transaction. Additionally, alphanumeric characters could be scanned, which specify a unique identifier for various advertisements." Sizer at 9:33-53 "For background purposes, a televised advertisement for a product could include additional ordering or "coupon" information embedded in a subliminal visual pattern within the video image displayed on a television. The coupon information describes the item(s) offered for sale, the regular price, the coupon value, and contains additional information needed for ordering. This embedded information could be transmitted between a television and capture device at a rate of about 60

Element	**'278 Claim Recitation**	**Exemplary Citations in Reference**
		bits per second, that is sufficient such that the necessary information can be transmitted in a very brief period of time. In the application just described, the embedded product information, dialing number, coupon information and necessary checkbits together comprise approximately 2,000 bits of information, so that between three and four seconds would be required for transmission. Once the coupon information is stored in capture device 110, the user of the capture device can supplement the coupon information with user entered data which can be entered using keypad 247. This supplemental information may include credit card or other billing information pertaining to the user, as an example." To the extent it is found that Sizer does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Sizer does not anticipate this claim, Sizer renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Sizer with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1D]	determining a validity of the document based at least in part on the image and the symbol information; and	Sizer discloses determining a validity of the document based at least in part on the image and the symbol information. Sizer at 6:18-38 • "If an outgoing telephone call in initiated by capture device 110, the call can illustratively be routed through telecommunications network 115 directly to an agent terminal 120, or to an interactive service platform 130, which is connected to another agent terminal 121. When a connection is

Element	'278 Claim Recitation	Exemplary Citations in Reference
		completed to platform 130, the call may be routed by a switch matrix 131 to a voice response unit 132 that operates under the control of a processor 134. Information included in the captured encoded can be provided to the attendants, or processor 134 can perform a look-up and retrieval operation in database 131, using encoded data as a key, in order to effectuate a desired transaction. For certain transactions, information connected with a transaction may also be transmitted from service platform 130 to capture device 110, either for display to a user, storage for later use, or output to a point of sale system or device. This information may, for example, be confirmation information, indicating that a transaction has been effected, or details regarding a just completed transaction, such as ticket information relating to a travel reservation." To the extent it is found that Sizer does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Sizer does not anticipate this claim, Sizer renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Sizer with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1E]	recognizing the document as a first target object based at least in part on	Sizer discloses this limitation. Sizer at 4:11-23

Element	'278 Claim Recitation	Exemplary Citations in Reference
		- "The scanner unit 50 of the present invention can be used in combination with marks that represents a unique code for transaction data. A standard bar code format could be used, which has encoded therein an 800 number to be called and other necessary information for completing a transaction. Xerox glyphs also could be appropriate because the glyphs can encode a large amount of data in a short space. FIG. 1 illustrates an object "M" in the form of a Magazine page having a bar code "B" adjacent an advertisement so that information can be gathered concerning the advertisement for initiating a transaction. Additionally, alphanumeric characters could be scanned, which specify a unique identifier for various advertisements." Sizer at 6:42-47 - "If the encoded non-perceptible data is inserted in the video portion of a television signal, the images displayed on a visual display area a monitor or television are captured by optical lens 117 and an associated photo-detector 203, which is arranged to supply an electrical signal to a filter/receiver 215 representing the image." Sizer at 9:33-53 - "For background purposes, a televised advertisement for a product could include additional ordering or "coupon" information embedded in a subliminal visual pattern within the video image displayed on a television. The coupon information describes the item(s) offered for sale, the regular price, the coupon value, and contains additional information needed for ordering. This embedded information could be transmitted between a television and capture device at a rate of about 60 bits per second, that is sufficient such that the necessary information can be transmitted in a very brief period of time. In the application just described, the embedded product information, dialing number, coupon information and necessary checkbits together comprise approximately 2,000 bits of information, so that between

Element	'278 Claim Recitation	Exemplary Citations in Reference
		three and four seconds would be required for transmission. Once the coupon information is stored in capture device 110, the user of the capture device can supplement the coupon information with user entered data which can be entered using keypad 247. This supplemental information may include credit card or other billing information pertaining to the user, as an example." Sizer at 10:27-37 - In the active mode, a user operates capture device 110 in real time, typically while scanning a product or bar code, or in more complicated devices viewing a video display or listening to an audio presentation. Scanning a product or bar code, or in more complicated devices the user activates the SAVE button 241. When button 241 is activated, a positive result occurs in step 503, thereby capturing the encoded data, displaying all or a portion of the data on display 225, and storing same in memory 221, all in step 507. Until the SAVE button is activated, capture device 110 may be in a "wait state" 505." Sizer at 11:11-15 - In accordance with one optional arrangement of the present invention, when data is captured and stored in step 507, they newly collected data is compared with previously stored data. When a match is found, any duplicate data is dropped. In addition, the information in various fields is examined, to assure that the data captured is in the appropriate format. If an error is detected in any data, the information may be deleted, and "new" data collected. This is accomplished by repeating capture step 507 several times To the extent it is found that Sizer does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Sizer does not anticipate this claim, Sizer renders it obvious, either alone, in combination with the

Element	'278 Claim Recitation	Exemplary Citations in Reference
		knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Sizer with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1E.i]	the image,	Sizer discloses this limitation. Sizer at 4:11-23 • "The scanner unit 50 of the present invention can be used in combination with marks that represents a unique code for transaction data. A standard bar code format could be used, which has encoded therein an 800 number to be called and other necessary information for completing a transaction. Xerox glyphs also could be appropriate because the glyphs can encode a large amount of data in a short space. FIG. 1 illustrates an object "M" in the form of a Magazine page having a bar code "B" adjacent an advertisement so that information can be gathered concerning the advertisement for initiating a transaction. Additionally, alphanumeric characters could be scanned, which specify a unique identifier for various advertisements." Sizer at 6:42-47 • "If the encoded non-perceptible data is inserted in the video portion of a television signal, the images displayed on a visual display area a monitor or television are captured by optical lens 117 and an associated photo-detector 203, which is arranged to supply an electrical signal to a filter/receiver 215 representing the image."

Element	'278 Claim Recitation	Exemplary Citations in Reference
		Sizer at 9:33-53 • "For background purposes, a televised advertisement for a product could include additional ordering or "coupon" information embedded in a subliminal visual pattern within the video image displayed on a television. The coupon information describes the item(s) offered for sale, the regular price, the coupon value, and contains additional information needed for ordering. This embedded information could be transmitted between a television and capture device at a rate of about 60 bits per second, that is sufficient such that the necessary information can be transmitted in a very brief period of time. In the application just described, the embedded product information, dialing number, coupon information and necessary checkbits together comprise approximately 2,000 bits of information, so that between three and four seconds would be required for transmission. Once the coupon information is stored in capture device 110, the user of the capture device can supplement the coupon information with user entered data which can be entered using keypad 247. This supplemental information may include credit card or other billing information pertaining to the user, as an example." Sizer at 10:27-37 • In the active mode, a user operates capture device 110 in real time, typically while scanning a product or bar code, or in more complicated devices viewing a video display or listening to an audio presentation. Scanning a product or bar code, or in more complicated devices the user activates the SAVE button 241. When button 241 is activated, a positive result occurs in step 503, thereby capturing the encoded data, displaying all or a portion of the data on display 225, and storing same in memory 221, all in step 507. Until the SAVE button is activated, capture device 110 may be in a "wait state" 505."

Element	'278 Claim Recitation	Exemplary Citations in Reference
		Sizer at 11:11-15 • In accordance with one optional arrangement of the present invention, when data is captured and stored in step 507, they newly collected data is compared with previously stored data. When a match is found, any duplicate data is dropped. In addition, the information in various fields is examined, to assure that the data captured is in the appropriate format. If an error is detected in any data, the information may be deleted, and "new" data collected. This is accomplished by repeating capture step 507 several times To the extent it is found that Sizer does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Sizer does not anticipate this claim, Sizer renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Sizer with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1E.ii]	the symbol information,	Sizer discloses this limitation. Sizer 2:57-59 • The marks to be scanned can include a conventional style bar code, alphanumeric characters, or Xerox glyphs on the surface of the object.

Element	'278 Claim Recitation	Exemplary Citations in Reference
		Sizer at 4:11-23 • "The scanner unit 50 of the present invention can be used in combination with marks that represents a unique code for transaction data. A standard bar code format could be used, which has encoded therein an 800 number to be called and other necessary information for completing a transaction. Xerox glyphs also could be appropriate because the glyphs can encode a large amount of data in a short space. FIG. 1 illustrates an object "M" in the form of a Magazine page having a bar code "B" adjacent an advertisement so that information can be gathered concerning the advertisement for initiating a transaction. Additionally, alphanumeric characters could be scanned, which specify a unique identifier for various advertisements." Sizer at 10:27-37 • In the active mode, a user operates capture device 110 in real time, typically while scanning a product or bar code, or in more complicated devices viewing a video display or listening to an audio presentation. Scanning a product or bar code, or in more complicated devices the user activates the SAVE button 241. When button 241 is activated, a positive result occurs in step 503, thereby capturing the encoded data, displaying all or a portion of the data on display 225, and storing same in memory 221, all in step 507. Until the SAVE button is activated, capture device 110 may be in a "wait state" 505." To the extent it is found that Sizer does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Sizer does not anticipate this claim, Sizer renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final

Element	'278 Claim Recitation	Exemplary Citations in Reference
		Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Sizer with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1E.iii]	and a query of a database storing target object information associated with a plurality of target objects including the first target object;	Sizer discloses this limitation. Sizer at 6:18-38 • "If an outgoing telephone call in initiated by capture device 110, the call can illustratively be routed through telecommunications network 115 directly to an agent terminal 120, or to an interactive service platform 130, which is connected to another agent terminal 121. When a connection is completed to platform 130, the call may be routed by a switch matrix 131 to a voice response unit 132 that operates under the control of a processor 134. Information included in the captured encoded can be provided to the attendants, or processor 134 can perform a look-up and retrieval operation in database 131, using encoded data as a key, in order to effectuate a desired transaction. For certain transactions, information connected with a transaction may also be transmitted from service platform 130 to capture device 110, either for display to a user, storage for later use, or output to a point of sale system or device. This information may, for example, be confirmation information, indicating that a transaction has been effected, or details regarding a just completed transaction, such as ticket information relating to a travel reservation." To the extent it is found that Sizer does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Sizer does not anticipate

Element	'278 Claim Recitation	Exemplary Citations in Reference
		this claim, Sizer renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Sizer with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1F]	receiving, via an address, first target object information associated with the first target object,	Sizer discloses receiving, via an address, first target object information associated with the first target object. Sizer at 6:18-38 • "If an outgoing telephone call in initiated by capture device 110, the call can illustratively be routed through telecommunications network 115 directly to an agent terminal 120, or to an interactive service platform 130, which is connected to another agent terminal 121. When a connection is completed to platform 130, the call may be routed by a switch matrix 131 to a voice response unit 132 that operates under the control of a processor 134. Information included in the captured encoded can be provided to the attendants, or processor 134 can perform a look-up and retrieval operation in database 131, using encoded data as a key, in order to effectuate a desired transaction. For certain transactions, information connected with a transaction may also be transmitted from service platform 130 to capture device 110, either for display to a user, storage for later use, or output to a point of sale system or device. This information may, for example, be confirmation information, indicating that a transaction has been effected, or details regarding a just completed transaction, such as ticket information relating

Element	'278 Claim Recitation	Exemplary Citations in Reference
		to a travel reservation." To the extent it is found that Sizer does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Sizer does not anticipate this claim, Sizer renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Sizer with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1F.i]	wherein the first target object information comprises a response regarding the validity of the document.	Sizer discloses the first target object information comprises a response regarding the validity of the document. Sizer at 6:18-38 • "If an outgoing telephone call in initiated by capture device 110, the call can illustratively be routed through telecommunications network 115 directly to an agent terminal 120, or to an interactive service platform 130, which is connected to another agent terminal 121. When a connection is completed to platform 130, the call may be routed by a switch matrix 131 to a voice response unit 132 that operates under the control of a processor 134. Information included in the captured encoded can be provided to the attendants, or processor 134 can perform a look-up and retrieval operation in database 131, using encoded data as a key, in order to effectuate a desired transaction. For certain transactions, information connected with a transaction may also be transmitted from service platform 130 to capture device 110, either

Element	'278 Claim Recitation	Exemplary Citations in Reference
		for display to a user, storage for later use, or output to a point of sale system or device. This information may, for example, be confirmation information, indicating that a transaction has been effected, or details regarding a just completed transaction, such as ticket information relating to a travel reservation." To the extent it is found that Sizer does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Sizer does not anticipate this claim, Sizer renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Sizer with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

2. Claim 3 (Non-asserted)

Claim	'278 Claim Recitation	Exemplary Citations in Reference
3	The method of claim 1, wherein the symbolic content comprises a unique code.	Sizer discloses the method of claim 1, wherein the symbolic content comprises a unique code. Sizer 2:57-59 • The marks to be scanned can include a conventional style bar code, alphanumeric characters, or Xerox glyphs on the surface of the object.

		Sizer at 4:11-23 • "The scanner unit 50 of the present invention can be used in combination with marks that represents a unique code for transaction data. A standard bar code format could be used, which has encoded therein an 800 number to be called and other necessary information for completing a transaction. Xerox glyphs also could be appropriate because the glyphs can encode a large amount of data in a short space. FIG. 1 illustrates an object "M" in the form of a Magazine page having a bar code "B" adjacent an advertisement so that information can be gathered concerning the advertisement for initiating a transaction. Additionally, alphanumeric characters could be scanned, which specify a unique identifier for various advertisements." Sizer at 9:33-53 • "For background purposes, a televised advertisement for a product could include additional ordering or "coupon" information embedded in a subliminal visual pattern within the video image displayed on a television. The coupon information describes the item(s) offered for sale, the regular price, the coupon value, and contains additional information needed for ordering. This embedded information could be transmitted between a television and capture device at a rate of about 60 bits per second, that is sufficient such that the necessary information can be transmitted in a very brief period of time. In the application just described, the embedded product information, dialing number, coupon information and necessary checkbits together comprise approximately 2,000 bits of information, so that between three and four seconds would be required for transmission. Once the coupon information is stored in capture device 110, the user of the capture device can supplement the coupon information with user entered data which can be entered using keypad 247. This supplemental information may include credit card or other

		billing information pertaining to the user, as an example." Sizer at 10:27-37 • In the active mode, a user operates capture device 110 in real time, typically while scanning a product or bar code, or in more complicated devices viewing a video display or listening to an audio presentation. Scanning a product or bar code, or in more complicated devices the user activates the SAVE button 241. When button 241 is activated, a positive result occurs in step 503, thereby capturing the encoded data, displaying all or a portion of the data on display 225, and storing same in memory 221, all in step 507. Until the SAVE button is activated, capture device 110 may be in a "wait state" 505." To the extent it is found that Sizer does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Sizer does not anticipate this claim, Sizer renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Sizer with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

3. Claim 4

Claim	'278 Claim Recitation	Exemplary Citations in Reference

4	The method of claim 3, wherein the image includes a representation of the code.	Sizer discloses the method of claim 3, wherein the image includes a representation of the code. Sizer 2:57-59 • The marks to be scanned can include a conventional style bar code, alphanumeric characters, or Xerox glyphs on the surface of the object. Sizer at 4:11-23 • "The scanner unit 50 of the present invention can be used in combination with marks that represents a unique code for transaction data. A standard bar code format could be used, which has encoded therein an 800 number to be called and other necessary information for completing a transaction. Xerox glyphs also could be appropriate because the glyphs can encode a large amount of data in a short space. FIG. 1 illustrates an object "M" in the form of a Magazine page having a bar code "B" adjacent an advertisement so that information can be gathered concerning the advertisement for initiating a transaction. Additionally, alphanumeric characters could be scanned, which specify a unique identifier for various advertisements." Sizer at 10:27-37 • In the active mode, a user operates capture device 110 in real time, typically while scanning a product or bar code, or in more complicated devices viewing a video display or listening to an audio presentation. Scanning a product or bar code, or in more complicated devices the user activates the SAVE button 241. When button 241 is activated, a positive result occurs in step 503, thereby capturing the encoded data, displaying all or a portion of the data on display 225, and storing same in memory 221, all in step 507. Until the SAVE button is activated, capture device 110 may be in a "wait state" 505."

		To the extent it is found that Sizer does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Sizer does not anticipate this claim, Sizer renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Sizer with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

4. Claim 5

Claim	'278 Claim Recitation	Exemplary Citations in Reference
5	The method of claim 3, further comprising detecting whether there is a match within a specified tolerance and returning a message indicating the absence of a match within a specified tolerance.	To the extent it is found that Sizer does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Sizer does not anticipate this claim, Sizer renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Sizer with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

Nantworks, LLC v. Bank of America Corporation
Central District of California, No. 2:20-cv-7872
Invalidity Chart for U.S. Patent 9,031,278 based on
U.S. Patent No. 6,666,337 to Harris ("Harris")

Exhibit F-6

Bank of America Corporation ("BofA" or "Defendant") provides this chart subject to all reservations, objections, statements, and disclaimers set forth herein and in BofA's Preliminary Invalidity Contentions Cover Pleading, as well as any amendment, supplement, or modification thereof, which are incorporated herein by reference in their entirety.

On information and belief, the application that matured into Harris was filed on July 18, 2000. Harris published at least as of December 23, 2003, and is available as prior art at least under 35 U.S.C. § 102(e).

As shown in the chart below, Harris anticipates asserted claims 1, 3, 4, 5 of the '278 Patent. To the extent Harris is found not to expressly disclose certain limitations in the asserted claims, such limitations are inherent. To the extent Harris is found not to anticipate any asserted claims or claim elements of the '278 Patent, the reference nevertheless renders those claims or claim elements obvious under 35 U.S.C. § 103, either alone or in combination with other art identified herein or in the cover pleading of BofA's Preliminary Invalidity Contentions, for the reasons identified in the cover pleading. These Preliminary Invalidity Contentions are not an admission by BofA that the accused products, including any current or past versions of the accused products, are covered by, or infringe the asserted claims, but are based instead on the recognition that if the claims are interpreted to be broad enough to encompass the accused products, the claims must also be construed to have that same scope when considering whether they are invalid.

The following chart is based on the Court's claim construction rulings (Dkt. Nos. 145, 153, 236), and to the extent not contradictory, Plaintiffs' apparent interpretations of the claims in Plaintiffs' "Preliminary" Final Infringement Contentions. BofA notes that in many instances, Plaintiffs' "Preliminary" Final Infringement Contentions ignore and/or contradict the Court's claim construction rulings. BofA does not accept the assumptions concerning the scope and meaning implicit in Plaintiffs' Preliminary Final Infringement Contentions, to the extent those assumptions are discernible, and reserves the right to challenge Plaintiffs' proposed (or implied) constructions. Furthermore, Plaintiffs' "Preliminary" Final Contentions fail to comply with the level of specificity required by SPR 2.1.2 and 2.1.3, and include improper reservation of rights to supplement with "additional ways in which BoA Accused Products infringe." Any alleged deficient discovery is of Plaintiffs own making. Most if not all of the deficiencies in Plaintiffs' Preliminary Final Contentions have been documented by Defendants through multiple letters since at least May 18, 2021 and up to Bank of America's most recent letter of August 19, 2023. Plaintiffs' shifting infringement contentions, deficiencies in its infringement contentions, and changing claim scope have prejudiced BofA from being able to reasonably prepare its defenses, including its invalidity contentions. BofA also reserves the right to revise and supplement these charts if and when Plaintiffs are permitted to provide revised or supplemental Infringement Contentions. Citations given in the chart below are merely representative of the respective elements to provide notice to Plaintiffs of BofA's invalidity theories and are not meant to be exhaustive.

1. Claim 1

Element	'278 Claim Recitation	Exemplary Citations in Reference
[1Pre]	A computer-assisted method, comprising:	Harris discloses a computer-assisted method.
[1A]	receiving, via a mobile device, an image comprising a representation of at least a portion of a document;	Harris discloses receiving, via a mobile device, an image comprising a representation of at least a portion of a document. Harris at FIG. 1A and FIG. 1B 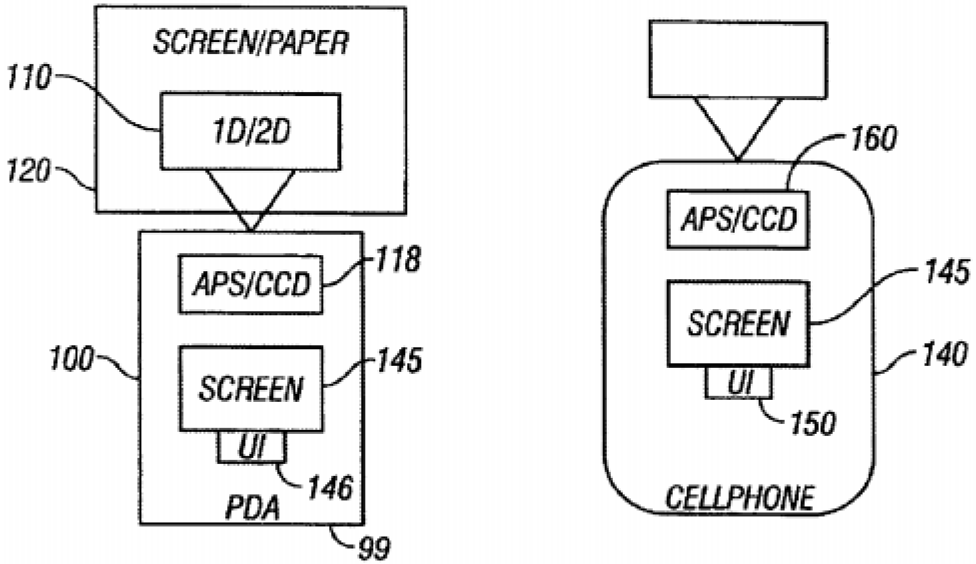 FIG. 1A FIG. 1B Harris at 2: 3-19 • "The bar codes can be imaged/scanned in a number of different ways. One

Element	'278 Claim Recitation	Exemplary Citations in Reference
		embodiment discloses using a camera to input and decode these bar codes. The embodiments are shown in FIGS. 1A and 1B. FIG. 1A describes using **a personal digital assistant 100** as the input device. One preferred input device of this type is the Palm V(TM) type hand held computer. A bar code scanner can be used, such as the commercially available Symbol Technology SPT 1700. Alternatively, a camera add-on unit can be added to the Palm V and used as described herein. The PDA includes a screen 145 and user interface 146 all within the same housing 99. FIG. 1B shows the client being **a cellular telephone** which also may include a screen 145 and user interface 150. The cellular telephone can be associated with either a dedicated bar code scanner or an image sensor 160 of the type used to obtain photographs for video telephony." Harris at FIG. 4

Element	'278 Claim Recitation	Exemplary Citations in Reference
		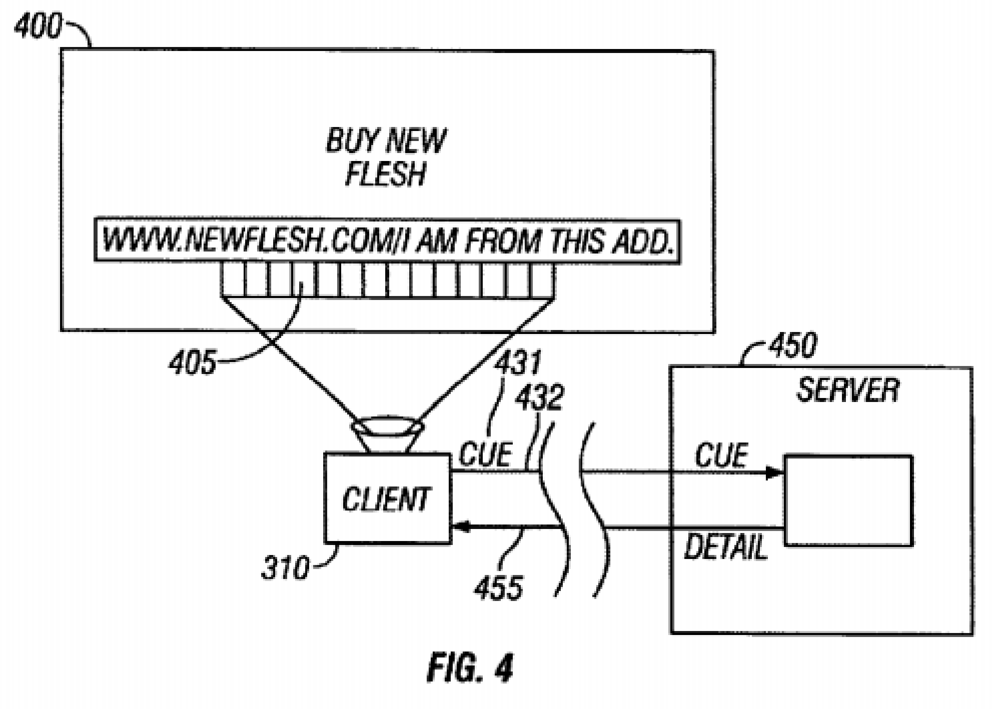 FIG. 4 Harris at 4: 9-15 • "Another embodiment shown in FIG. 4 uses bar codes to enter information into the computer. The FIG. 4 embodiment stores scannable non-alphanumeric information, e.g. bar code information, as some part of a communication—here an advertisement. The advertisement can be a print advertisement, a television advertisement, or an Internet advertisement for example." Harris at 4:15-30

Element	**'278 Claim Recitation**	**Exemplary Citations in Reference**
		- "The advertisement 400 includes a bar code 405 therein. The bar code 405 is associated with the advertisement, and includes some additional information about the advertisement. For example, the bar code may include the web site address of the company preparing the advertisement, or appointment information about the advertisement, or a "vcf file" or the like which is an importable file with address information about the company or author sponsoring the advertisement. In operation, the user brings one of the clients, either the cellular telephone 310 or PDA 315, into range of the bar code 405. The client reads the bar code and decodes it as noted above. The decoded information can represent ASCII information, a compressed file such as a zip file, G code information, or any other compressed or non-compressed information." To the extent it is found that Harris does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Harris does not anticipate this claim, Harris renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Preliminary Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Harris with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1B]	determining that symbolic content is on the at least the portion of the document based on the image;	Harris discloses determining that symbolic content is on the at least the portion of the document based on the image. Harris at FIG. 4

Element	'278 Claim Recitation	Exemplary Citations in Reference
		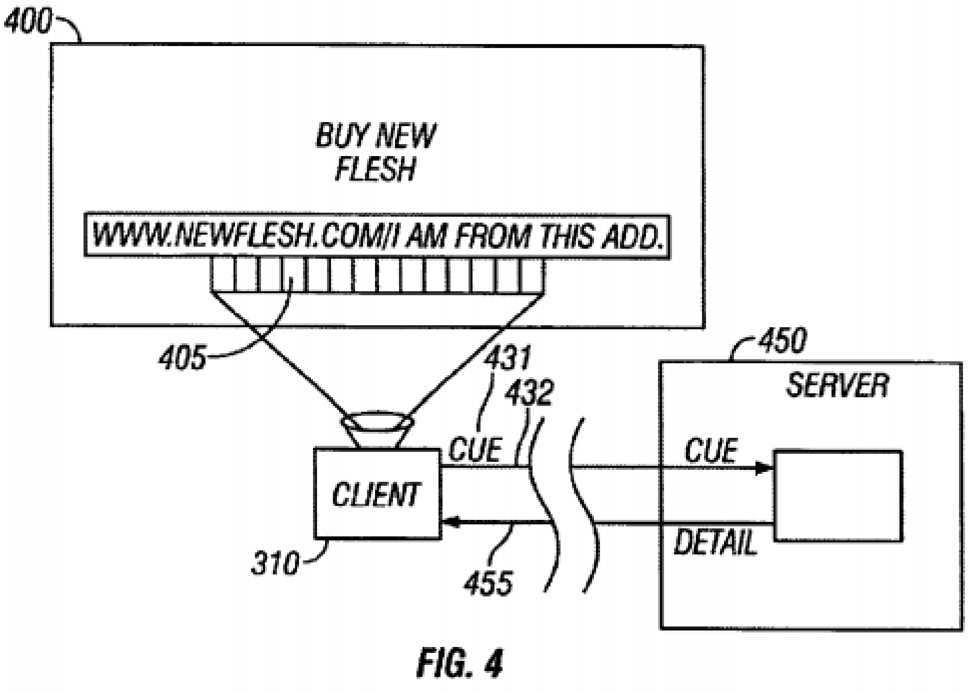 FIG. 4 Harris at 4: 9-15 • "Another embodiment shown in FIG. 4 uses bar codes to enter information into the computer. The FIG. 4 embodiment stores scannable non-alphanumeric information, e.g. bar code information, as some part of a communication—here an advertisement. The advertisement can be a print advertisement, a television advertisement, or an Internet advertisement for example." Harris at 4:15-30

Element	'278 Claim Recitation	Exemplary Citations in Reference
		• "The advertisement 400 includes a bar code 405 therein. The bar code 405 is associated with the advertisement, and includes some additional information about the advertisement. For example, the bar code may include the web site address of the company preparing the advertisement, or appointment information about the advertisement, or a "vcf file" or the like which is an importable file with address information about the company or author sponsoring the advertisement. In operation, the user brings one of the clients, either the cellular telephone 310 or PDA 315, into range of the bar code 405. The client reads the bar code and decodes it as noted above. The decoded information can represent ASCII information, a compressed file such as a zip file, G code information, or any other compressed or non-compressed information." To the extent it is found that Harris does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Harris does not anticipate this claim, Harris renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Preliminary Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Harris with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1C]	extracting symbol information based on the symbolic content according to symbol type;	Harris discloses extracting symbol information based on the symbolic content according to symbol type. Harris at FIG. 4

Element	'278 Claim Recitation	Exemplary Citations in Reference
		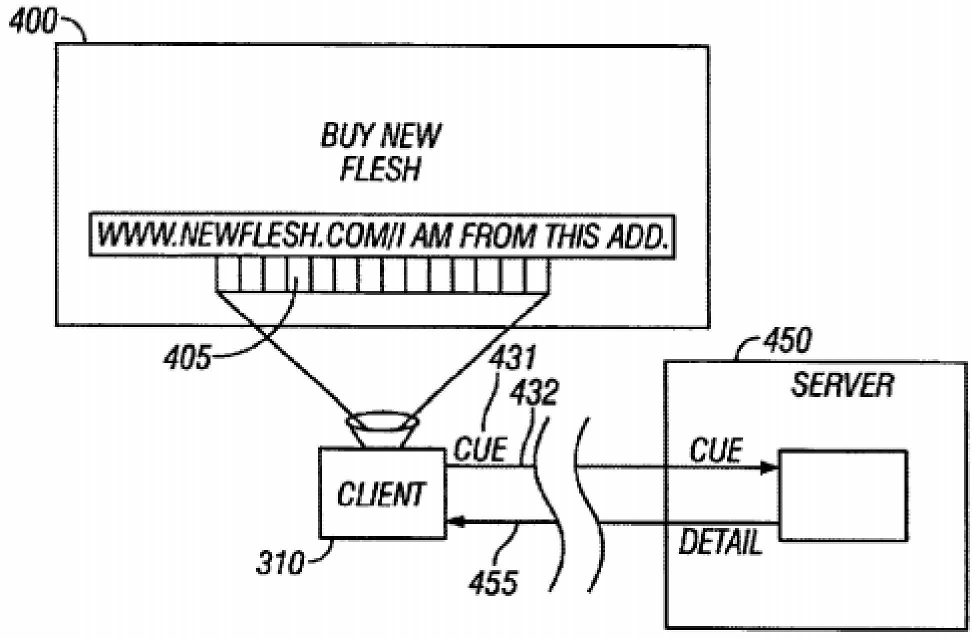 FIG. 4 Harris at 4: 9-15 • "Another embodiment shown in FIG. 4 uses bar codes to enter information into the computer. The FIG. 4 embodiment stores scannable non-alphanumeric information, e.g. bar code information, as some part of a communication—here an advertisement. The advertisement can be a print advertisement, a television advertisement, or an Internet advertisement for example." Harris at 4:15-30

Element	**'278 Claim Recitation**	**Exemplary Citations in Reference**
		- "The advertisement 400 includes a bar code 405 therein. The bar code 405 is associated with the advertisement, and includes some additional information about the advertisement. For example, the bar code may include the web site address of the company preparing the advertisement, or appointment information about the advertisement, or a "vcf file" or the like which is an importable file with address information about the company or author sponsoring the advertisement. In operation, the user brings one of the clients, either the cellular telephone 310 or PDA 315, into range of the bar code 405. The client reads the bar code and decodes it as noted above. The decoded information can represent ASCII information, a compressed file such as a zip file, G code information, or any other compressed or non-compressed information." To the extent it is found that Harris does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Harris does not anticipate this claim, Harris renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Preliminary Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Harris with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1D]	determining a validity of the document based at least in part on the image and the symbol information; and	Harris discloses determining a validity of the document based at least in part on the image and the symbol information. Harris at FIG. 7

Element	'278 Claim Recitation	Exemplary Citations in Reference
		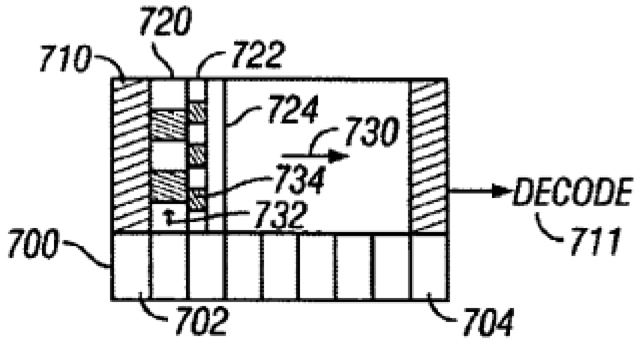 FIG. 7A 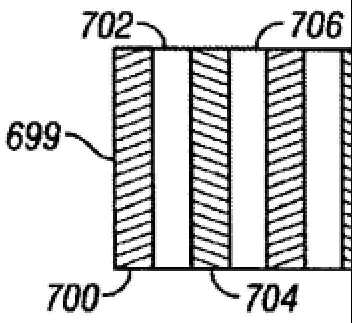 FIG. 7B 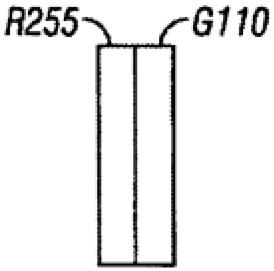 FIG. 7C Harris at 5:56-62 - "The present system teaches, in FIG. 7A, a special bar code which includes increased capacity bar code information, as well as backwards compatibility with previous bar code scanners. The code shown in FIG. 7A includes two parts. A first part 710 is found as legal when scanned for

Element	'278 Claim Recitation	Exemplary Citations in Reference
		linear codes. A second part 710 registers as invalid /illegal when scanned in this way." To the extent it is found that Harris does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Harris does not anticipate this claim, Harris renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Preliminary Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Harris with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1E]	recognizing the document as a first target object based at least in part on	Harris discloses this limitation. Harris at FIG. 4

Element	'278 Claim Recitation	Exemplary Citations in Reference
		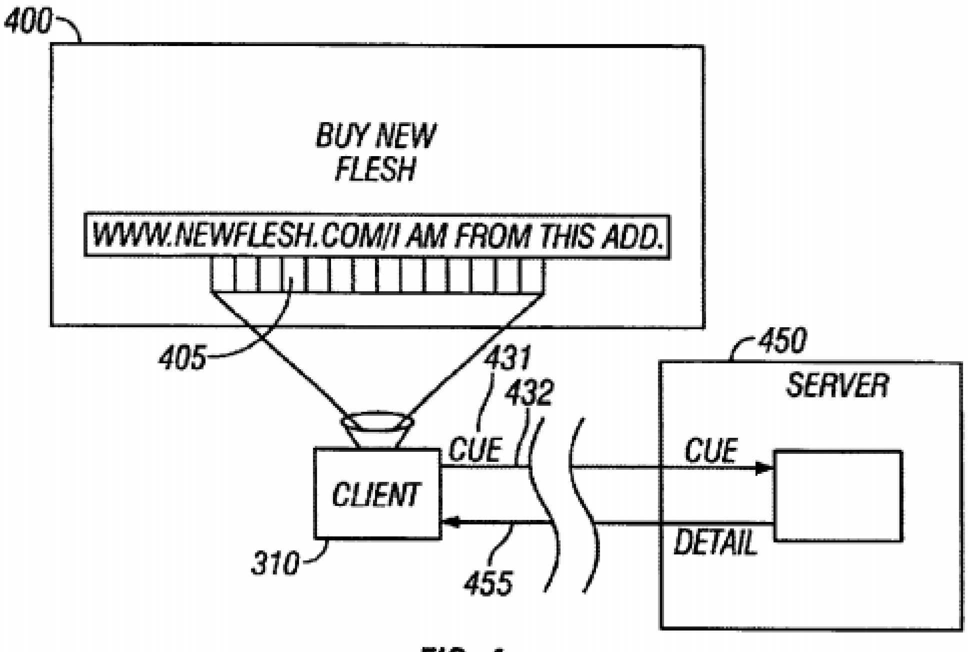 FIG. 4 Harris at 4:41-48 • "In this system, the information can represent a cue, e.g. a number, that is associated with a special function. The cue is used during a hot sync to access a publicly available information network, e.g. by direct connection or by the Internet. The cue 431 is sent over the Internet 432 to **the server 450**, and addresses more detailed information in a memory of the server. The server returns that information as 455, and the client receives the more detailed information." To the extent it is found that Harris does not expressly disclose this limitation, the

Element	'278 Claim Recitation	Exemplary Citations in Reference
		limitation is inherent. And to the extent it is found that Harris does not anticipate this claim, Harris renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Preliminary Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Harris with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1E.i]	the image,	Harris discloses this limitation. Harris at FIG. 4

Element	'278 Claim Recitation	Exemplary Citations in Reference
		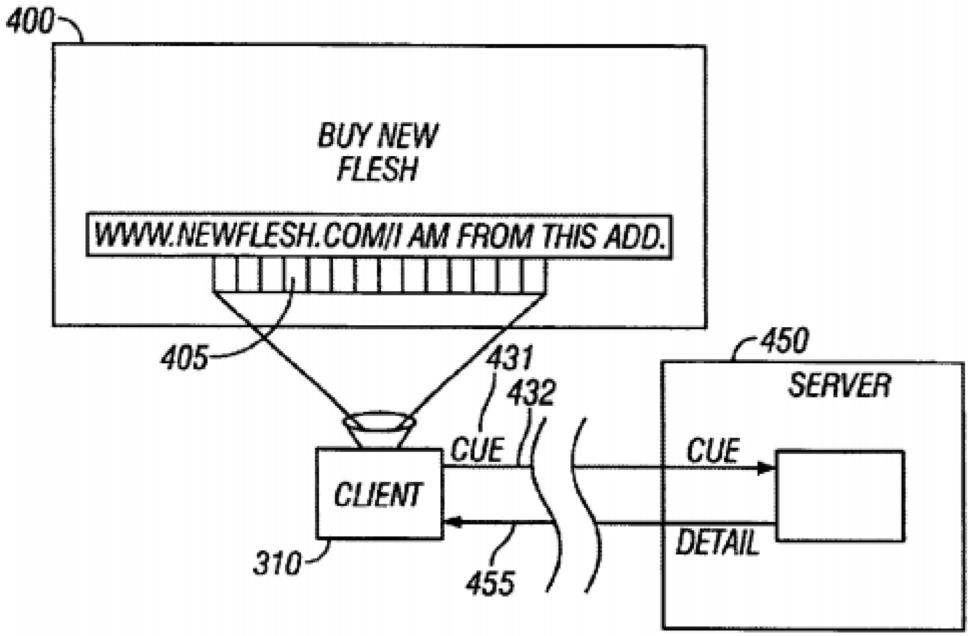 FIG. 4 Harris at 4:41-48 • "In this system, the information can represent a cue, e.g. a number, that is associated with a special function. The cue is used during a hot sync to access a publicly available information network, e.g. by direct connection or by the Internet. The cue 431 is sent over the Internet 432 to **the server 450**, and addresses more detailed information in a memory of the server. The server returns that information as 455, and the client receives the more detailed information." To the extent it is found that Harris does not expressly disclose this limitation, the

Element	**'278 Claim Recitation**	**Exemplary Citations in Reference**
		limitation is inherent. And to the extent it is found that Harris does not anticipate this claim, Harris renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Preliminary Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Harris with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1E.ii]	the symbol information,	Harris discloses this limitation. Harris at FIG. 4

Element	'278 Claim Recitation	Exemplary Citations in Reference
		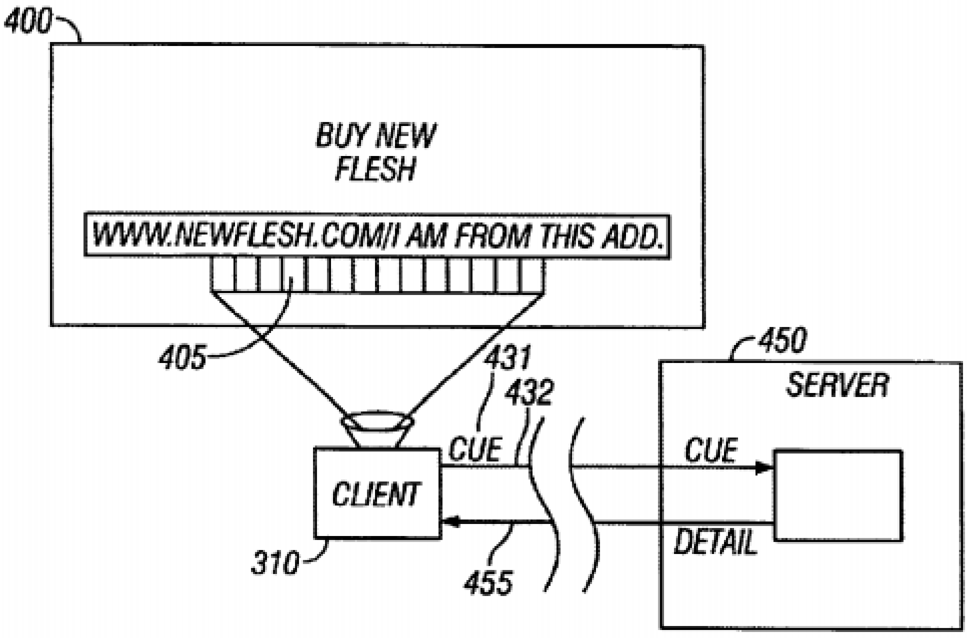 FIG. 4 Harris at 4:41-48 • "In this system, the information can represent a cue, e.g. a number, that is associated with a special function. The cue is used during a hot sync to access a publicly available information network, e.g. by direct connection or by the Internet. The cue 431 is sent over the Internet 432 to **the server 450**, and addresses more detailed information in a memory of the server. The server returns that information as 455, and the client receives the more detailed information." To the extent it is found that Harris does not expressly disclose this limitation, the

Element	**'278 Claim Recitation**	**Exemplary Citations in Reference**
		limitation is inherent. And to the extent it is found that Harris does not anticipate this claim, Harris renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Preliminary Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Harris with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1E.iii]	and a query of a database storing target object information associated with a plurality of target objects including the first target object;	Harris discloses this limitation. Harris at FIG. 4

Element	'278 Claim Recitation	Exemplary Citations in Reference
		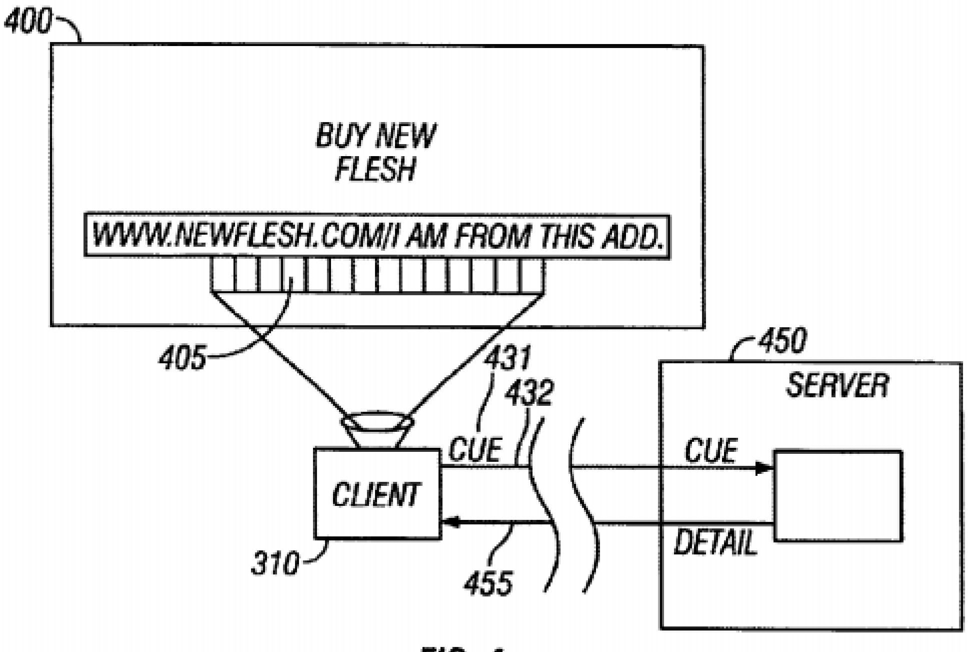 FIG. 4 Harris at 4:41-48 • "In this system, the information can represent a cue, e.g. a number, that is associated with a special function. The cue is used during a hot sync to access a publicly available information network, e.g. by direct connection or by the Internet. The cue 431 is sent over the Internet 432 to **the server 450**, and addresses more detailed information in a memory of the server. The server returns that information as 455, and the client receives the more detailed information." Harris at 4:55-58

Element	'278 Claim Recitation	Exemplary Citations in Reference
		• "During hot sync, the database returns the full text of the detailed information, e.g., "visit the website at http://www.pdascan.com/˜more to get a free gift." Any desired length or size of information can be returned." To the extent it is found that Harris does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Harris does not anticipate this claim, Harris renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Preliminary Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Harris with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1F]	receiving, via an address, first target object information associated with the first target object,	Harris discloses receiving, via an address, first target object information associated with the first target object. Harris at FIG. 4

Element	'278 Claim Recitation	Exemplary Citations in Reference
		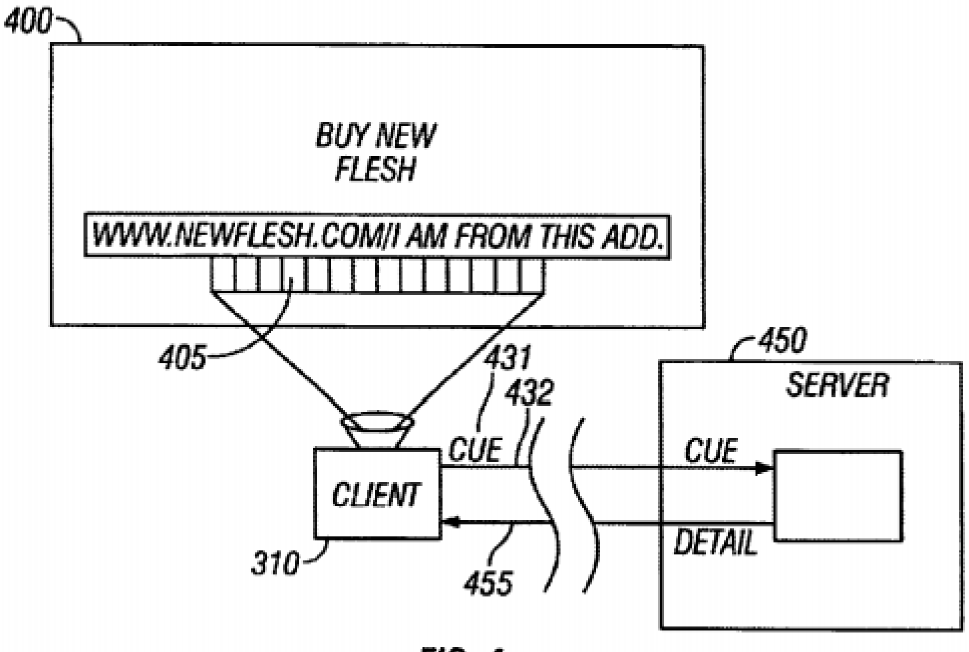 FIG. 4 Harris at 4:41-48 - "In this system, the information can represent a cue, e.g. a number, that is associated with a special function. The cue is used during a hot sync to access a publicly available information network, e.g. by direct connection or by the Internet. The cue 431 is sent over the Internet 432 to **the server 450**, and addresses more detailed information in a memory of the server. The server returns that information as 455, and the client receives the more detailed information." Harris at 4:55-58

Element	**'278 Claim Recitation**	**Exemplary Citations in Reference**
		• "During hot sync, the database returns the full text of the detailed information, e.g., "visit the website at http://www.pdascan.com/˜more to get a free gift." Any desired length or size of information can be returned." To the extent it is found that Harris does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Harris does not anticipate this claim, Harris renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Preliminary Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Harris with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1F.i]	wherein the first target object information comprises a response regarding the validity of the document.	Harris discloses the first target object information comprises a response regarding the validity of the document. Harris at FIG. 7

Element	'278 Claim Recitation	Exemplary Citations in Reference
		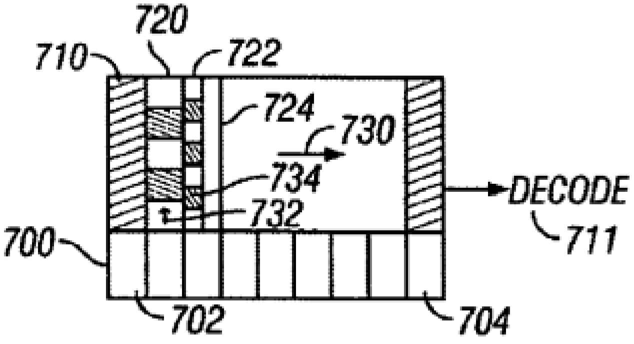 FIG. 7A 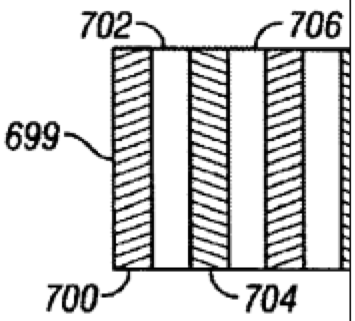 FIG. 7B 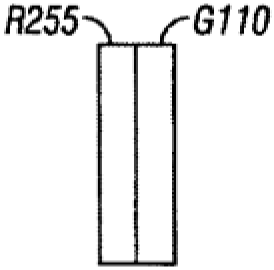 FIG. 7C Harris at 5:56-62 • "The present system teaches, in FIG. 7A, a special bar code which includes increased capacity bar code information, as well as backwards compatibility with previous bar code scanners. The code shown in FIG. 7A includes two parts. A first part 710 is found as legal when scanned for

Element	'278 Claim Recitation	Exemplary Citations in Reference
		linear codes. A second part 710 registers as invalid /illegal when scanned in this way." To the extent it is found that Harris does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Harris does not anticipate this claim, Harris renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Preliminary Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Harris with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

2. Claim 3 (Non-asserted)

Claim	'278 Claim Recitation	Exemplary Citations in Reference
3	The method of claim 1, wherein the symbolic content comprises a unique code.	Harris discloses the method of claim 1, wherein the symbolic content comprises a unique code. Harris at 4: 9-15 • "Another embodiment shown in FIG. 4 uses bar codes to enter information into the computer. The FIG. 4 embodiment stores scannable non-alphanumeric information, e.g. bar code information, as some part of a communication—here an advertisement. The advertisement can be a print advertisement, a television advertisement, or an Internet advertisement for example." Harris at 4:15-30

		• "The advertisement 400 includes a bar code 405 therein. The bar code 405 is associated with the advertisement, and includes some additional information about the advertisement. For example, the bar code may include the web site address of the company preparing the advertisement, or appointment information about the advertisement, or a "vcf file" or the like which is an importable file with address information about the company or author sponsoring the advertisement. In operation, the user brings one of the clients, either the cellular telephone 310 or PDA 315, into range of the bar code 405. The client reads the bar code and decodes it as noted above. The decoded information can represent ASCII information, a compressed file such as a zip file, G code information, or any other compressed or non-compressed information." Harris at FIG. 7

		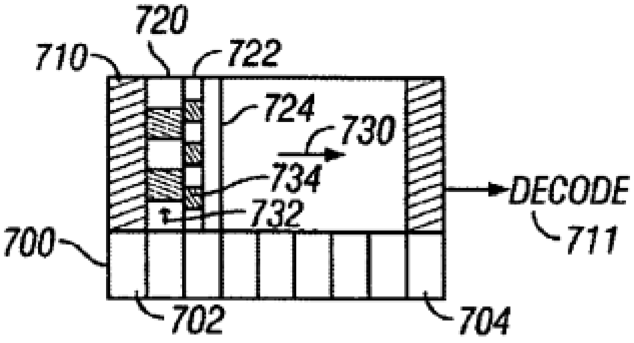 FIG. 7A 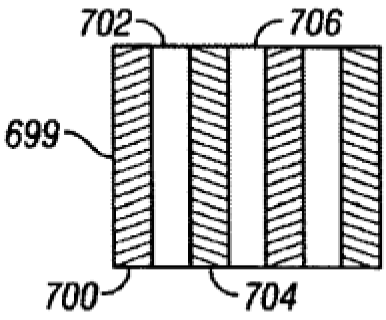 FIG. 7B 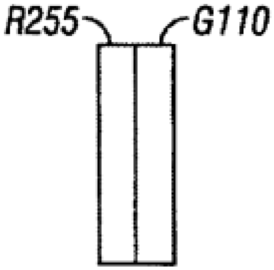 FIG. 7C Harris at 6:36-54 • "In the context of this system, the linear bar code information can represent an address for look up code of the type described above with reference to FIG. 4. The enhanced or non-linear information 710 can represent the total information. A person with a sufficiently advanced bar code scanner can read the entire information. A person with only a linear bar code scanner, however, may scan only the information 700, and then updates the information via a hot sync.

		An additional way of using this information, for example for scanning products, is also contemplated. Scan part 700 may include a Universal Product Code or UPC. Scanning part 710 may include additional information about the product, such as a description, or picture. The picture of the product may be displayed to the sales clerk, so that the sales clerk can verify that the product being purchased is actually the product that the user is presenting. Both parts of the bar code may represent information about the product. The part 700 represents basic information and the part 710 represents advanced information." To the extent it is found that Harris does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Harris does not anticipate this claim, Harris renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Preliminary Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Harris with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

3. Claim 4

Claim	'278 Claim Recitation	Exemplary Citations in Reference
4	The method of claim 3, wherein the image includes a representation of the code.	Harris discloses the method of claim 3, wherein the image includes a representation of the code. Harris at 4: 9-15 • "Another embodiment shown in FIG. 4 uses bar codes to enter information into the computer. The FIG. 4 embodiment stores scannable non-alphanumeric information, e.g. bar code information, as some part of a communication—here

		an advertisement. The advertisement can be a print advertisement, a television advertisement, or an Internet advertisement for example." Harris at 4:15-30 • "The advertisement 400 includes a bar code 405 therein. The bar code 405 is associated with the advertisement, and includes some additional information about the advertisement. For example, the bar code may include the web site address of the company preparing the advertisement, or appointment information about the advertisement, or a "vcf file" or the like which is an importable file with address information about the company or author sponsoring the advertisement. In operation, the user brings one of the clients, either the cellular telephone 310 or PDA 315, into range of the bar code 405. The client reads the bar code and decodes it as noted above. The decoded information can represent ASCII information, a compressed file such as a zip file, G code information, or any other compressed or non-compressed information." Harris at FIG. 7

		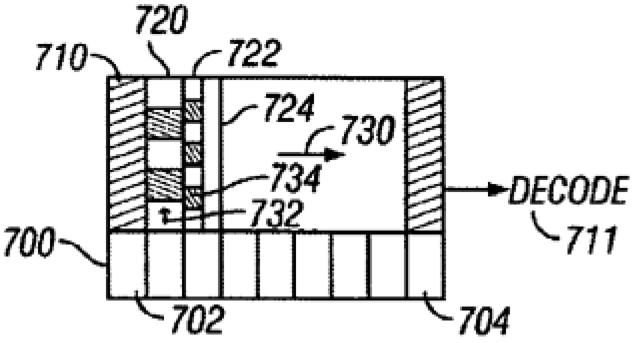 FIG. 7A 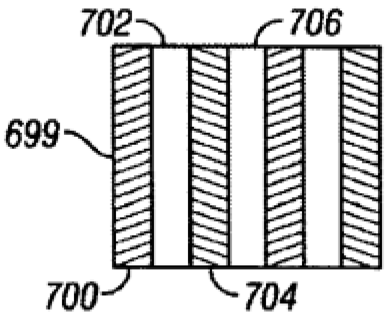 FIG. 7B 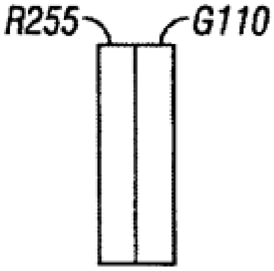 FIG. 7C Harris at 6:36-54 • "In the context of this system, the linear bar code information can represent an address for look up code of the type described above with reference to FIG. 4. The enhanced or non-linear information 710 can represent the total information. A person with a sufficiently advanced bar code scanner can read the entire information. A person with only a linear bar code scanner, however, may scan only the information 700, and then updates the information via a hot sync.

		An additional way of using this information, for example for scanning products, is also contemplated. Scan part 700 may include a Universal Product Code or UPC. Scanning part 710 may include additional information about the product, such as a description, or picture. The picture of the product may be displayed to the sales clerk, so that the sales clerk can verify that the product being purchased is actually the product that the user is presenting. Both parts of the bar code may represent information about the product. The part 700 represents basic information and the part 710 represents advanced information." To the extent it is found that Harris does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Harris does not anticipate this claim, Harris renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Preliminary Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Harris with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

4. Claim 5

Claim	'278 Claim Recitation	Exemplary Citations in Reference
5	The method of claim 3, further comprising detecting whether there is a match within a specified tolerance and returning a message indicating the absence of a match within a specified tolerance.	Harris discloses the method of claim 4, further comprising detecting whether there is a match within a specified tolerance and returning a message indicating the absence of a match within a specified tolerance. To the extent it is found that Harris does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Harris does not anticipate this claim, Harris renders it obvious, either alone, in combination with

		the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Preliminary Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Harris with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

Nantworks, LLC v. Bank of America Corporation
Central District of California, No. 2:20-cv-7872
Invalidity Chart for U.S. Patent 9,031,278 based on
U.S. Patent No. 6,722,569 to Ehrhart et al. ("Ehrhart")

Exhibit F-11

Bank of America Corporation ("BofA" or "Defendant") provides this chart subject to all reservations, objections, statements, and disclaimers set forth herein and in BofA's Final Invalidity Contentions Cover Pleading, as well as any amendment, supplement, or modification thereof, which are incorporated herein by reference in their entirety.

On information and belief, the application that matured into Ehrhart was filed on July 13, 2001, claiming the priority date of July 13, 2001. Ehrhart published at least as of April 20, 2004, and is available as prior art at least under 35 U.S.C. § 102(e).

As shown in the chart below, Ehrhart anticipates asserted claims 1, 3, 4, 5 of the '278 Patent. To the extent Ehrhart is found not to expressly disclose certain limitations in the asserted claims, such limitations are inherent. To the extent Ehrhart is found not to anticipate any asserted claims or claim elements of the '278 Patent, the reference nevertheless renders those claims or claim elements obvious under 35 U.S.C. § 103, either alone or in combination with other art identified in the cover pleading of BofA's Invalidity Contentions, for the reasons identified in the cover pleading. These Invalidity Contentions are not an admission by BofA that the accused products, including any current or past versions of the accused products, are covered by, or infringe the asserted claims, but are based instead on the recognition that if the claims are interpreted to be broad enough to encompass the accused products, the claims must also be construed to have that same scope when considering whether they are invalid.

The following chart is based on the Court's claim construction rulings (Dkt. Nos. 145, 153, 236), and to the extent not contradictory, Plaintiffs' apparent interpretations of the claims in Plaintiffs' "Preliminary" Final Infringement Contentions. BofA notes that in many instances, Plaintiffs' "Preliminary" Final Infringement Contentions ignore and/or contradict the Court's claim construction rulings. BofA does not accept the assumptions concerning the scope and meaning implicit in Plaintiffs' Preliminary Final Infringement Contentions, to the extent those assumptions are discernible, and reserves the right to challenge Plaintiffs' proposed (or implied) constructions. Furthermore, Plaintiffs' "Preliminary" Final Contentions fail to comply with the level of specificity required by SPR 2.1.2 and 2.1.3, and include improper reservation of rights to supplement with "additional ways in which BoA Accused Products infringe." Any alleged deficient discovery is of Plaintiffs own making. Most if not all of the deficiencies in Plaintiffs' Preliminary Final Contentions have been documented by Defendants through multiple letters since at least May 18, 2021 and up to Bank of America's most recent letter of August 19, 2023. Plaintiffs' shifting infringement contentions, deficiencies in its infringement contentions, and changing claim scope have prejudiced BofA from being able to reasonably prepare its defenses, including its invalidity contentions. BofA also reserves the right to revise and supplement these charts if and when Plaintiffs are permitted to provide revised or supplemental Infringement Contentions. Citations given in the chart below are merely representative of the respective elements to provide notice to Plaintiffs of BofA's invalidity theories and are not meant to be exhaustive.

1. Claim 1

Element	'278 Claim Recitation	Exemplary Citations in Reference
[1Pre]	A computer-assisted method, comprising:	Ehrhart discloses a computer-assisted method. "The present invention relates to an optical reader that includes a color imaging assembly that generates color imaging data. An image analysis circuit determines if the acquired image should be characterized as a color photograph or as including a graphical symbol. A processing circuit processes the imaging data based on the image analysis circuit's determination of whether the image is a graphical symbol or a color photograph. The present invention allows a user to acquire and process both color images and graphical symbols, such as bar codes, text, OCR symbols or signatures. The optical reader of the present invention is also configured to associate an acquired image with at least one other acquired image." Ehrhart at Abstract. "In another aspect, the present invention includes a method for acquiring an image of an object with an optical reader." Ehrhart at 2:62-64.
[1A]	receiving, via a mobile device, an image comprising a representation of at least a portion of a document;	Ehrhart discloses receiving, via a mobile device, an image comprising a representation of at least a portion of a document. *Id.* at [1Pre]. "The optical reader of the present invention automatically, or through manual Selection, determines whether a captured image is a color photographic image or, a color image that includes a graphical Symbol. Subsequently, the optical reader of the present invention processes the acquired imaging data in accordance with that determination. The optical reader of the present invention is operative to acquire and associate a plurality of acquired images.

Element	'278 Claim Recitation	Exemplary Citations in Reference
		One aspect of the present invention is an optical reader. The optical reader includes a color imaging assembly for acquiring an image of an object, the color imaging assembly generating imaging data corresponding to the image. An image analysis circuit is coupled to the color imaging assembly. The image analysis circuit being configured to determine if the color imaging data includes at least one graphical Symbol. The image is classified as a graphical Symbol, or the image is classified as a color photograph if the color imaging data does not include at least one graphical Symbol. A processing circuit is coupled to the image analysis circuit. The processing circuit is operative to process the imaging data based on the determination. In another aspect, the present invention includes an optical reader for capturing an image of an object. The optical reader includes a color imaging assembly for converting the image of the object into color digital data corresponding to the image." Ehrhart at 1:57-2:16. "In another aspect, the present invention includes an optical reader for capturing an image of an object. The optical reader includes a color imaging assembly for capturing the image as color imaging data. A classification circuit is coupled to the color imaging assembly, the classification circuit being configured to process at least a portion of the color imaging data to thereby Select one of a plurality of classifications, whereby the image is classified as a color photographic image, or as an image that includes at least one graphical Symbol. An automatic mode Selector is coupled to the classification circuit, the automatic mode Selector being configured to Select an optical reader mode in accordance with the Selected classification. A processor is coupled to the classification circuit, the processor being programmed to process the color imaging data in accordance with the optical reader mode Selected by the automatic mode Selector." Ehrhart at 2:26-42.

Element	'278 Claim Recitation	Exemplary Citations in Reference
		"As embodied herein, and depicted in FIGS. 1A-1D, perspective views of the optical reader in accordance with various embodiments of the present invention are disclosed. FIG. 1A shows the underside of hand held wireless optical reader 10. FIG. 1B shows the top of the optical reader depicted in FIG. 1A. Optical reader 10 includes housing 100, antenna 102, window 104 and trigger 12. Window 104 accommodates illumination assembly 20 and imaging assembly 30. As shown in FIG. 1B, the top side of reader 10 includes function keys 14, alphanumeric key pad 16, and display 60. In one embodiment, function keys 14 include an enter key and up and down cursor keys. FIG. 1C is also a hand held wireless optical reader 10. Reader 10 includes function keys 14, alphanumeric key pad 16, writing stylus 18, display 60, and signature block 62. Stylus 18 is employed by a user to write his signature in signature block 62. FIG. 1D shows yet another embodiment of optical reader 10 of the present invention. In this embodiment, reader 10 includes a gun-shaped housing 100. Display 60 and keypad 16 are disposed on a top portion of gun-shaped housing 100, whereas trigger 12 is disposed on the underside of the top portion of housing 100. Housing 100 also includes window 104 that accommodates illumination assembly 20 and imaging assembly 30. Wire 106 is disposed at the butt-end of housing 100. Wire 106 provides optical reader 10 with a hard wired communication link for external devices such as a host processor or other data collection devices." Ehrhart at 5:51-6:10. To the extent it is found that Ehrhart does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Ehrhart does not anticipate this claim, Ehrhart renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Ehrhart with any of the other prior art references identified in the cover

Element	'278 Claim Recitation	Exemplary Citations in Reference
		pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1B]	determining that symbolic content is on the at least the portion of the document based on the image;	Ehrhart discloses determining that symbolic content is on the at least the portion of the document based on the image. *Id.* at [1Pre-A]. "In accordance with the invention, the present invention for an optical reader includes a color imaging assembly for acquiring color imaging data. An image analysis circuit determines if the acquired image includes at least one graphical symbol. A processing circuit processes the imaging data based on the determination of whether the image includes at least one graphical symbol. The present invention allows a user to read graphical symbols, such as bar codes, text, OCR characters or signatures using a color imager. The color optical reader of the present invention is configured to automatically determine whether a color image includes a graphical symbol, or is merely a color photographic image. The optical reader of the present invention also is operative to associate one acquired image with at least one subsequently acquired image." Ehrhart at 5:36-50. To the extent it is found that Ehrhart does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Ehrhart does not anticipate this claim, Ehrhart renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Ehrhart with any of the other prior art references identified in the cover

Element	'278 Claim Recitation	Exemplary Citations in Reference
		pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1C]	extracting symbol information based on the symbolic content according to symbol type;	Ehrhart discloses extracting symbol information based on the symbolic content according to symbol type. "In step 1004, processor 40 identifies the page features. Processor 40 analyzes the page and divides it into blank and non-blank portions. The non-blank portions are analyzed to distinguish text regions from non-text regions. After determining the layout of the page, processor 40 uses black-to-white transitions to determine degrees of skew. In step 1008, horizontal white spaces are identified to separate lines of text. In step 1010, vertical white spaces are identified within each line of text to thereby separate individual words and characters from each other. In Step 1014, a character recognition algorithm is used in an attempt to recognize each individual character." Ehrhart at 12:55-66. "As embodied herein and depicted in FIG. 13, an example of image association in accordance with the present invention is disclosed. One or ordinary skill in the art will recognize that associated images 1300 can be disposed on paper, displayed electronically on display 60, or displayed electronically sing other electronic means, such as a computer monitor. In this example, the first image captured is color photograph 1302 which shows a damaged parcel. The second image captured is bar code 1304 affixed to the side of the damaged parcel. Processor 40 decodes bar code 1304 and associates decoded bar code data 1306 with color photograph 1302. In this example, the user elected to associate a third image, signature 1308. Thus, personnel viewing record 1300 may reasonably conclude that a damaged parcel was delivered to Company XYZ, and that the person signing for the parcel delivery was someone named John W. Smith." Ehrhart at 13:27-43.

Element	'278 Claim Recitation	Exemplary Citations in Reference
		"In the Automatic imaging mode, processor 40 is programmed to analyze the color imaging data to determine if an acquired image includes a graphical symbol or is merely a color photographic image. If it makes the determination that the color image includes a graphical symbol, it further analyzes the acquired image and classifies it as a bar code, OCR symbol, text, or a signature. Based on the classification, optical reader 10 jumps to the appropriate routine in EROM 48. The semi-automatic mode is similar. Thus, in the automatic or semi-automatic modes, the bar code scanning mode, the OCR/text mode, the signature capture mode, the color photography mode, and the association mode are controlled by the application program, not by the user." Ehrhart at 8:36-49.

		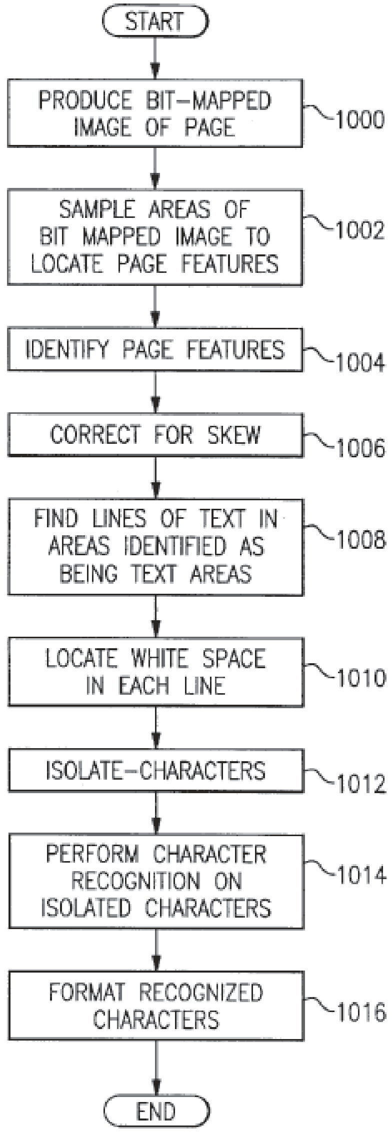 START PRODUCE BIT-MAPPED IMAGE OF PAGE 1000 SAMPLE AREAS OF BIT MAPPED IMAGE TO LOCATE PAGE FEATURES 1002 IDENTIFY PAGE FEATURES 1004 CORRECT FOR SKEW 1006 FIND LINES OF TEXT IN AREAS IDENTIFIED AS BEING TEXT AREAS 1008 LOCATE WHITE SPACE IN EACH LINE 1010 ISOLATE-CHARACTERS 1012 PERFORM CHARACTER RECOGNITION ON ISOLATED CHARACTERS 1014 FORMAT RECOGNIZED CHARACTERS 1016 END FIG.10

Element	'278 Claim Recitation	Exemplary Citations in Reference
		Ehrhart at Fig. 10.

		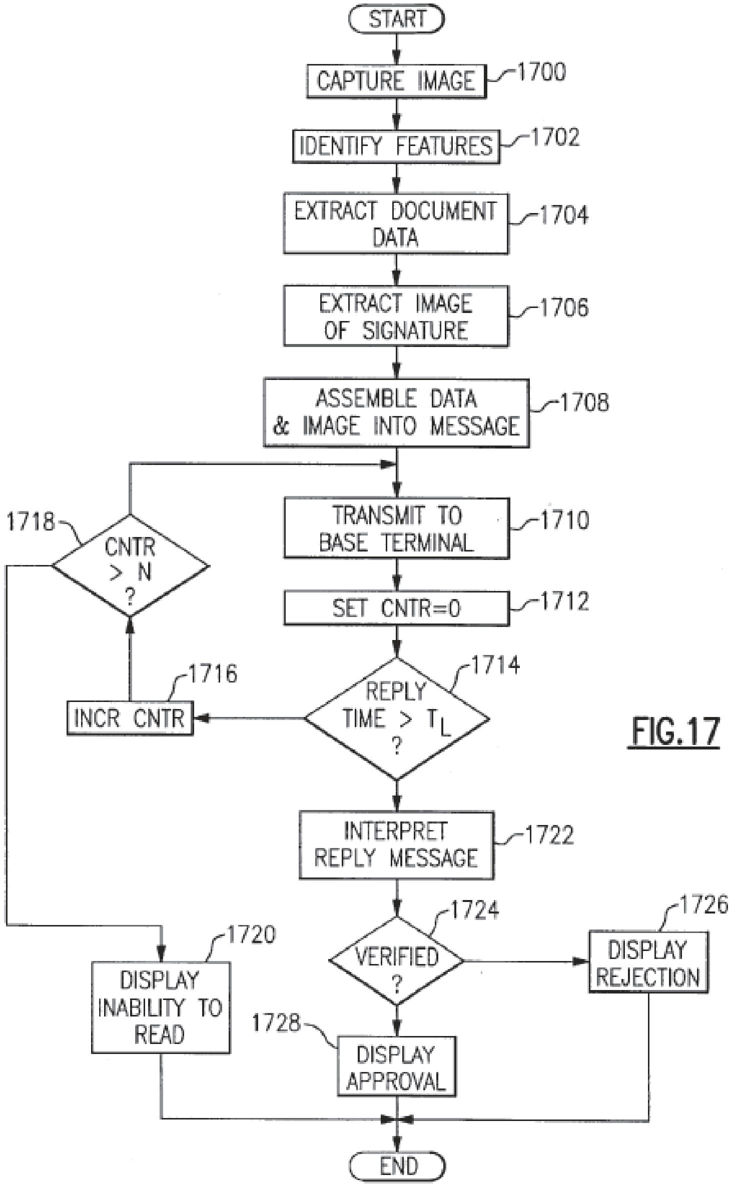START CAPTURE IMAGE 1700 IDENTIFY FEATURES 1702 EXTRACT DOCUMENT DATA 1704 EXTRACT IMAGE OF SIGNATURE 1706 ASSEMBLE DATA & IMAGE INTO MESSAGE 1708 TRANSMIT TO BASE TERMINAL 1710 SET CNTR=0 1712 REPLY TIME > T_L ? 1714 INCR CNTR 1716 CNTR > N ? 1718 INTERPRET REPLY MESSAGE 1722 VERIFIED ? 1724 DISPLAY REJECTION 1726 DISPLAY INABILITY TO READ 1720 DISPLAY APPROVAL 1728 END FIG.17

Element	'278 Claim Recitation	Exemplary Citations in Reference
		Ehrhart at Fig. 17. To the extent it is found that Ehrhart does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Ehrhart does not anticipate this claim, Ehrhart renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Ehrhart with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1D]	determining a validity of the document based at least in part on the image and the symbol information; and	Ehrhart discloses determining a validity of the document based at least in part on the image and the symbol information. *Id.* at [1Pre-C]. "The user may also click on OCR/Text icon 660. Clicking icon 660 provides the user with a check validation mode, a text scanning mode, or a bi-tonal image capture mode. The check validation mode is performed in conjunction with network services." Ehrhart at 8:58-62. "In steps 1712 and 1714, processor 40 initializes a counter and begins waiting for a reply from the host computer. In steps 1714-1718, if the reply is not received within time limit TL, the counter CNTR is incremented and the message is re-transmitted. After several attempts, if CNTR>N (N being an integer), processor 40 outputs a fault message. If the reply message is received within time limit TL, processor interprets

Element	'278 Claim Recitation	Exemplary Citations in Reference
		the reply in step 1722. If the extracted data and the signature match information stored in the database accessible by the host computer, an approval message is displayed. If the extracted data and the signature do not match information stored in the database accessible by the host computer, a disapproval message is displayed. The dynamic signature verification embodiment is similar to the static embodiment described immediately above. In the dynamic version, the user provides his signature using stylus 18 and signature block 62, as shown in FIG. 1C. Signature block 62 provides processor 40 with the dynamic parameters recorded during signature. The dynamic parameters are transmitted to a host processor, as described above." Ehrhart at 15:31-50. To the extent it is found that Ehrhart does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Ehrhart does not anticipate this claim, Ehrhart renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Ehrhart with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1E]	recognizing the document as a first target object based at least in part on	Ehrhart discloses this limitation. *Id.* at [1Pre-D]. "In accordance with the invention, the present invention for an optical reader includes a color imaging assembly for acquiring color imaging data. An image analysis circuit determines if the acquired image includes at least one graphical symbol. A processing circuit processes the imaging data based on the determination

Element	'278 Claim Recitation	Exemplary Citations in Reference
		of whether the image includes at least one graphical symbol. The present invention allows a user to read graphical symbols, such as bar codes, text, OCR characters or signatures using a color imager. The color optical reader of the present invention is configured to automatically determine whether a color image includes a graphical symbol, or is merely a color photographic image. The optical reader of the present invention also is operative to associate one acquired image with at least one subsequently acquired image." Ehrhart at 5:36-50. "In the embodiment depicted in FIG. 2, ASIC 44 is implemented using a programmable logic array (PLA) device. In a similar embodiment, ASIC 44 is implemented using a field programmable gate array (FPGA) device. ASIC 44 is tasked with controlling the image acquisition process, and the storage of image data. As part of the image acquisition process, ASIC 44 performs various timing and control functions including control of light source 24, control of color imager 34, and control of external interface 56. It will be apparent to those of ordinary skill in the pertinent art that modifications and variations can be made to processor 40 of the present invention depending on the cost, availability, and performance of off-the-shelf microprocessors, and the type of color imager used. In one embodiment, microprocessor 42 and ASIC 44 are replaced by a single microprocessor 40. In one embodiment, microprocessor 40 is implemented using a single RISC processor. In yet another embodiment, microprocessor 40 is implemented using a RISC and DSP hybrid processor." Ehrhart at 7:32-50. To the extent it is found that Ehrhart does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Ehrhart does not anticipate this claim, Ehrhart renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to

Element	'278 Claim Recitation	Exemplary Citations in Reference
		combine Ehrhart with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1E.i]	the image,	Ehrhart discloses this limitation. *Id.* at [1Pre-E]. "In the Automatic imaging mode, processor 40 is programmed to analyze the color imaging data to determine if an acquired image includes a graphical symbol or is merely a color photographic image. If it makes the determination that the color image includes a graphical symbol, it further analyzes the acquired image and classifies it as a bar code, OCR symbol, text, or a signature. Based on the classification, optical reader 10 jumps to the appropriate routine in EROM 48. The semi-automatic mode is similar. Thus, in the automatic or semi-automatic modes, the bar code scanning mode, the OCR/text mode, the signature capture mode, the color photography mode, and the association mode are controlled by the application program, not by the user." Ehrhart at 8:36-49. "If processor 40 does not locate a bar code symbol it moves onto step 420 and looks for OCR-A or OCR-B characters. If it finds these characters it performs optical character recognition in step 422. If it does not, processor evaluates the image for the presence of text. If text is located, the image is cropped, and the text is compressed and stored in steps 428 and 430. If the image does not include text, processor 40 evaluates the image for the presence of a signature. If one is present, the image is cropped, and the data is compressed and stored in steps 428 and 430. In another embodiment, optical reader 10 is networked, and processor 40 communicates with remote network resources to provide signature verification services. If processor 40 cannot detect a bar code symbol, OCR symbols, text, or a signature, the user is asked

Element	'278 Claim Recitation	Exemplary Citations in Reference
		in step 432 if he desires to store the image. If he does, the color photographic image is stored in memory in step 434." Ehrhart at 10:5-20. "As embodied herein and depicted in FIG. 11, a flow chart showing a method for performing OCR in accordance with yet another embodiment of the present invention is disclosed. In step 1100, reader 10 produces a bit-mapped image of the page. Subsequently, processor 40 finds lines of text in the image, locates the white spaces in each line, and isolates the characters. In step 1108, processor 40 performs character recognition, either OCR-A or OCR-B, as desired. The decoded characters are stored in memory." Ehrhart at 13:1-9. To the extent it is found that Ehrhart does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Ehrhart does not anticipate this claim, Ehrhart renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Ehrhart with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1E.ii]	the symbol information,	Ehrhart discloses this limitation. *Id.* at [1Pre-E]. "In accordance with the invention, the present invention for an optical reader includes a color imaging assembly for acquiring color imaging data. An image analysis circuit determines if the acquired image includes at least one graphical

Element	'278 Claim Recitation	Exemplary Citations in Reference
		symbol. A processing circuit processes the imaging data based on the determination of whether the image includes at least one graphical symbol. The present invention allows a user to read graphical symbols, such as bar codes, text, OCR characters or signatures using a color imager. The color optical reader of the present invention is configured to automatically determine whether a color image includes a graphical symbol, or is merely a color photographic image. The optical reader of the present invention also is operative to associate one acquired image with at least one subsequently acquired image." Ehrhart at 5:36-50. To the extent it is found that Ehrhart does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Ehrhart does not anticipate this claim, Ehrhart renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Ehrhart with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1E.iii]	and a query of a database storing target object information associated with a plurality of target objects including the first target object;	Ehrhart discloses this limitation. *Id.* at [1Pre-E]. "In the Automatic imaging mode, processor 40 is programmed to analyze the color imaging data to determine if an acquired image includes a graphical symbol or is merely a color photographic image. If it makes the determination that the color image includes a graphical symbol, it further analyzes the acquired image and

Element	'278 Claim Recitation	Exemplary Citations in Reference
		classifies it as a bar code, OCR symbol, text, or a signature. Based on the classification, optical reader 10 jumps to the appropriate routine in EROM 48. The semi-automatic mode is similar. Thus, in the automatic or semi-automatic modes, the bar code scanning mode, the OCR/text mode, the signature capture mode, the color photography mode, and the association mode are controlled by the application program, not by the user." Ehrhart at 8:36-49. "In one embodiment, this mode includes a Signature verification program wherein the user may select between a Static verification or a dynamic verification. In the Static mode, the user captures the image of a Signature. The captured image is compared with a reference image Stored in a remote database. In the dynamic mode, optical reader 10 uses the Stylus and Signature block to capture the Signature. In this mode, Signature block 62 measures unique dynamic parameters, Such as applied pressure, direction and timing of movements, or a combination of these parameters. One of ordinary skill in the art will recognize that this list is not meant to be all-inclusive, but rather, is a representative example. The captured dynamic parameters are compared with a reference data Stored in a remote database." Ehrhart at 8:64-9:10. "If processor 40 does not locate a bar code symbol it moves onto step 420 and looks for OCR-A or OCR-B characters. If it finds these characters it performs optical character recognition in step 422. If it does not, processor evaluates the image for the presence of text. If text is located, the image is cropped, and the text is compressed and stored in steps 428 and 430. If the image does not include text, processor 40 evaluates the image for the presence of a signature. If one is present, the image is cropped, and the data is compressed and stored in steps 428 and 430. In another embodiment, optical reader 10 is networked, and processor 40 communicates with remote network resources to provide signature verification services. If processor 40 cannot detect a bar code symbol, OCR symbols, text, or a signature, the user is asked in step 432 if he desires to store the image. If he does, the color photographic image is stored in memory in step 434." Ehrhart at 10:5-20.

Element	'278 Claim Recitation	Exemplary Citations in Reference
		"As embodied herein and depicted in FIG. 11, a flow chart showing a method for performing OCR in accordance with yet another embodiment of the present invention is disclosed. In step 1100, reader 10 produces a bit-mapped image of the page. Subsequently, processor 40 finds lines of text in the image, locates the white spaces in each line, and isolates the characters. In step 1108, processor 40 performs character recognition, either OCR-A or OCR-B, as desired. The decoded characters are stored in memory." Ehrhart at 13:1-9. To the extent it is found that Ehrhart does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Ehrhart does not anticipate this claim, Ehrhart renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Ehrhart with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1F]	receiving, via an address, first target object information associated with the first target object,	Ehrhart discloses receiving, via an address, first target object information associated with the first target object. *Id.* at [1Pre-E]. "LAN 1820 includes server 1822, computer 1824, database 1826, and a plurality of optical readers 10. Database 1826 is used to store associated images along with other data fields. For example, it would be rather useful to store additional information with the associated images shown in FIG. 13. One may want to associate the

Element	'278 Claim Recitation	Exemplary Citations in Reference
		delivery means, route, driver, and other related information for subsequent analysis. Network 1810 allows reader 10, PAN 1850, and wireless system 1400 a way to store such data in database 1826. System analysts can access this information via personal computer 1802 connected to network 1810. In one embodiment, LAN 1820 includes an Internet website. In this embodiment, users are authenticated before gaining access to database 1826." Ehrhart at 16:1-12. To the extent it is found that Ehrhart does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Ehrhart does not anticipate this claim, Ehrhart renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Ehrhart with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1F.i]	wherein the first target object information comprises a response regarding the validity of the document.	Ehrhart discloses the first target object information comprises a response regarding the validity of the document. *Id.* at [1Pre-F]. "The user may also click on OCR/Text icon 660. Clicking icon 660 provides the user with a check validation mode, a text scanning mode, or a bi-tonal image capture mode. The check validation mode is performed in conjunction with network services." Ehrhart at 8:58-62.

Element	'278 Claim Recitation	Exemplary Citations in Reference
		"LAN 1820 includes server 1822, computer 1824, database 1826, and a plurality of optical readers 10. Database 1826 is used to store associated images along with other data fields. For example, it would be rather useful to store additional information with the associated images shown in FIG. 13. One may want to associate the delivery means, route, driver, and other related information for subsequent analysis. Network 1810 allows reader 10, PAN 1850, and wireless system 1400 a way to store such data in database 1826. System analysts can access this information via personal computer 1802 connected to network 1810. In one embodiment, LAN 1820 includes an Internet website. In this embodiment, users are authenticated before gaining access to database 1826." Ehrhart at 16:1-12.

		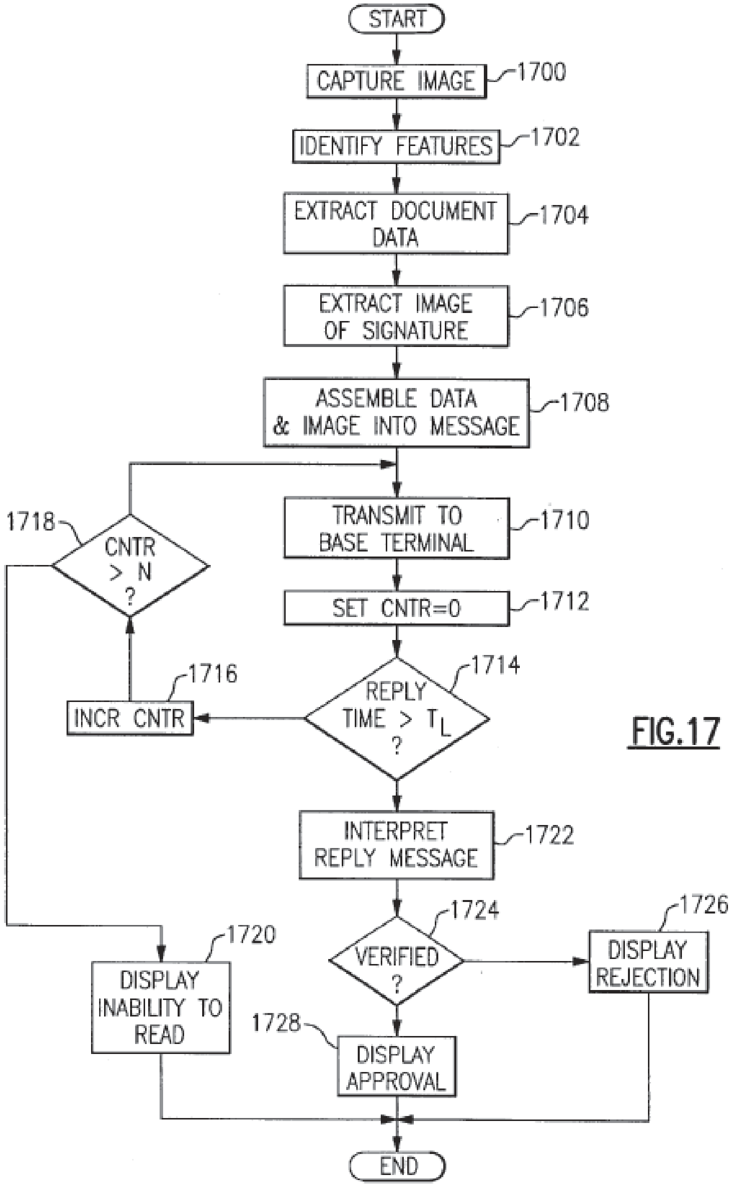START CAPTURE IMAGE 1700 IDENTIFY FEATURES 1702 EXTRACT DOCUMENT DATA 1704 EXTRACT IMAGE OF SIGNATURE 1706 ASSEMBLE DATA & IMAGE INTO MESSAGE 1708 1718 CNTR > N ? TRANSMIT TO BASE TERMINAL 1710 SET CNTR=0 1712 1716 INCR CNTR REPLY TIME > T_L ? 1714 INTERPRET REPLY MESSAGE 1722 1724 1726 1720 VERIFIED ? DISPLAY REJECTION DISPLAY INABILITY TO READ 1728 DISPLAY APPROVAL END FIG.17

Element	'278 Claim Recitation	Exemplary Citations in Reference
		Ehrhart at Fig. 17. To the extent it is found that Ehrhart does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Ehrhart does not anticipate this claim, Ehrhart renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Ehrhart with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

2. Claim 3

Claim	'278 Claim Recitation	Exemplary Citations in Reference
3	The method of claim 1, wherein the symbolic content comprises a unique code.	Ehrhart discloses the method of claim 1, wherein the symbolic content comprises a unique code. *Id.* at [1]. "In accordance with the invention, the present invention for an optical reader includes a color imaging assembly for acquiring color imaging data. An image analysis circuit determines if the acquired image includes at least one graphical symbol. A processing circuit processes the imaging data based on the determination of whether the image includes at least one graphical symbol. The present invention allows a user to read graphical symbols, such as bar codes, text,

		OCR characters or signatures using a color imager. The color optical reader of the present invention is configured to automatically determine whether a color image includes a graphical symbol, or is merely a color photographic image. The optical reader of the present invention also is operative to associate one acquired image with at least one subsequently acquired image." Ehrhart at 5:36-50. To the extent it is found that Ehrhart does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Ehrhart does not anticipate this claim, Ehrhart renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Ehrhart with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

3. Claim 4

Claim	'278 Claim Recitation	Exemplary Citations in Reference
4	The method of claim 3, wherein the image includes a representation of the code.	Ehrhart discloses the method of claim 3, wherein the image includes a representation of the code. *Id.* at [3]. "In accordance with the invention, the present invention for an optical reader includes a color imaging assembly for acquiring color imaging data. An image analysis circuit determines if the acquired image includes at least one graphical

		symbol. A processing circuit processes the imaging data based on the determination of whether the image includes at least one graphical symbol. The present invention allows a user to read graphical symbols, such as bar codes, text, OCR characters or signatures using a color imager. The color optical reader of the present invention is configured to automatically determine whether a color image includes a graphical symbol, or is merely a color photographic image. The optical reader of the present invention also is operative to associate one acquired image with at least one subsequently acquired image." Ehrhart at 5:36-50. To the extent it is found that Ehrhart does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Ehrhart does not anticipate this claim, Ehrhart renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Ehrhart with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

4. Claim 5

Claim	'278 Claim Recitation	Exemplary Citations in Reference
5	The method of claim 3, further comprising detecting whether there is a match within a specified tolerance and returning a message	Ehrhart discloses the method of claim 3, further comprising detecting whether there is a match within a specified tolerance and returning a message indicating the absence of a match within a specified tolerance.

	indicating the absence of a match within a specified tolerance.	"The user may also click on OCR/Text icon 660. Clicking icon 660 provides the user with a check validation mode, a text scanning mode, or a bi-tonal image capture mode. The check validation mode is performed in conjunction with network services." Ehrhart at 8:58-62. "In steps 1712 and 1714, processor 40 initializes a counter and begins waiting for a reply from the host computer. In steps 1714-1718, if the reply is not received within time limit TL, the counter CNTR is incremented and the message is re-transmitted. After several attempts, if CNTR>N (N being an integer), processor 40 outputs a fault message. If the reply message is received within time limit TL, processor interprets the reply in step 1722. If the extracted data and the signature match information stored in the database accessible by the host computer, an approval message is displayed. If the extracted data and the signature do not match information stored in the database accessible by the host computer, a disapproval message is displayed. The dynamic signature verification embodiment is similar to the static embodiment described immediately above. In the dynamic version, the user provides his signature using stylus 18 and signature block 62, as shown in FIG. 1C. Signature block 62 provides processor 40 with the dynamic parameters recorded during signature. The dynamic parameters are transmitted to a host processor, as described above." Ehrhart at 15:31-50.

		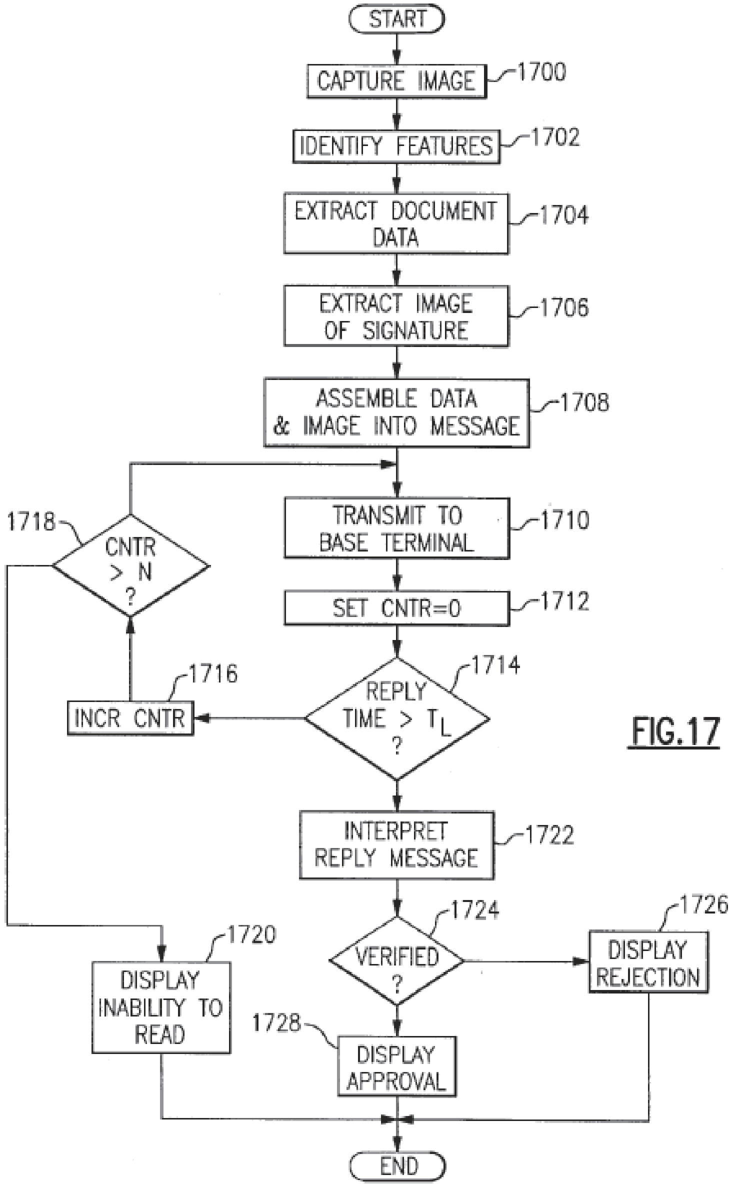START CAPTURE IMAGE 1700 IDENTIFY FEATURES 1702 EXTRACT DOCUMENT DATA 1704 EXTRACT IMAGE OF SIGNATURE 1706 ASSEMBLE DATA & IMAGE INTO MESSAGE 1708 TRANSMIT TO BASE TERMINAL 1710 SET CNTR=0 1712 1718 CNTR > N ? REPLY TIME > T_L ? 1714 INCR CNTR 1716 INTERPRET REPLY MESSAGE 1722 VERIFIED ? 1724 DISPLAY REJECTION 1726 DISPLAY INABILITY TO READ 1720 1728 DISPLAY APPROVAL END FIG.17

		Ehrhart at Fig. 17. To the extent it is found that Ehrhart does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Ehrhart does not anticipate this claim, Ehrhart renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Ehrhart with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

Exhibit F-18

Bank of America Corporation ("BofA" or "Defendant") provides this chart subject to all reservations, objections, statements, and disclaimers set forth herein and in BofA's Final Invalidity Contentions Cover Pleading, as well as any amendment, supplement, or modification thereof, which are incorporated herein by reference in their entirety.

On information and belief, the application that matured into Krouse was filed on June 13, 1997, claiming the priority date of June 13, 1997. Krouse published at least as of August 1, 2000, and is available as prior art at least under 35 U.S.C. § 102(a).

As shown in the chart below, Krouse anticipates asserted claims 1, 4, 5 of the '278 Patent. To the extent Krouse is found not to expressly disclose certain limitations in the asserted claims, such limitations are inherent. To the extent Krouse is found not to anticipate any asserted claims or claim elements of the '278 Patent, the reference nevertheless renders those claims or claim elements obvious under 35 U.S.C. § 103, either alone or in combination with other art identified herein or in the cover pleading of BofA's Invalidity Contentions, for the reasons identified in the cover pleading. These Invalidity Contentions are not an admission by BofA that the accused products, including any current or past versions of the accused products, are covered by, or infringe the asserted claims, but are based instead on the recognition that if the claims are interpreted to be broad enough to encompass the accused products, the claims must also be construed to have that same scope when considering whether they are invalid.

The following chart is based on the Court's claim construction rulings (Dkt. Nos. 145, 153, 236), and to the extent not contradictory, Plaintiffs' apparent interpretations of the claims in Plaintiffs' "Preliminary" Final Infringement Contentions. BofA notes that in many instances, Plaintiffs' "Preliminary" Final Infringement Contentions ignore and/or contradict the Court's claim construction rulings. BofA does not accept the assumptions concerning the scope and meaning implicit in Plaintiffs' Preliminary Final Infringement Contentions, to the extent those assumptions are discernible, and reserves the right to challenge Plaintiffs' proposed (or implied) constructions. Furthermore, Plaintiffs' "Preliminary" Final Contentions fail to comply with the level of specificity required by SPR 2.1.2 and 2.1.3, and include improper reservation of rights to supplement with "additional ways in which BoA Accused Products infringe." Any alleged deficient discovery is of Plaintiffs own making. Most if not all of the deficiencies in Plaintiffs' Preliminary Final Contentions have been documented by Defendants through multiple letters since at least May 18, 2021 and up to Bank of America's most recent letter of August 19, 2023. Plaintiffs' shifting infringement contentions, deficiencies in its infringement contentions, and changing claim scope have prejudiced BofA from being able to reasonably prepare its defenses, including its invalidity contentions. BofA also reserves the right to revise and supplement these charts if and when Plaintiffs are permitted to provide revised or supplemental Infringement Contentions. Citations given in the chart below are merely representative of the respective elements to provide notice to Plaintiffs of BofA's invalidity theories and are not meant to be exhaustive.

1. Claim 1

Element	'278 Claim Recitation	Exemplary Citations in Reference
[1Pre]	A computer-assisted method, comprising:	Krouse discloses a computer-assisted method. Krouse at Abstract • "In accordance with the present invention, financial transaction processing systems and methods are provided. One preferred embodiment of a method according to one aspect of the present invention includes generating an optically scanned image of at least a portion of document containing visual data, in a particular format, representing information related to the financial transaction. Recognition characteristics are generated from the scanned image and are compared to respective sets of reference recognition characteristics generated from respective other transaction documents having different respective formats to determine therefrom whether the particular format of the visual data matches one of the respective formats of the other documents. When such a match is found to exist, location is determined of a field in the scanned image to which optical character recognition may be applied to generate therefrom the information, based upon the respective format found to match the particular format of the visual data. Optical character recognition is then utilized to generate said visual data from said location."
[1A]	receiving, via a mobile device, an image comprising a representation of at least a portion of a document;	Krouse discloses receiving, via a mobile device, an image comprising a representation of at least a portion of a document. *Id.* at [1Pre]. Krouse at Abstract • "In accordance with the present invention, financial transaction processing systems and methods are provided. One preferred embodiment of a method according to one aspect of the present invention

Element	'278 Claim Recitation	Exemplary Citations in Reference
		includes generating an optically scanned image of at least a portion of document containing visual data, in a particular format, representing information related to the financial transaction. Recognition characteristics are generated from the scanned image and are compared to respective sets of reference recognition characteristics generated from respective other transaction documents having different respective formats to determine therefrom whether the particular format of the visual data matches one of the respective formats of the other documents. When such a match is found to exist, location is determined of a field in the scanned image to which optical character recognition may be applied to generate therefrom the information, based upon the respective format found to match the particular format of the visual data. Optical character recognition is then utilized to generate said visual data from said location." Krouse at 3:43–4:12 • "One preferred embodiment of the system of this aspect of the present invention includes an optical scanner for generating a computer-readable scanned image of the document, and an optical character recognizer for generating, from the scanned image, transaction data indicative of at least a portion of the visual information. This embodiment of the system of this aspect of the present invention also includes an authorization record generator for generating, based upon the transaction data, a record of the transaction for being assented to by the party tendering the document prior to completion of the transaction, and for generating from the record once it is assented to by the party, a computer-readable record of assent of the party to the transaction. A collection system is also included in this embodiment of the system of this aspect of the present invention for generating in response to the transaction data and the computer-readable record, an electronic financial transaction request for being used by an electronic transaction system to execute the electronic financial

Element	'278 Claim Recitation	Exemplary Citations in Reference
		transaction. One embodiment of the method of the first aspect of the present invention includes generating a computer-readable scanned image of the document, and generating, from the scanned image, transaction data indicative of at least a portion of the visual information contained on the document. A record is generated, from the transaction data, of the electronic financial transaction being processed, which record is for being assented to by the party tendering the document, prior to completion of said transaction. Once the record has been assented to by the party tendering the document, a computer-readable record of the party's assent to the transaction is generated from the assented-to record. An electronic financial transaction request is then generated in response to the transaction data and the computer-readable assented-to record, which request is for being used by an electronic financial transaction system to execute the electronic financial transaction." Krouse at 4:13–57 • "Also, in a second aspect of the present invention, a system and method are provided for use in processing a financial transaction based at least in part upon visual data formed on a document. The visual data represents information related to the transaction and is in a particular format. One preferred embodiment of the system of the second aspect of the present invention includes an optical scanner for generating a scanned image of at least a portion of the document containing the visual representation. An image characterization generator is provided for generating recognition characteristics from the scanned image, and a recognition characteristic comparitor is provided for comparing the recognition characteristics to respective sets of reference recognition characteristics generated from respective other transaction documents

Element	'278 Claim Recitation	Exemplary Citations in Reference
		having different respective formats and for determining therefrom whether the particular format of the visual data on the scanned document matches one of the respective formats of the other documents. A field location generator is also provided for determining, based upon the one respective format when the particular format is determined to match the one respective format, location of a field in the scanned image to which optical character recognition (OCR) may be applied to generate therefrom the transaction-related information represented by the visual data. OCR is then applied by an image processor to this location to generate the visual data for use in processing the transaction. One preferred embodiment of the method according to the second aspect of the present invention includes generating an optically scanned image of at least a portion of document containing visual data, in a particular format, representing information related to the financial transaction. Recognition characteristics are generated from the scanned image and are compared to respective sets of reference recognition characteristics generated from respective other transaction documents having different respective formats to determine therefrom whether the particular format of the visual data matches one of the respective formats of the other documents. When such a match is found to exist, location is determined of a field in the scanned image to which optical character recognition may be applied to generate therefrom the information, based upon the respective format found to match the particular format of the visual data. Optical character recognition is then utilized to generate said visual data from said location." Krouse at 6:4–14 • "Each of the terminals 18, 20, 22 includes an optical scanner 34 for scanning negotiable instrument documents, (exemplified by check 110 shown in FIG. 6) tendered at point of sale in payment for goods or services

Element	'278 Claim Recitation	Exemplary Citations in Reference
		provided by the merchant or retailer associated with that point of sale, to generate respective computer-readable scanned images of the documents. Preferably, scanner 34 comprises a conventional monotonal image scanner having a resolution of at least 300 dots-per-inch and being adapted for being controlled by local controller/processor 40 to scan one or both sides of a document 110 placed in the scanner 34." To the extent it is found that Krouse does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Krouse does not anticipate this claim, Krouse renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Krouse with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1B]	determining that symbolic content is on the at least the portion of the document based on the image;	Krouse discloses determining that symbolic content is on the at least the portion of the document based on the image. *Id.* at [1Pre-A]. Krouse at Abstract • "In accordance with the present invention, financial transaction processing systems and methods are provided. One preferred embodiment of a method according to one aspect of the present invention includes generating an optically scanned image of at least a portion of document containing visual data, in a particular format, representing information related to the financial transaction. Recognition

Element	'278 Claim Recitation	Exemplary Citations in Reference
		characteristics are generated from the scanned image and are compared to respective sets of reference recognition characteristics generated from respective other transaction documents having different respective formats to determine therefrom whether the particular format of the visual data matches one of the respective formats of the other documents. When such a match is found to exist, location is determined of a field in the scanned image to which optical character recognition may be applied to generate therefrom the information, based upon the respective format found to match the particular format of the visual data. Optical character recognition is then utilized to generate said visual data from said location." Krouse at 7:38–67 • "One such standardized portion in check **110** is Magnetic Ink Character Recognizable (MICR) line **94**, which is positioned at the bottom of the check **110** and in check **110**, includes three fields **114, 116, 118** of printed digits separated by special control characters. Fields **114, 116** contain the financial institution transaction routing number and account number on which the check **110** is drawn, and field **118** contains the check number of the check **110**. It should be noted, however, that other than the left most nine digits (and associated separator characters) of the MICR line **94**, the contents of said line may vary, depending upon the financial institution upon which the negotiable instrument is drawn. System **10** may be adapted to recognize and parse information contained in different types of negotiable instrument MICR lines (not shown). Based upon the standardized location, size, style, etc. of the MICR line **94** in document **110** and the fields **114, 116, 118** contained therein, recognizer **32** utilizes optical character recognition and image processing techniques to extract from the scanned image data the data generated by the scanner **34** from scanning the MICR line portion **94** of the document

Element	'278 Claim Recitation	Exemplary Citations in Reference
		110, and once thus extracted, to generate numeric transaction data indicative of the visual information (i.e., the numbers printed in fields **114, 116, 118**) in MICR line **94**. This numeric data is then transmitted to the processor **40**, which transmits a copy of the numeric data to the storage system **42** where it is stored in association with the scanned image data and the payee and payment amount data previously stored therein." Krouse at 3:43–4:12 • "One preferred embodiment of the system of this aspect of the present invention includes an optical scanner for generating a computer-readable scanned image of the document, and an optical character recognizer for generating, from the scanned image, transaction data indicative of at least a portion of the visual information. This embodiment of the system of this aspect of the present invention also includes an authorization record generator for generating, based upon the transaction data, a record of the transaction for being assented to by the party tendering the document prior to completion of the transaction, and for generating from the record once it is assented to by the party, a computer-readable record of assent of the party to the transaction. A collection system is also included in this embodiment of the system of this aspect of the present invention for generating in response to the transaction data and the computer-readable record, an electronic financial transaction request for being used by an electronic transaction system to execute the electronic financial transaction. One embodiment of the method of the first aspect of the present invention includes generating a computer-readable scanned image of the document, and generating, from the scanned image, transaction data indicative of at least a portion of the visual information contained on the document. A record is generated, from the transaction data, of the electronic financial

Element	'278 Claim Recitation	Exemplary Citations in Reference
		transaction being processed, which record is for being assented to by the party tendering the document, prior to completion of said transaction. Once the record has been assented to by the party tendering the document, a computer-readable record of the party's assent to the transaction is generated from the assented-to record. An electronic financial transaction request is then generated in response to the transaction data and the computer-readable assented-to record, which request is for being used by an electronic financial transaction system to execute the electronic financial transaction." To the extent it is found that Krouse does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Krouse does not anticipate this claim, Krouse renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Krouse with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1C]	extracting symbol information based on the symbolic content according to symbol type;	Krouse discloses extracting symbol information based on the symbolic content according to symbol type. *Id.* at [1Pre-B]. Krouse at Abstract • "In accordance with the present invention, financial transaction processing systems and methods are provided. One preferred embodiment of a method according to one aspect of the present invention

Element	'278 Claim Recitation	Exemplary Citations in Reference
		includes generating an optically scanned image of at least a portion of document containing visual data, in a particular format, representing information related to the financial transaction. Recognition characteristics are generated from the scanned image and are compared to respective sets of reference recognition characteristics generated from respective other transaction documents having different respective formats to determine therefrom whether the particular format of the visual data matches one of the respective formats of the other documents. When such a match is found to exist, location is determined of a field in the scanned image to which optical character recognition may be applied to generate therefrom the information, based upon the respective format found to match the particular format of the visual data. Optical character recognition is then utilized to generate said visual data from said location." Krouse at 7:38–67 • "One such standardized portion in check **110** is Magnetic Ink Character Recognizable (MICR) line **94**, which is positioned at the bottom of the check **110** and in check **110**, includes three fields **114, 116, 118** of printed digits separated by special control characters. Fields **114, 116** contain the financial institution transaction routing number and account number on which the check **110** is drawn, and field **118** contains the check number of the check **110**. It should be noted, however, that other than the left most nine digits (and associated separator characters) of the MICR line **94**, the contents of said line may vary, depending upon the financial institution upon which the negotiable instrument is drawn. System **10** may be adapted to recognize and parse information contained in different types of negotiable instrument MICR lines (not shown). Based upon the standardized location, size, style, etc. of the MICR line **94** in document **110** and the fields **114, 116, 118** contained therein,

Element	'278 Claim Recitation	Exemplary Citations in Reference
		recognizer **32** utilizes optical character recognition and image processing techniques to extract from the scanned image data the data generated by the scanner **34** from scanning the MICR line portion **94** of the document **110**, and once thus extracted, to generate numeric transaction data indicative of the visual information (i.e., the numbers printed in fields **114, 116, 118**) in MICR line **94**. This numeric data is then transmitted to the processor **40**, which transmits a copy of the numeric data to the storage system **42** where it is stored in association with the scanned image data and the payee and payment amount data previously stored therein." To the extent it is found that Krouse does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Krouse does not anticipate this claim, Krouse renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Krouse with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1D]	determining a validity of the document based at least in part on the image and the symbol information; and	Krouse discloses determining a validity of the document based at least in part on the image and the symbol information. *Id.* at [1Pre-C]. Krouse at Abstract • "In accordance with the present invention, financial transaction processing systems and methods are provided. One preferred embodiment of a method

Element	'278 Claim Recitation	Exemplary Citations in Reference
		according to one aspect of the present invention includes generating an optically scanned image of at least a portion of document containing visual data, in a particular format, representing information related to the financial transaction. Recognition characteristics are generated from the scanned image and are compared to respective sets of reference recognition characteristics generated from respective other transaction documents having different respective formats to determine therefrom whether the particular format of the visual data matches one of the respective formats of the other documents. When such a match is found to exist, location is determined of a field in the scanned image to which optical character recognition may be applied to generate therefrom the information, based upon the respective format found to match the particular format of the visual data. Optical character recognition is then utilized to generate said visual data from said location." Krouse at 3:49–57 • "This embodiment of the system of this aspect of the present invention also includes an authorization record generator for generating, based upon the transaction data, a record of the transaction for being assented to by the party tendering the document prior to completion of the transaction, and for generating from the record once it is assented to by the party, a computer-readable record of assent of the party to the transaction. A collection system is also included in this embodiment of the system of this aspect of the present invention for generating in response to the transaction data and the computer-readable record, an electronic financial transaction request for being used by an electronic transaction system to execute the electronic financial transaction." Krouse at 8:1–12 • "A second copy of the numeric data is also transmitted from the processor **40** to the authorization record generator **38**, which also retrieves from the

Element	'278 Claim Recitation	Exemplary Citations in Reference
		storage system **42** the payee and payment amount data associated with this transaction and transmits same to the generator **38**. Processor **40** also supplies to the generator 38 from the storage system **42** data indicative of the initiation of the current transaction. Preferably, each of the terminals **18, 20, 22** of the system **10** is associated with a unique station/terminal identification code or number which is stored in the terminal's respective storage system **42**. This terminal identification number is also supplied to the generator **38** from the system **42**." To the extent it is found that Krouse does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Krouse does not anticipate this claim, Krouse renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Krouse with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1E]	recognizing the document as a first target object based at least in part on	Krouse discloses this limitation. *Id.* at [1Pre-D]. Krouse at Abstract • "In accordance with the present invention, financial transaction processing systems and methods are provided. One preferred embodiment of a method according to one aspect of the present invention includes generating an optically scanned image of at least a portion of document containing visual data, in a particular format, representing

Element	'278 Claim Recitation	Exemplary Citations in Reference
		information related to the financial transaction. Recognition characteristics are generated from the scanned image and are compared to respective sets of reference recognition characteristics generated from respective other transaction documents having different respective formats to determine therefrom whether the particular format of the visual data matches one of the respective formats of the other documents. When such a match is found to exist, location is determined of a field in the scanned image to which optical character recognition may be applied to generate therefrom the information, based upon the respective format found to match the particular format of the visual data. Optical character recognition is then utilized to generate said visual data from said location." Krouse at 3:43–4:12 • "One preferred embodiment of the system of this aspect of the present invention includes an optical scanner for generating a computer-readable scanned image of the document, and an optical character recognizer for generating, from the scanned image, transaction data indicative of at least a portion of the visual information. This embodiment of the system of this aspect of the present invention also includes an authorization record generator for generating, based upon the transaction data, a record of the transaction for being assented to by the party tendering the document prior to completion of the transaction, and for generating from the record once it is assented to by the party, a computer-readable record of assent of the party to the transaction. A collection system is also included in this embodiment of the system of this aspect of the present invention for generating in response to the transaction data and the computer-readable record, an electronic financial transaction request for being used by an electronic transaction system to execute the electronic financial transaction.

Element	'278 Claim Recitation	Exemplary Citations in Reference
		One embodiment of the method of the first aspect of the present invention includes generating a computer-readable scanned image of the document, and generating, from the scanned image, transaction data indicative of at least a portion of the visual information contained on the document. A record is generated, from the transaction data, of the electronic financial transaction being processed, which record is for being assented to by the party tendering the document, prior to completion of said transaction. Once the record has been assented to by the party tendering the document, a computer-readable record of the party's assent to the transaction is generated from the assented-to record. An electronic financial transaction request is then generated in response to the transaction data and the computer-readable assented-to record, which request is for being used by an electronic financial transaction system to execute the electronic financial transaction." Krouse at 4:13–57 • "Also, in a second aspect of the present invention, a system and method are provided for use in processing a financial transaction based at least in part upon visual data formed on a document. The visual data represents information related to the transaction and is in a particular format. One preferred embodiment of the system of the second aspect of the present invention includes an optical scanner for generating a scanned image of at least a portion of the document containing the visual representation. An image characterization generator is provided for generating recognition characteristics from the scanned image, and a recognition characteristic comparitor is provided for comparing the recognition characteristics to respective sets of reference recognition characteristics generated from respective other transaction documents having different respective formats and for determining therefrom whether the particular format of the visual data on the scanned document

Element	'278 Claim Recitation	Exemplary Citations in Reference
		matches one of the respective formats of the other documents. A field location generator is also provided for determining, based upon the one respective format when the particular format is determined to match the one respective format, location of a field in the scanned image to which optical character recognition (OCR) may be applied to generate therefrom the transaction-related information represented by the visual data. OCR is then applied by an image processor to this location to generate the visual data for use in processing the transaction. One preferred embodiment of the method according to the second aspect of the present invention includes generating an optically scanned image of at least a portion of document containing visual data, in a particular format, representing information related to the financial transaction. Recognition characteristics are generated from the scanned image and are compared to respective sets of reference recognition characteristics generated from respective other transaction documents having different respective formats to determine therefrom whether the particular format of the visual data matches one of the respective formats of the other documents. When such a match is found to exist, location is determined of a field in the scanned image to which optical character recognition may be applied to generate therefrom the information, based upon the respective format found to match the particular format of the visual data. Optical character recognition is then utilized to generate said visual data from said location." To the extent it is found that Krouse does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Krouse does not anticipate this claim, Krouse renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been

Element	'278 Claim Recitation	Exemplary Citations in Reference
		motivated to combine Krouse with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1E.i]	the image,	Krouse discloses this limitation. *Id.* at [1Pre-E]. Krouse at 1:6–26 • "The present invention relates generally to financial transaction processing systems and methods, and more specifically, to such systems and methods wherein at least one document containing financial transaction-related information is optically scanned to generate at least one computer-readable image from which the information may be extracted for use in processing the transaction. Particular utility for the present invention is found in the area of processing electronic financial transactions (e.g., electronic funds transfer (EFT) and Automatic Clearing House (ACH) transactions), although other utilities are also contemplated for the present invention, including other financial transaction processing and accounting applications." Krouse at 3:43–4:12 • "One preferred embodiment of the system of this aspect of the present invention includes an optical scanner for generating a computer-readable scanned image of the document, and an optical character recognizer for generating, from the scanned image, transaction data indicative of at least a portion of the visual information. This embodiment of the system of this aspect of the present invention also includes an authorization record generator for generating, based upon the transaction data, a record of the

Element	'278 Claim Recitation	Exemplary Citations in Reference
		transaction for being assented to by the party tendering the document prior to completion of the transaction, and for generating from the record once it is assented to by the party, a computer-readable record of assent of the party to the transaction. A collection system is also included in this embodiment of the system of this aspect of the present invention for generating in response to the transaction data and the computer-readable record, an electronic financial transaction request for being used by an electronic transaction system to execute the electronic financial transaction. One embodiment of the method of the first aspect of the present invention includes generating a computer-readable scanned image of the document, and generating, from the scanned image, transaction data indicative of at least a portion of the visual information contained on the document. A record is generated, from the transaction data, of the electronic financial transaction being processed, which record is for being assented to by the party tendering the document, prior to completion of said transaction. Once the record has been assented to by the party tendering the document, a computer-readable record of the party's assent to the transaction is generated from the assented-to record. An electronic financial transaction request is then generated in response to the transaction data and the computer-readable assented-to record, which request is for being used by an electronic financial transaction system to execute the electronic financial transaction." Krouse at 4:13–57 • "Also, in a second aspect of the present invention, a system and method are provided for use in processing a financial transaction based at least in part upon visual data formed on a document. The visual data represents information related to the transaction and is in a particular format.

Element	'278 Claim Recitation	Exemplary Citations in Reference
		One preferred embodiment of the system of the second aspect of the present invention includes an optical scanner for generating a scanned image of at least a portion of the document containing the visual representation. An image characterization generator is provided for generating recognition characteristics from the scanned image, and a recognition characteristic comparitor is provided for comparing the recognition characteristics to respective sets of reference recognition characteristics generated from respective other transaction documents having different respective formats and for determining therefrom whether the particular format of the visual data on the scanned document matches one of the respective formats of the other documents. A field location generator is also provided for determining, based upon the one respective format when the particular format is determined to match the one respective format, location of a field in the scanned image to which optical character recognition (OCR) may be applied to generate therefrom the transaction-related information represented by the visual data. OCR is then applied by an image processor to this location to generate the visual data for use in processing the transaction. One preferred embodiment of the method according to the second aspect of the present invention includes generating an optically scanned image of at least a portion of document containing visual data, in a particular format, representing information related to the financial transaction. Recognition characteristics are generated from the scanned image and are compared to respective sets of reference recognition characteristics generated from respective other transaction documents having different respective formats to determine therefrom whether the particular format of the visual data matches one of the respective formats of the other documents. When such a match is found to exist, location is determined of a field in the scanned image to which optical character recognition may be applied to generate therefrom the information, based upon the

Element	'278 Claim Recitation	Exemplary Citations in Reference
		respective format found to match the particular format of the visual data. Optical character recognition is then utilized to generate said visual data from said location." Krouse at 7:38–67 - "One such standardized portion in check **110** is Magnetic Ink Character Recognizable (MICR) line **94**, which is positioned at the bottom of the check **110** and in check **110**, includes three fields **114, 116, 118** of printed digits separated by special control characters. Fields **114, 116** contain the financial institution transaction routing number and account number on which the check **110** is drawn, and field **118** contains the check number of the check **110**. It should be noted, however, that other than the left most nine digits (and associated separator characters) of the MICR line **94**, the contents of said line may vary, depending upon the financial institution upon which the negotiable instrument is drawn. System **10** may be adapted to recognize and parse information contained in different types of negotiable instrument MICR lines (not shown). Based upon the standardized location, size, style, etc. of the MICR line **94** in document **110** and the fields **114, 116, 118** contained therein, recognizer **32** utilizes optical character recognition and image processing techniques to extract from the scanned image data the data generated by the scanner **34** from scanning the MICR line portion **94** of the document **110**, and once thus extracted, to generate numeric transaction data indicative of the visual information (i.e., the numbers printed in fields **114, 116, 118**) in MICR line **94**. This numeric data is then transmitted to the processor **40**, which transmits a copy of the numeric data to the storage system **42** where it is stored in association with the scanned image data and the payee and payment amount data previously stored therein."

Element	'278 Claim Recitation	Exemplary Citations in Reference
		To the extent it is found that Krouse does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Krouse does not anticipate this claim, Krouse renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Krouse with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1E.ii]	the symbol information,	Krouse discloses this limitation. *Id.* at [1Pre-E]. Krouse at 7:38–67 • "One such standardized portion in check **110** is Magnetic Ink Character Recognizable (MICR) line **94**, which is positioned at the bottom of the check **110** and in check **110**, includes three fields **114, 116, 118** of printed digits separated by special control characters. Fields **114, 116** contain the financial institution transaction routing number and account number on which the check **110** is drawn, and field **118** contains the check number of the check **110**. It should be noted, however, that other than the left most nine digits (and associated separator characters) of the MICR line **94**, the contents of said line may vary, depending upon the financial institution upon which the negotiable instrument is drawn. System **10** may be adapted to recognize and parse information contained in different types of negotiable instrument MICR lines (not shown). Based upon the standardized location, size, style, etc. of the MICR line

Element	'278 Claim Recitation	Exemplary Citations in Reference
		94 in document **110** and the fields **114, 116, 118** contained therein, recognizer **32** utilizes optical character recognition and image processing techniques to extract from the scanned image data the data generated by the scanner **34** from scanning the MICR line portion **94** of the document **110**, and once thus extracted, to generate numeric transaction data indicative of the visual information (i.e., the numbers printed in fields **114, 116, 118**) in MICR line **94**. This numeric data is then transmitted to the processor **40**, which transmits a copy of the numeric data to the storage system **42** where it is stored in association with the scanned image data and the payee and payment amount data previously stored therein." To the extent it is found that Krouse does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Krouse does not anticipate this claim, Krouse renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Krouse with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1E.iii]	and a query of a database storing target object information associated with a plurality of target objects including the first target object;	Krouse discloses this limitation. *Id.* at [1Pre-E]. Krouse at Abstract • "In accordance with the present invention, financial transaction processing systems and methods are provided. One preferred embodiment of a method

Element	'278 Claim Recitation	Exemplary Citations in Reference
		according to one aspect of the present invention includes generating an optically scanned image of at least a portion of document containing visual data, in a particular format, representing information related to the financial transaction. Recognition characteristics are generated from the scanned image and are compared to respective sets of reference recognition characteristics generated from respective other transaction documents having different respective formats to determine therefrom whether the particular format of the visual data matches one of the respective formats of the other documents. When such a match is found to exist, location is determined of a field in the scanned image to which optical character recognition may be applied to generate therefrom the information, based upon the respective format found to match the particular format of the visual data. Optical character recognition is then utilized to generate said visual data from said location." Krouse at 4:13–57 • "Also, in a second aspect of the present invention, a system and method are provided for use in processing a financial transaction based at least in part upon visual data formed on a document. The visual data represents information related to the transaction and is in a particular format. One preferred embodiment of the system of the second aspect of the present invention includes an optical scanner for generating a scanned image of at least a portion of the document containing the visual representation. An image characterization generator is provided for generating recognition characteristics from the scanned image, and a recognition characteristic comparitor is provided for comparing the recognition characteristics to respective sets of reference recognition characteristics generated from respective other transaction documents having different respective formats and for determining therefrom whether the particular format of the visual data on the scanned document

Element	'278 Claim Recitation	Exemplary Citations in Reference
		matches one of the respective formats of the other documents. A field location generator is also provided for determining, based upon the one respective format when the particular format is determined to match the one respective format, location of a field in the scanned image to which optical character recognition (OCR) may be applied to generate therefrom the transaction-related information represented by the visual data. OCR is then applied by an image processor to this location to generate the visual data for use in processing the transaction. One preferred embodiment of the method according to the second aspect of the present invention includes generating an optically scanned image of at least a portion of document containing visual data, in a particular format, representing information related to the financial transaction. Recognition characteristics are generated from the scanned image and are compared to respective sets of reference recognition characteristics generated from respective other transaction documents having different respective formats to determine therefrom whether the particular format of the visual data matches one of the respective formats of the other documents. When such a match is found to exist, location is determined of a field in the scanned image to which optical character recognition may be applied to generate therefrom the information, based upon the respective format found to match the particular format of the visual data. Optical character recognition is then utilized to generate said visual data from said location." Krouse at 14:64–15:67 • "Processor **220** compares the recognition characteristics of the particular document **400** being processed with the sets of reference recognition characteristics stored in the archive **228**. More specifically, the vertex coordinates and dimensions of the imaginary rectangles generated by the generator **204** from the scanned image of the document **400** are compared

Element	'278 Claim Recitation	Exemplary Citations in Reference
		with respective reference recognition characteristics (i.e., vertex coordinates and dimensions of imaginary rectangles) generated by the process used to generate the recognition characteristics from the scanned image of document **400**, to determine whether a match exists between the recognition characteristics from the scanned image of document **400** and any one of the sets of reference recognition characteristics stored in the database archive system **228**. Preferably, this is accomplished by first comparing the recognition characteristics from the scanned image of document **400** to the sets of reference characteristics, and unless the respective dimensions of at least three of the imaginary rectangles from the scanned image of document **400** match to within a predetermined error tolerance (e.g., 2 percent) of corresponding respective vertices and dimensions in a respective set of reference recognition characteristics stored in the system **228**, processor **220** signals that the format of the scanned document 400 does not match any of the formats of the reference documents from which the reference recognition characteristics stored in the database archive **228** were generated. If such a "no match" condition is determined by the processor **220** to exist, the bill document **400** may be sent to the reference characterization generator station to permit the human operator to determine whether to input the reference characteristics and other bill document identification information (i.e., the other types of information stored in association with each set of reference recognition characteristics previously stored in the database archive **228**) manually for storage in the archive system **228** as a new set of reference recognition characteristics to permit other bills having the same format as document **400** to be processable by terminal **202** in the future, or to process the bill payment transaction as a "one-of-a-kind" transaction (i.e., without storing such recognition characteristics and information in the database archive **228** for future use by the processor **220**). Otherwise, if the predetermined match error tolerance conditions are found to exist, processor **220** selects those sets of reference recognition characteristics that are found to satisfy such

Element	'278 Claim Recitation	Exemplary Citations in Reference
		conditions, and calculates the displacement differences that exist between each of the imaginary rectangles of the scanned image of the document **400** and those described in these selected sets of reference recognition characteristics (i.e., the displacements that must be applied to the imaginary rectangles of the scanned image of document **400** to accomplish a "best fit" match of same with the corresponding imaginary rectangles described in these selected sets of reference recognition characteristics). Such "best fit" match conditions are determined to exist for given corresponding rectangles of the recognition characteristics of the scanned image of the document **400** and a given respective set of reference recognition characteristics, when the absolute pixel displacement differences between the respective abscissa and ordinate coordinates of the vertices of the rectangles are within 40 pixels and 60 pixels of each other, respectively, and the respective lengths and widths of the rectangles are within 20 pixels and 2 pixels of each other, respectively. Once these absolute displacement differences have been calculated by the processor **220**, processor **220** "scores" the degree of best fit match condition that exists between the imaginary rectangles of the scanned image of the document **400** and those described in the aforesaid selected sets of reference recognition characteristics, whereby to determine which one of the aforesaid selected sets of reference recognition characteristics best matches the recognition characteristics of the scanned document **400**." To the extent it is found that Krouse does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Krouse does not anticipate this claim, Krouse renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Krouse with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the

Element	'278 Claim Recitation	Exemplary Citations in Reference
		reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1F]	receiving, via an address, first target object information associated with the first target object,	Krouse discloses receiving, via an address, first target object information associated with the first target object. *Id.* at [1Pre-E]. Krouse at Abstract • "In accordance with the present invention, financial transaction processing systems and methods are provided. One preferred embodiment of a method according to one aspect of the present invention includes generating an optically scanned image of at least a portion of document containing visual data, in a particular format, representing information related to the financial transaction. Recognition characteristics are generated from the scanned image and are compared to respective sets of reference recognition characteristics generated from respective other transaction documents having different respective formats to determine therefrom whether the particular format of the visual data matches one of the respective formats of the other documents. When such a match is found to exist, location is determined of a field in the scanned image to which optical character recognition may be applied to generate therefrom the information, based upon the respective format found to match the particular format of the visual data. Optical character recognition is then utilized to generate said visual data from said location." Krouse at 3:43–4:12 • "One preferred embodiment of the system of this aspect of the present invention includes an optical scanner for generating a computer-readable

Element	'278 Claim Recitation	Exemplary Citations in Reference
		scanned image of the document, and an optical character recognizer for generating, from the scanned image, transaction data indicative of at least a portion of the visual information. This embodiment of the system of this aspect of the present invention also includes an authorization record generator for generating, based upon the transaction data, a record of the transaction for being assented to by the party tendering the document prior to completion of the transaction, and for generating from the record once it is assented to by the party, a computer-readable record of assent of the party to the transaction. A collection system is also included in this embodiment of the system of this aspect of the present invention for generating in response to the transaction data and the computer-readable record, an electronic financial transaction request for being used by an electronic transaction system to execute the electronic financial transaction. One embodiment of the method of the first aspect of the present invention includes generating a computer-readable scanned image of the document, and generating, from the scanned image, transaction data indicative of at least a portion of the visual information contained on the document. A record is generated, from the transaction data, of the electronic financial transaction being processed, which record is for being assented to by the party tendering the document, prior to completion of said transaction. Once the record has been assented to by the party tendering the document, a computer-readable record of the party's assent to the transaction is generated from the assented-to record. An electronic financial transaction request is then generated in response to the transaction data and the computer-readable assented-to record, which request is for being used by an electronic financial transaction system to execute the electronic financial transaction." Krouse at 4:13–57

Element	'278 Claim Recitation	Exemplary Citations in Reference
		- "Also, in a second aspect of the present invention, a system and method are provided for use in processing a financial transaction based at least in part upon visual data formed on a document. The visual data represents information related to the transaction and is in a particular format. One preferred embodiment of the system of the second aspect of the present invention includes an optical scanner for generating a scanned image of at least a portion of the document containing the visual representation. An image characterization generator is provided for generating recognition characteristics from the scanned image, and a recognition characteristic comparitor is provided for comparing the recognition characteristics to respective sets of reference recognition characteristics generated from respective other transaction documents having different respective formats and for determining therefrom whether the particular format of the visual data on the scanned document matches one of the respective formats of the other documents. A field location generator is also provided for determining, based upon the one respective format when the particular format is determined to match the one respective format, location of a field in the scanned image to which optical character recognition (OCR) may be applied to generate therefrom the transaction-related information represented by the visual data. OCR is then applied by an image processor to this location to generate the visual data for use in processing the transaction. One preferred embodiment of the method according to the second aspect of the present invention includes generating an optically scanned image of at least a portion of document containing visual data, in a particular format, representing information related to the financial transaction. Recognition characteristics are generated from the scanned image and are compared to respective sets of reference recognition characteristics generated from respective other transaction documents having different

Element	**'278 Claim Recitation**	**Exemplary Citations in Reference**
		respective formats to determine therefrom whether the particular format of the visual data matches one of the respective formats of the other documents. When such a match is found to exist, location is determined of a field in the scanned image to which optical character recognition may be applied to generate therefrom the information, based upon the respective format found to match the particular format of the visual data. Optical character recognition is then utilized to generate said visual data from said location." To the extent it is found that Krouse does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Krouse does not anticipate this claim, Krouse renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Krouse with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1F.i]	wherein the first target object information comprises a response regarding the validity of the document.	Krouse discloses the first target object information comprises a response regarding the validity of the document. *Id.* at [1Pre-E]. Krouse at Abstract • "In accordance with the present invention, financial transaction processing systems and methods are provided. One preferred embodiment of a method according to one aspect of the present invention includes generating an optically scanned image of at least a portion of document containing visual

Element	'278 Claim Recitation	Exemplary Citations in Reference
		data, in a particular format, representing information related to the financial transaction. Recognition characteristics are generated from the scanned image and are compared to respective sets of reference recognition characteristics generated from respective other transaction documents having different respective formats to determine therefrom whether the particular format of the visual data matches one of the respective formats of the other documents. When such a match is found to exist, location is determined of a field in the scanned image to which optical character recognition may be applied to generate therefrom the information, based upon the respective format found to match the particular format of the visual data. Optical character recognition is then utilized to generate said visual data from said location." Krouse at 3:49–57 • "This embodiment of the system of this aspect of the present invention also includes an authorization record generator for generating, based upon the transaction data, a record of the transaction for being assented to by the party tendering the document prior to completion of the transaction, and for generating from the record once it is assented to by the party, a computer-readable record of assent of the party to the transaction. A collection system is also included in this embodiment of the system of this aspect of the present invention for generating in response to the transaction data and the computer-readable record, an electronic financial transaction request for being used by an electronic transaction system to execute the electronic financial transaction." Krouse at 8:1–12 • "A second copy of the numeric data is also transmitted from the processor **40** to the authorization record generator **38**, which also retrieves from the storage system **42** the payee and payment amount data associated with this transaction and transmits same to the generator **38**. Processor **40** also

Element	'278 Claim Recitation	Exemplary Citations in Reference
		supplies to the generator 38 from the storage system **42** data indicative of the initiation of the current transaction. Preferably, each of the terminals **18, 20, 22** of the system **10** is associated with a unique station/terminal identification code or number which is stored in the terminal's respective storage system **42**. This terminal identification number is also supplied to the generator **38** from the system **42**." Krouse at 7:38–67 • "One such standardized portion in check **110** is Magnetic Ink Character Recognizable (MICR) line **94**, which is positioned at the bottom of the check **110** and in check **110**, includes three fields **114, 116, 118** of printed digits separated by special control characters. Fields **114, 116** contain the financial institution transaction routing number and account number on which the check **110** is drawn, and field **118** contains the check number of the check **110**. It should be noted, however, that other than the left most nine digits (and associated separator characters) of the MICR line **94**, the contents of said line may vary, depending upon the financial institution upon which the negotiable instrument is drawn. System **10** may be adapted to recognize and parse information contained in different types of negotiable instrument MICR lines (not shown). Based upon the standardized location, size, style, etc. of the MICR line **94** in document **110** and the fields **114, 116, 118** contained therein, recognizer **32** utilizes optical character recognition and image processing techniques to extract from the scanned image data the data generated by the scanner **34** from scanning the MICR line portion **94** of the document **110**, and once thus extracted, to generate numeric transaction data indicative of the visual information (i.e., the numbers printed in fields **114, 116, 118**) in MICR line **94**. This numeric data is then transmitted to the processor **40**, which transmits a copy of the numeric data to the storage system **42** where it is stored in association with the scanned image data and the payee and payment amount data previously stored

Element	**'278 Claim Recitation**	**Exemplary Citations in Reference**
		therein." To the extent it is found that Krouse does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Krouse does not anticipate this claim, Krouse renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Krouse with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

2. Claim 3

Claim	'278 Claim Recitation	Exemplary Citations in Reference
3	The method of claim 1, wherein the symbolic content comprises a unique code.	Krouse discloses the method of claim 1, wherein the symbolic content comprises a unique code. *Id.* at [1]. Krouse at Fig. 6 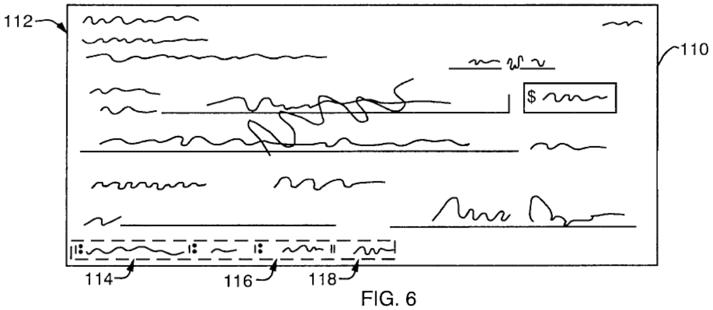 FIG. 6 Krouse at 7:38–67 - "One such standardized portion in check **110** is Magnetic Ink Character Recognizable (MICR) line **94**, which is positioned at the bottom of the check **110** and in check **110**, includes three fields **114, 116, 118** of printed digits separated by special control characters. Fields **114, 116** contain the financial institution transaction routing number and account number on which the check **110** is drawn, and field **118** contains the check number of the check **110**. It should be noted, however, that other than the left most nine digits (and associated separator characters) of the MICR line **94**, the contents of said line may vary, depending upon the financial institution upon which the negotiable instrument is drawn. System **10** may be adapted to recognize and parse information contained in different types of negotiable instrument MICR lines (not

		shown). Based upon the standardized location, size, style, etc. of the MICR line **94** in document **110** and the fields **114, 116, 118** contained therein, recognizer **32** utilizes optical character recognition and image processing techniques to extract from the scanned image data the data generated by the scanner **34** from scanning the MICR line portion **94** of the document **110**, and once thus extracted, to generate numeric transaction data indicative of the visual information (i.e., the numbers printed in fields **114, 116, 118**) in MICR line **94**. This numeric data is then transmitted to the processor **40**, which transmits a copy of the numeric data to the storage system **42** where it is stored in association with the scanned image data and the payee and payment amount data previously stored therein." To the extent it is found that Krouse does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Krouse does not anticipate this claim, Krouse renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Krouse with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

3. Claim 4

Claim	'278 Claim Recitation	Exemplary Citations in Reference

4	The method of claim 3, wherein the image includes a representation of the code.	Krouse discloses the method of claim 3, wherein the image includes a representation of the code. *Id.* at [3]. Krouse at Fig. 6 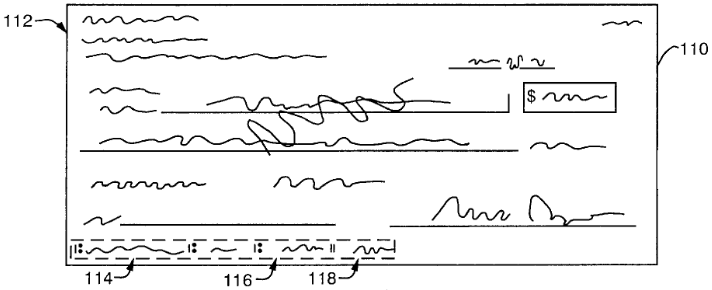 FIG. 6 Krouse at 7:38–67 - "One such standardized portion in check **110** is Magnetic Ink Character Recognizable (MICR) line **94**, which is positioned at the bottom of the check **110** and in check **110**, includes three fields **114, 116, 118** of printed digits separated by special control characters. Fields **114, 116** contain the financial institution transaction routing number and account number on which the check **110** is drawn, and field **118** contains the check number of the check **110**. It should be noted, however, that other than the left most nine digits (and associated separator characters) of the MICR line **94**, the contents of said line may vary, depending upon the financial institution upon which the negotiable instrument is drawn. System **10** may be adapted to recognize and parse information contained in different types of negotiable instrument MICR lines (not shown). Based upon the standardized location, size, style, etc. of the MICR line

		94 in document **110** and the fields **114, 116, 118** contained therein, recognizer **32** utilizes optical character recognition and image processing techniques to extract from the scanned image data the data generated by the scanner **34** from scanning the MICR line portion **94** of the document **110**, and once thus extracted, to generate numeric transaction data indicative of the visual information (i.e., the numbers printed in fields **114, 116, 118**) in MICR line **94**. This numeric data is then transmitted to the processor **40**, which transmits a copy of the numeric data to the storage system **42** where it is stored in association with the scanned image data and the payee and payment amount data previously stored therein." To the extent it is found that Krouse does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Krouse does not anticipate this claim, Krouse renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Krouse with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

4. Claim 5

Claim	'278 Claim Recitation	Exemplary Citations in Reference
5	The method of claim 3, further comprising detecting whether there is a match within a specified tolerance and returning a message	Krouse discloses the method of claim 3, further comprising detecting whether there is a match within a specified tolerance and returning a message indicating the absence of a match within a specified tolerance.

	indicating the absence of a match within a specified tolerance.	*Id.* at [3]. Krouse at Abstract • "In accordance with the present invention, financial transaction processing systems and methods are provided. One preferred embodiment of a method according to one aspect of the present invention includes generating an optically scanned image of at least a portion of document containing visual data, in a particular format, representing information related to the financial transaction. Recognition characteristics are generated from the scanned image and are compared to respective sets of reference recognition characteristics generated from respective other transaction documents having different respective formats to determine therefrom whether the particular format of the visual data matches one of the respective formats of the other documents. When such a match is found to exist, location is determined of a field in the scanned image to which optical character recognition may be applied to generate therefrom the information, based upon the respective format found to match the particular format of the visual data. Optical character recognition is then utilized to generate said visual data from said location." Krouse at 4:13–57 • "Also, in a second aspect of the present invention, a system and method are provided for use in processing a financial transaction based at least in part upon visual data formed on a document. The visual data represents information related to the transaction and is in a particular format. One preferred embodiment of the system of the second aspect of the present invention includes an optical scanner for generating a scanned image of at least a portion of the document containing the visual representation. An image characterization generator is provided for

		generating recognition characteristics from the scanned image, and a recognition characteristic comparitor is provided for comparing the recognition characteristics to respective sets of reference recognition characteristics generated from respective other transaction documents having different respective formats and for determining therefrom whether the particular format of the visual data on the scanned document matches one of the respective formats of the other documents. A field location generator is also provided for determining, based upon the one respective format when the particular format is determined to match the one respective format, location of a field in the scanned image to which optical character recognition (OCR) may be applied to generate therefrom the transaction-related information represented by the visual data. OCR is then applied by an image processor to this location to generate the visual data for use in processing the transaction. One preferred embodiment of the method according to the second aspect of the present invention includes generating an optically scanned image of at least a portion of document containing visual data, in a particular format, representing information related to the financial transaction. Recognition characteristics are generated from the scanned image and are compared to respective sets of reference recognition characteristics generated from respective other transaction documents having different respective formats to determine therefrom whether the particular format of the visual data matches one of the respective formats of the other documents. When such a match is found to exist, location is determined of a field in the scanned image to which optical character recognition may be applied to generate therefrom the information, based upon the respective format found to match the particular format of the visual data. Optical character recognition is then utilized to generate said visual data from said location."

		Krouse at 14:64–15:67 • “Processor **220** compares the recognition characteristics of the particular document **400** being processed with the sets of reference recognition characteristics stored in the archive **228**. More specifically, the vertex coordinates and dimensions of the imaginary rectangles generated by the generator **204** from the scanned image of the document **400** are compared with respective reference recognition characteristics (i.e., vertex coordinates and dimensions of imaginary rectangles) generated by the process used to generate the recognition characteristics from the scanned image of document **400**, to determine whether a match exists between the recognition characteristics from the scanned image of document **400** and any one of the sets of reference recognition characteristics stored in the database archive system **228**. Preferably, this is accomplished by first comparing the recognition characteristics from the scanned image of document **400** to the sets of reference characteristics, and unless the respective dimensions of at least three of the imaginary rectangles from the scanned image of document **400** match to within a predetermined error tolerance (e.g., 2 percent) of corresponding respective vertices and dimensions in a respective set of reference recognition characteristics stored in the system **228**, processor **220** signals that the format of the scanned document 400 does not match any of the formats of the reference documents from which the reference recognition characteristics stored in the database archive **228** were generated. If such a "no match" condition is determined by the processor **220** to exist, the bill document **400** may be sent to the reference characterization generator station to permit the human operator to determine whether to input the reference characteristics and other bill document identification information (i.e., the other types of information stored in association with each set of reference recognition characteristics previously stored in the database archive **228**) manually for storage in the archive system **228** as a new set of reference recognition characteristics to permit other bills having the same format as document **400** to be processable by terminal **202** in the future, or to

		process the bill payment transaction as a "one-of-a-kind" transaction (i.e., without storing such recognition characteristics and information in the database archive **228** for future use by the processor **220**). Otherwise, if the predetermined match error tolerance conditions are found to exist, processor **220** selects those sets of reference recognition characteristics that are found to satisfy such conditions, and calculates the displacement differences that exist between each of the imaginary rectangles of the scanned image of the document **400** and those described in these selected sets of reference recognition characteristics (i.e., the displacements that must be applied to the imaginary rectangles of the scanned image of document **400** to accomplish a "best fit" match of same with the corresponding imaginary rectangles described in these selected sets of reference recognition characteristics). Such "best fit" match conditions are determined to exist for given corresponding rectangles of the recognition characteristics of the scanned image of the document **400** and a given respective set of reference recognition characteristics, when the absolute pixel displacement differences between the respective abscissa and ordinate coordinates of the vertices of the rectangles are within 40 pixels and 60 pixels of each other, respectively, and the respective lengths and widths of the rectangles are within 20 pixels and 2 pixels of each other, respectively. Once these absolute displacement differences have been calculated by the processor **220**, processor **220** "scores" the degree of best fit match condition that exists between the imaginary rectangles of the scanned image of the document **400** and those described in the aforesaid selected sets of reference recognition characteristics, whereby to determine which one of the aforesaid selected sets of reference recognition characteristics best matches the recognition characteristics of the scanned document **400**." To the extent it is found that Krouse does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Krouse does not anticipate this claim, Krouse renders it obvious, either alone, in combination with

		the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Krouse with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

Nantworks, LLC v. Bank of America Corporation
Central District of California, No. 2:20-cv-7872
Invalidity Chart for U.S. Patent 9,031,278 based on the Query by Image and Video Content ("QBIC") System

Exhibit F-23

Bank of America Corporation ("BofA" or "Defendant") provides this chart subject to all reservations, objections, statements, and disclaimers set forth herein and in BofA's Final Invalidity Contentions Cover Pleading, as well as any amendment, supplement, or modification thereof, which are incorporated herein by reference in their entirety. A statement of reasons and good cause for each supplement in this chart is identified in the cover pleading.

QBIC was known or used by others in the United States and publicly used or on sale in the United States by November 5, 2000. *See generally* W. Niblack et al., The QBIC Project: Querying Images By Content Using Color, Texture, and Shape, SPIE Vol. 1908 (1993), pp. 173–187 (IBM0002390–404); M. Flickner et al., Query by Image and Video Content: The QBIC System ("QBIC"), IEEE Sept. 1995, pp. 23–32 (IBM0000893–902); W. Niblack et al, Updates to the QBIC System, SPIE Vol. 3312 (1997), pp. 150–161 (IBM 0002419–430); and Query By Image Content QBIC Demonstration Program, IBM Research, Sept. 1998 (IBM000747). QBIC is available as prior art at least under 35 U.S.C. § 102(a), (b), and (g), having been known or used by others in the United States and publicly used or on sale in the United States by at least 1995 and certainly by November 5, 2000, including, for example, at http://libra.uc.davis.edu through the Art History Department at U.C. Davis; at www..thinker.org/imagebase/indox-2 .htm3 through the Fine Arts Museum of San Francisco; through demos at http://wwwqbic.almaden.ibm.com; and through IBM's commercial products, including in IBM's Ultimedia Manager, the DB2 series of products such as the Image Extender in DB2 Universal Database, and the Enterprise Information Portal. *See generally* Deposition of IBM Corporate Representative Myron Flickner, *Nantworks, LLC and Nant Holdings IP, LLC v. Bank of America Corporation and Bank of America, N.A.*, 2:20-cv-07872-GW-PVC (C.D.C.A. 2020) ("Flickner 8/29/23 Depo") at e.g., 32–33, 39–40, 44, 62–63, 66–67,72–74, 103–107, 110–112, 127, 163–174, 183–189, 191–195, 215, 219–220, 241, 307–309, 311–317, Exs. 3–13; *see also* Deposition of IBM Corporate Representative Myron Flickner, *Network-1 Technologies, Inc., v. Google LLC and Youtube LLC*, 1:14-cv-09558-PGG (S.D.N.Y. 2019) ("Flickner Depo"), including, e.g., at 13, 14, 15–18, 29–30, 34, 36, 37, 43–46, 48, 58, 60, 61–62, 67–69, 72, 74–76, 84–86, 88–89, 98–100, 102, 106, 113, 125, 135, 152–154, 160–161, 170–172, 180–181, 194, 198–201, 203–205, 215–216, 220–221, and Exs. 1–16.

As shown in the chart below, QBIC anticipates asserted claims 1, 4, and 5 of the '278 Patent. The disclosures/exemplary citations for each element incorporate the disclosures/exemplary citations of the proceeding limitations. To the extent QBIC is found not to expressly disclose certain limitations in the asserted claims, such limitations are inherent. To the extent QBIC is found not to anticipate any asserted claims or claim elements of the '278 Patent, the reference nevertheless renders those claims or claim elements obvious under 35 U.S.C. § 103, either alone or in combination with other art identified herein or in the cover pleading of BofA's Final Invalidity Contentions, for the reasons identified in the cited references and/or in the cover pleading. These Invalidity Contentions are not an admission by BofA that the accused products, including any current or past versions of the accused products, are covered by, or infringe the asserted claims, but are based instead on the recognition that if the claims are interpreted to be broad enough to encompass the accused products, the claims must also be construed to have that same scope when considering whether they are invalid. The following chart incorporates by reference the analysis from Exhibit F-21 for each limitation.

The following chart is based on the Court's claim construction rulings (Dkt. Nos. 145, 153, 236), and to the extent not contradictory, Plaintiffs' apparent interpretations of the claims in Plaintiffs' "Preliminary" Final Infringement Contentions. BofA notes that in many instances, Plaintiffs' "Preliminary" Final Infringement Contentions ignore and/or contradict the Court's claim construction rulings. BofA does not accept the assumptions concerning the scope and meaning implicit in Plaintiffs' Preliminary Final Infringement Contentions, to the extent those assumptions are discernible, and reserves the right to challenge Plaintiffs' proposed (or implied) constructions. Furthermore, Plaintiffs' "Preliminary" Final Contentions fail to comply with the level of specificity required by SPR 2.1.2 and 2.1.3, and include improper reservation of rights to supplement with "additional ways in which BoA Accused Products infringe." Any alleged deficient discovery is of Plaintiffs own making. Most if not all of the deficiencies in Plaintiffs' Preliminary Final Contentions have been documented by Defendants through multiple letters since at least May 18, 2021 and up to Bank of America's most recent letter of August 19, 2023. Plaintiffs' shifting infringement contentions, deficiencies in its infringement contentions, and changing claim scope have prejudiced BofA from being able to reasonably prepare its defenses, including its invalidity contentions. BofA also reserves the right to revise and supplement these charts if and when Plaintiffs are permitted to provide revised or supplemental Infringement Contentions. Citations given in the chart below are merely representative of the respective elements to provide notice to Plaintiffs of BofA's invalidity theories and are not meant to be exhaustive.

1. Claim 1

Element	'278 Claim Recitation	Exemplary Citations in Reference
[1Pre]	A computer-assisted method, comprising:	QBIC discloses a computer-assisted method. The QBIC Project: Querying Images by Content Using Color, Texture, and Shape (IBM 0002390–404) 2390 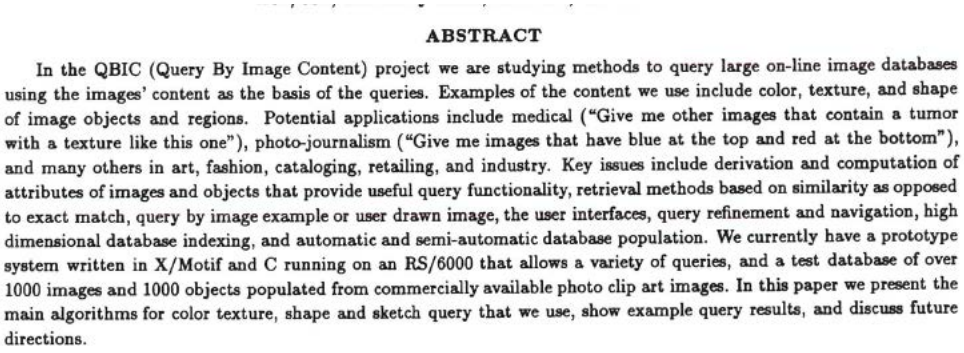 **ABSTRACT** In the QBIC (Query By Image Content) project we are studying methods to query large on-line image databases using the images' content as the basis of the queries. Examples of the content we use include color, texture, and shape of image objects and regions. Potential applications include medical ("Give me other images that contain a tumor with a texture like this one"), photo-journalism ("Give me images that have blue at the top and red at the bottom"), and many others in art, fashion, cataloging, retailing, and industry. Key issues include derivation and computation of attributes of images and objects that provide useful query functionality, retrieval methods based on similarity as opposed to exact match, query by image example or user drawn image, the user interfaces, query refinement and navigation, high dimensional database indexing, and automatic and semi-automatic database population. We currently have a prototype system written in X/Motif and C running on an RS/6000 that allows a variety of queries, and a test database of over 1000 images and 1000 objects populated from commercially available photo clip art images. In this paper we present the main algorithms for color texture, shape and sketch query that we use, show example query results, and discuss future directions. QBIC Demo (IBM000747)  Windows 98 - Console - VMware ESXi - Google Chrome Not secure | https://172.16.85.32/ui/#/console/5 Netscape - [Message: 200] Location: http://172.16.85.176:3456/ **QBIC(TM) -- IBM's Query By Image Content** Goto: How to Use QBIC IBM Image Search Demo Stop QBIC Http Server QBIC allows you to make queries of image databases large and small based on visual characteristics -- properties such as color percentages, color layout and textures occurring in the images. The queries match image colors, textures and their positions without having to describe them with in words. QBIC's content-based queries can also be combined with keywords to provide a powerful multimedia search and retrieval tool. As image collections grow larger and become more prolific, QBIC is the best tool to more efficiently search and navigate through them. In this demo, we use QBIC to index about 470 color images based on their color and texture information. You can evaluate the QBIC Search capability by following the link located at left called IBM Image Search Demo. For license and business related questins, please contact Ted Loewenberg by phone 408-927-1202, or by email QBIC Demo (IBM000747)

Element	'278 Claim Recitation	Exemplary Citations in Reference
		 Flickner Depo at 13 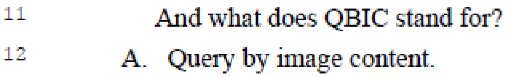 And what does QBIC stand for? A. Query by image content. Flickner Depo at 84

Element	'278 Claim Recitation	Exemplary Citations in Reference
		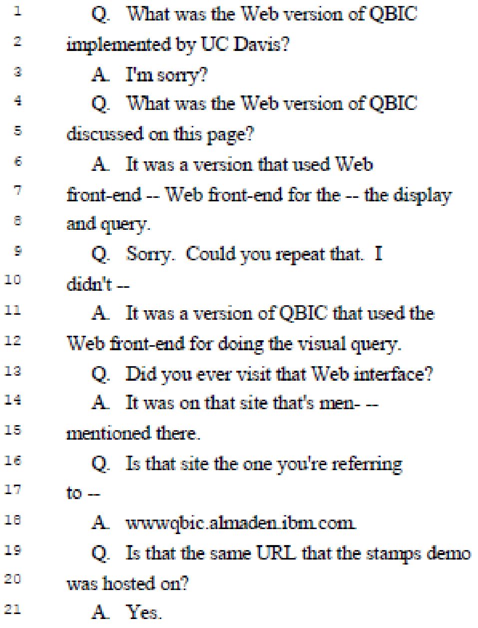 Q. What was the Web version of QBIC implemented by UC Davis? A. I'm sorry? Q. What was the Web version of QBIC discussed on this page? A. It was a version that used Web front-end -- Web front-end for the -- the display and query. Q. Sorry. Could you repeat that. I didn't -- A. It was a version of QBIC that used the Web front-end for doing the visual query. Q. Did you ever visit that Web interface? A. It was on that site that's men- -- mentioned there. Q. Is that site the one you're referring to -- A. wwwqbic.almaden.ibm.com. Q. Is that the same URL that the stamps demo was hosted on? A. Yes. Flickner Depo at, e.g., 24, 41-42, 62-63 (user could take a picture or video and submit it to the QBIC system).

Element	'278 Claim Recitation	Exemplary Citations in Reference
		Q. And you mentioned that you would extract features from the images? A. We did. Q. How would you do that? A. We ran algorithms on the pixels of the image and from that we created a feature. Q. What sort of features would you extract from the images? A. A QBIC system did color, shape and texture. Q. Okay. So let's take color. How would you, say, extract a color feature from an image? A. We typically would compute -- compute a color histogram. Q. And what's a color histogram? A. It's a measure of -- it's an estimate of the probability density of a color being in an image. Q. So you would take an image and extract that color histogram from it? A. Yup.

		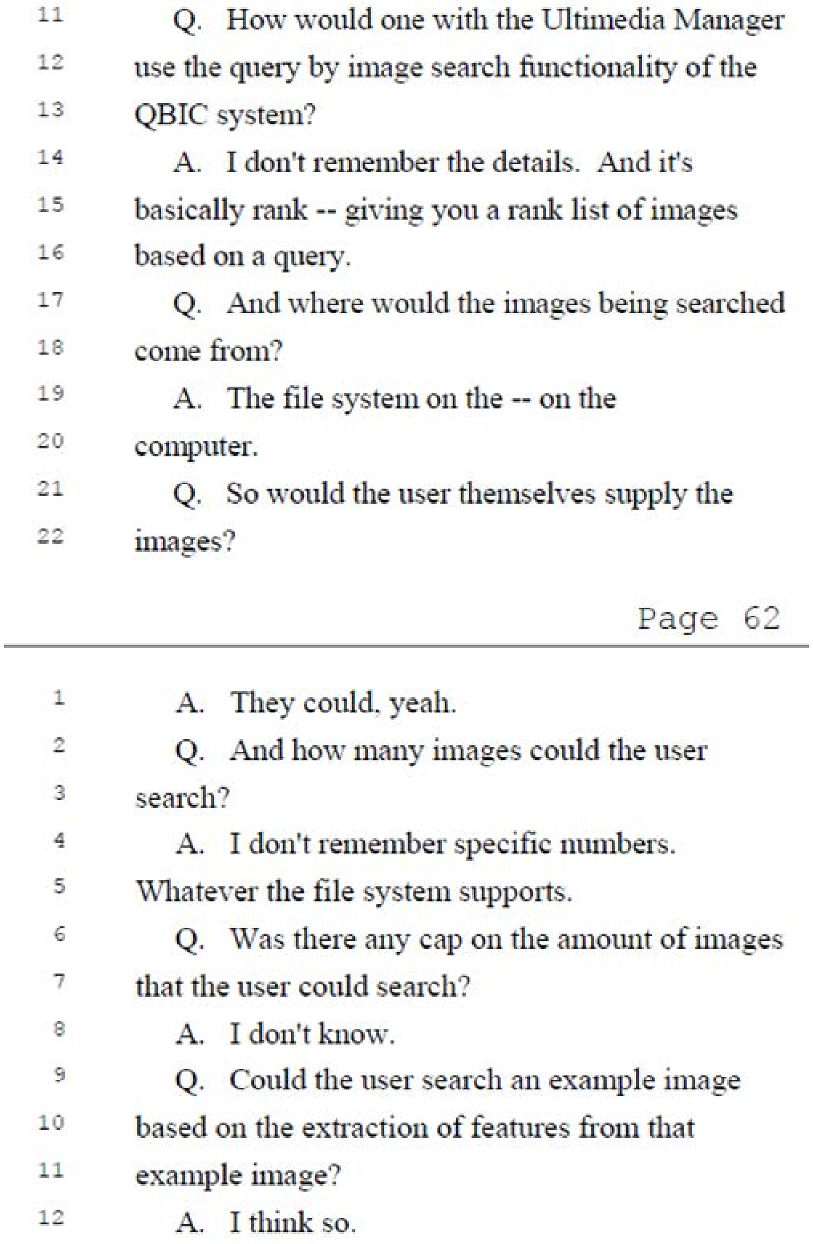 Q. How would one with the Ultimedia Manager use the query by image search functionality of the QBIC system? A. I don't remember the details. And it's basically rank -- giving you a rank list of images based on a query. Q. And where would the images being searched come from? A. The file system on the -- on the computer. Q. So would the user themselves supply the images? Page 62 A. They could, yeah. Q. And how many images could the user search? A. I don't remember specific numbers. Whatever the file system supports. Q. Was there any cap on the amount of images that the user could search? A. I don't know. Q. Could the user search an example image based on the extraction of features from that example image? A. I think so.

Element	'278 Claim Recitation	Exemplary Citations in Reference
		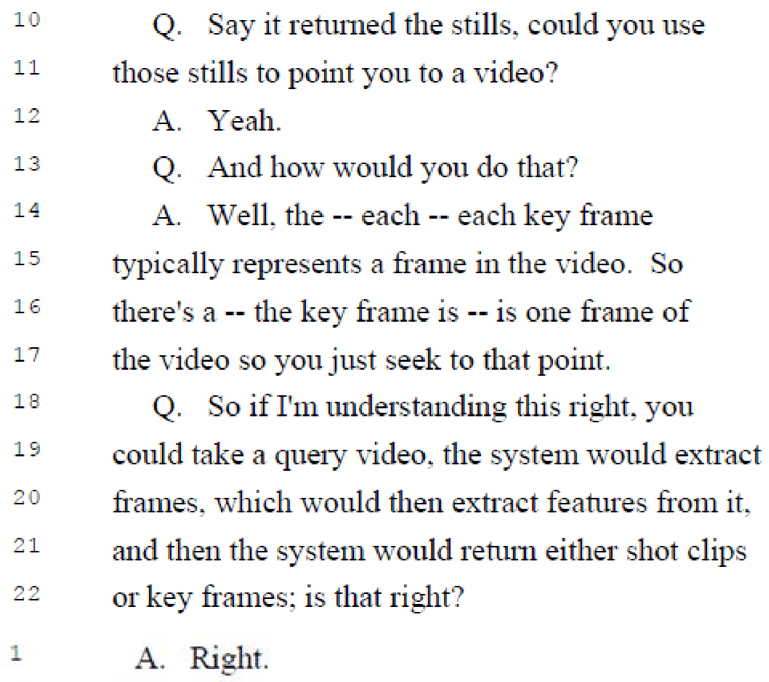 Q. Say it returned the stills, could you use those stills to point you to a video? A. Yeah. Q. And how would you do that? A. Well, the -- each -- each key frame typically represents a frame in the video. So there's a -- the key frame is -- is one frame of the video so you just seek to that point. Q. So if I'm understanding this right, you could take a query video, the system would extract frames, which would then extract features from it, and then the system would return either shot clips or key frames; is that right? A. Right. Flickner 8/29/23 Depo at 31

Element	'278 Claim Recitation	Exemplary Citations in Reference
		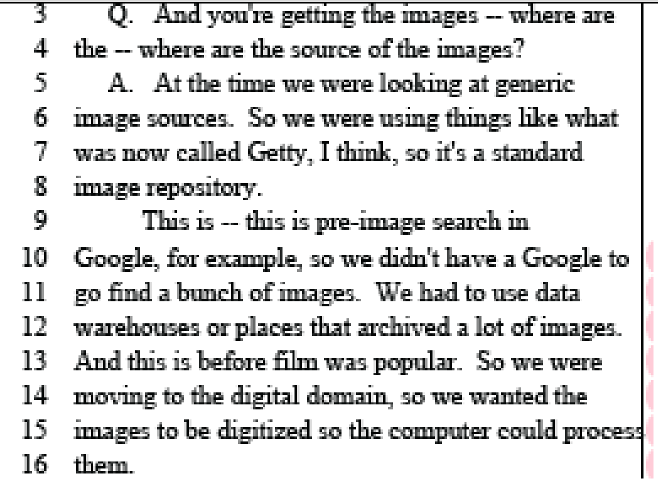 Q. And you're getting the images -- where are the -- where are the source of the images? A. At the time we were looking at generic image sources. So we were using things like what was now called Getty, I think, so it's a standard image repository. This is -- this is pre-image search in Google, for example, so we didn't have a Google to go find a bunch of images. We had to use data warehouses or places that archived a lot of images. And this is before film was popular. So we were moving to the digital domain, so we wanted the images to be digitized so the computer could process them. *Id.* at 42 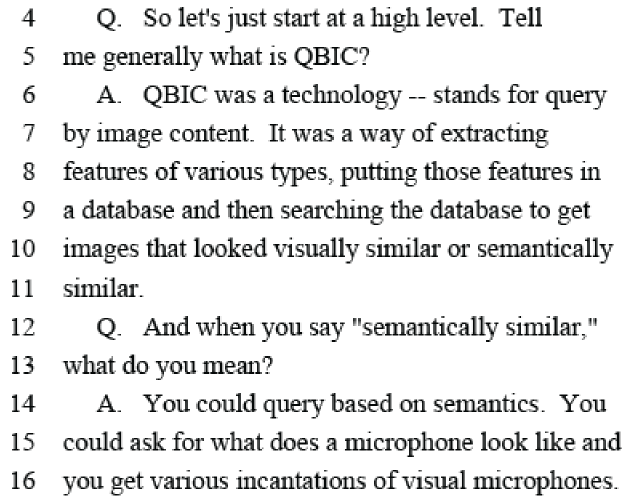 Q. So let's just start at a high level. Tell me generally what is QBIC? A. QBIC was a technology -- stands for query by image content. It was a way of extracting features of various types, putting those features in a database and then searching the database to get images that looked visually similar or semantically similar. Q. And when you say "semantically similar," what do you mean? A. You could query based on semantics. You could ask for what does a microphone look like and you get various incantations of visual microphones. *Id.* at 205–206

Element	'278 Claim Recitation	Exemplary Citations in Reference
		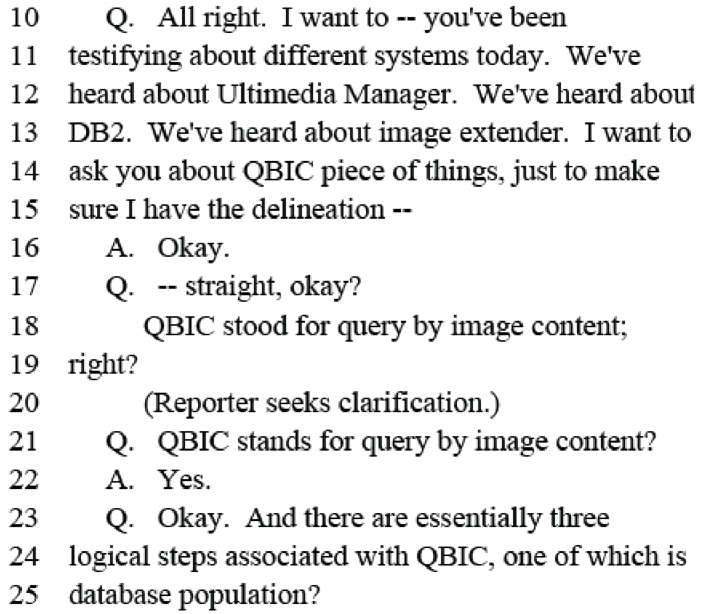 Q. All right. I want to -- you've been testifying about different systems today. We've heard about Ultimedia Manager. We've heard about DB2. We've heard about image extender. I want to ask you about QBIC piece of things, just to make sure I have the delineation -- A. Okay. Q. -- straight, okay? QBIC stood for query by image content; right? (Reporter seeks clarification.) Q. QBIC stands for query by image content? A. Yes. Q. Okay. And there are essentially three logical steps associated with QBIC, one of which is database population? 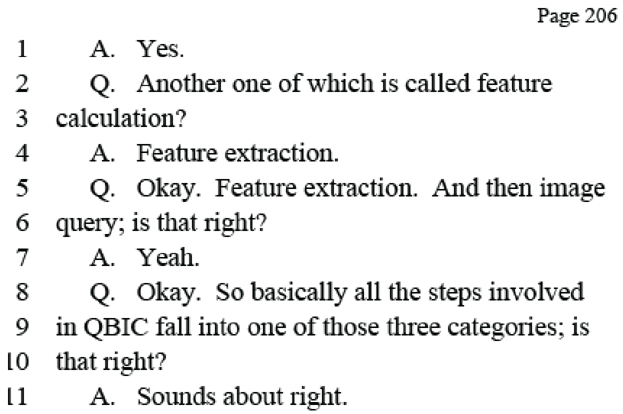 Page 206 A. Yes. Q. Another one of which is called feature calculation? A. Feature extraction. Q. Okay. Feature extraction. And then image query; is that right? A. Yeah. Q. Okay. So basically all the steps involved in QBIC fall into one of those three categories; is that right? A. Sounds about right. The QBIC system was implemented in IBM commercial products, including in

Element	’278 Claim Recitation	Exemplary Citations in Reference
		IBM’s Ultimedia Manager, the DB2 series of products such as the Image Extender in DB2 Universal Database, and the Enterprise Information Portal. All of the features referenced in this chart were incorporated into these commercial implementations before November 6, 2000, as reflected in a combination of scientific papers, IBM’s internal documentation, and the QBIC demo software applications. *See, e.g.*, Flickner 8/29/23 Depo Tr. at 33:9-24, 71:21-72:10, 164:21-166:2, 166:11-175:24, 178:5-186:17, 186:21-192:4, 192:14-194:1, 195:16-196:4, 194:1-308:1; IBM 000001-617; IBM 000747; IBM 000777-880; IBM 000893-902; IBM 0002315-320; IBM 0002390-404; IBM 0002418-430.
[1A]	receiving, via a mobile device, an image comprising a representation of at least a portion of a document;	QBIC discloses receiving an image comprising a representation of at least a portion of a document. The QBIC Project: Querying Images by Content Using Color, Texture, and Shape (IBM 0002390–404) at 2396 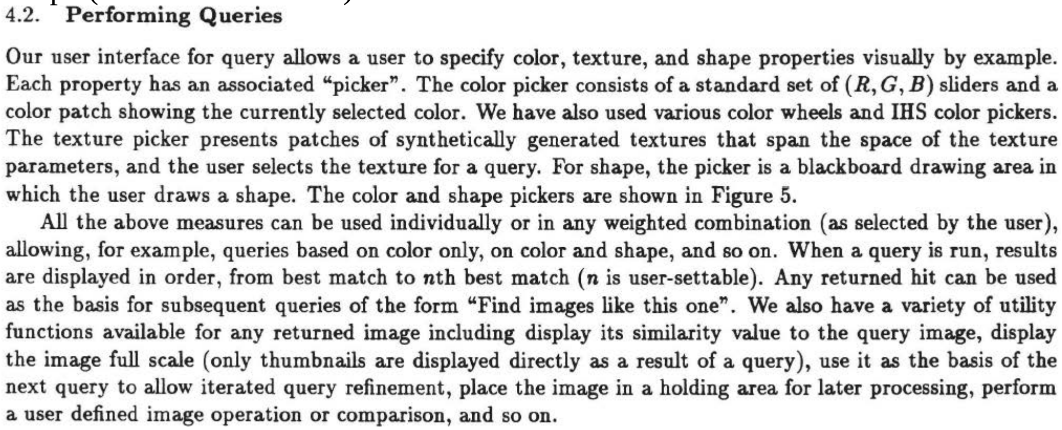 4.2. **Performing Queries** Our user interface for query allows a user to specify color, texture, and shape properties visually by example. Each property has an associated “picker”. The color picker consists of a standard set of (R, G, B) sliders and a color patch showing the currently selected color. We have also used various color wheels and IHS color pickers. The texture picker presents patches of synthetically generated textures that span the space of the texture parameters, and the user selects the texture for a query. For shape, the picker is a blackboard drawing area in which the user draws a shape. The color and shape pickers are shown in Figure 5. All the above measures can be used individually or in any weighted combination (as selected by the user), allowing, for example, queries based on color only, on color and shape, and so on. When a query is run, results are displayed in order, from best match to nth best match (n is user-settable). Any returned hit can be used as the basis for subsequent queries of the form “Find images like this one”. We also have a variety of utility functions available for any returned image including display its similarity value to the query image, display the image full scale (only thumbnails are displayed directly as a result of a query), use it as the basis of the next query to allow iterated query refinement, place the image in a holding area for later processing, perform a user defined image operation or comparison, and so on. IBM0002390 at 2391-92:

Element	'278 Claim Recitation	Exemplary Citations in Reference
		2.1. Database Population The first step in population is to simply load the images into the system. This involves adding the image to the database, preparing a reduced 100x100 "thumbnail", and adding any available text information to the database. Object/element identification is an optional part of this step. It allows a user to manually or semi-automatically identify objects, areas, or significant edges in a scene using a mouse. Internally, each identified object or area becomes a binary mask. There can be an arbitrary number of objects per image, objects may overlap, and objects may consist of multiple disconnected components (e.g. the set of dots on a polka dot dress). Text (e.g. "baby on beach") can be added to an outlined object, or to the scene as a whole. Ideally this step would be done automatically but current automatic methods to identify and outline image objects are not sufficiently robust so we allow this step to be done manually or semi-automatically. Because database population is the most labor intensive task in many QBIC applications, we try to automate it as much as possible. We call these semi-automatic methods interactive outlining or shrink-wrapping. *Interactive Outlining:* In this operation, the user provides some starting information which the system uses, along with image analysis methods, to compute an object outline. An example is shown in Figure 1. One technique, described in these proceedings [6], uses "snakes" or active contours. The user enters an approximate object outline, and the method iteratively adjusts this outline using an algorithm that minimizes an expression containing two main terms: total image edge strength along the outline, and the outline curvature. Another method, flood fill, can take as starting information a single mouse-clicked image point. Starting from this object point, all adjacent pixels are included in the object whose color is within a specified color distance from the original point. The perimeter of all included points becomes the object boundary. For reasonably uniform objects that are distinct from the background, this operation provides a fast object outliner. In interactive outlining, as in other parts of our system, we have attempted to build the functions in such a way as to carefully distribute the tasks between the human and machine. For example, the tasks of giving the starting outline for an object or of attaching semantic meaning to objects is left to the user since they are difficult for machines. On the other hand, computation of quantitative image/object features is ideally suited for machines and so is done automatically.  Figure 1: Original image, approximate user outline, and automatically refined outline of sun. Query by Image and Video Content: The QBIC System (IBM 0000893–902): "We have developed the QBIC (Query by Image Content) system to explore

Element	'278 Claim Recitation	Exemplary Citations in Reference
		content-based retrieval methods. QBIC allows queries on large image and video databases based on • example images, • user-constructed sketches and drawings, • selected color and texture patterns, • camera and object motion, and • other graphical information." *Id*. at pp. 1–3. "INTEGRATION WITH TEXT AND PARAMETRIC ANNOTATION. Query by visual content complements and extends existing query methods. Systems must be able to integrate queries combining date, subject matter, price, and availability with content properties such as color, texture, and shape." *Id*. at p. 8. Image, Audio, and Video Extenders Administration and Programming (IBM 000001–617): "DB2 (DB2) Universal Database (UDB) is a powerful, object-relational database manager. It stores and protects traditional numeric and character data, as well as large, complex objects (LOBs). DB2 extenders help you exploit DB2 's object-relational features. The extenders define distinct data types and special functions for image, audio, video, and text objects. … Because of that, the extenders give your applications a single point of access to any or all of these types of data, along with traditional numeric and character data." *Id*. at p. 3. Managing Enterprise Information Portal (IBM 000777–880): "Many paper-intensive enterprises, such as insurance companies and financial institutions, administer large volumes of business-related content. The need for an enterprise solution for managing and accessing business information spans many industries. A content server is a software system that stores multimedia, business forms, documents, and related data, along with metadata that allows employees to process and work with the content." *Id*. at p. 3. "The image search server uses IBM's QBIC (query by image content) technology to help you search for objects by certain visual properties, such as

Element	'278 Claim Recitation	Exemplary Citations in Reference
		color and texture. The image search server analyzes images and stores the image information in a database. Then users can run image queries, which use the visual properties of images, to match colors, textures, and their positions without describing them in words. You can combine content-based queries with text and keyword searches for more powerful retrieval of image and multimedia data." *Id.* at p. 63. Flickner Depo at 88 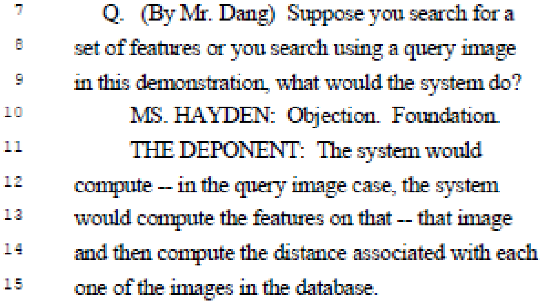 Q. (By Mr. Dang) Suppose you search for a set of features or you search using a query image in this demonstration, what would the system do? MS. HAYDEN: Objection. Foundation. THE DEPONENT: The system would compute -- in the query image case, the system would compute the features on that -- that image and then compute the distance associated with each one of the images in the database. Flickner Depo at 89-90

Element	'278 Claim Recitation	Exemplary Citations in Reference
		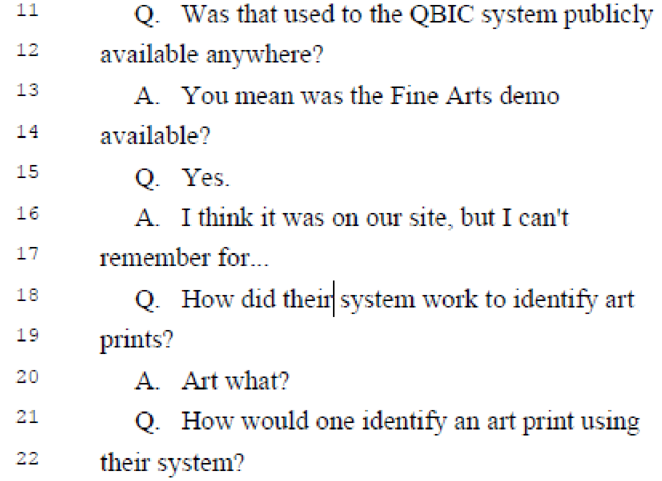 Q. Was that used to the QBIC system publicly available anywhere? A. You mean was the Fine Arts demo available? Q. Yes. A. I think it was on our site, but I can't remember for... Q. How did their system work to identify art prints? A. Art what? Q. How would one identify an art print using their system? 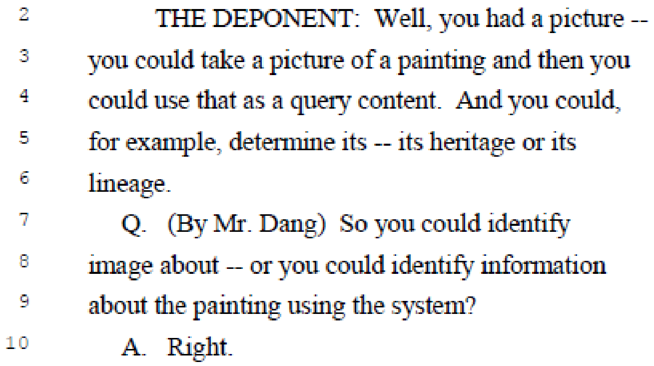 THE DEPONENT: Well, you had a picture -- you could take a picture of a painting and then you could use that as a query content. And you could, for example, determine its -- its heritage or its lineage. Q. (By Mr. Dang) So you could identify image about -- or you could identify information about the painting using the system? A. Right. Flickner Depo at 39-40 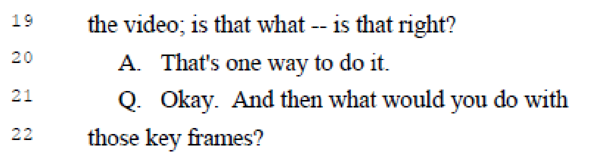 Q. Okay. So you would take -- you would use a histogram methodology to extract key frames from the video; is that what -- is that right? A. That's one way to do it. Q. Okay. And then what would you do with those key frames?

Element	'278 Claim Recitation	Exemplary Citations in Reference
		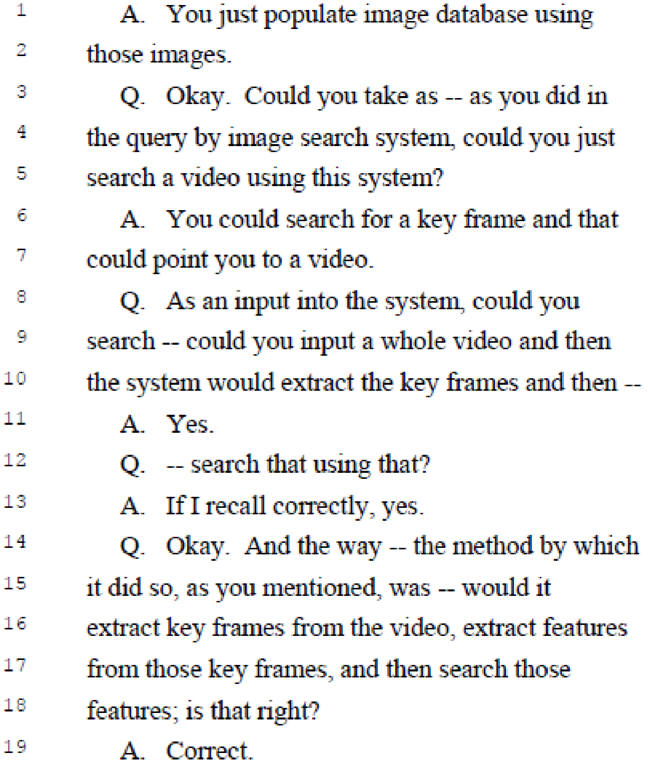 A. You just populate image database using those images. Q. Okay. Could you take as -- as you did in the query by image search system, could you just search a video using this system? A. You could search for a key frame and that could point you to a video. Q. As an input into the system, could you search -- could you input a whole video and then the system would extract the key frames and then -- A. Yes. Q. -- search that using that? A. If I recall correctly, yes. Q. Okay. And the way -- the method by which it did so, as you mentioned, was -- would it extract key frames from the video, extract features from those key frames, and then search those features; is that right? A. Correct. Flickner Depo at 54-55

Element	'278 Claim Recitation	Exemplary Citations in Reference
		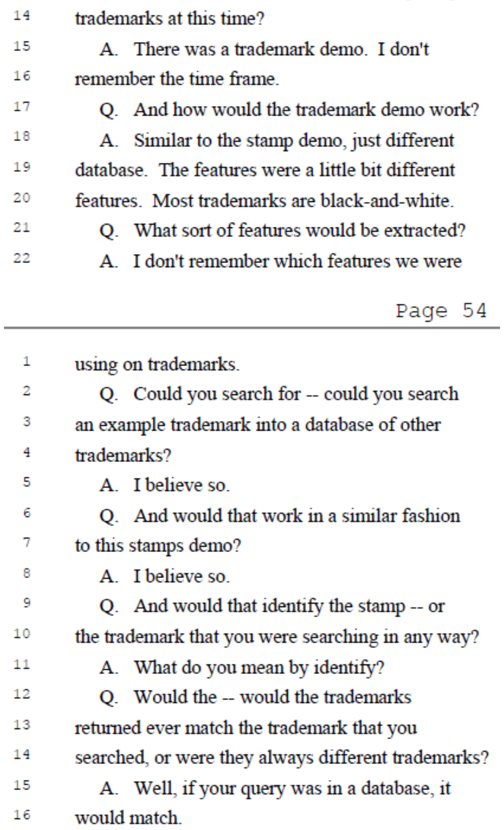 Q. Was there also a demo online comparing trademarks at this time? A. There was a trademark demo. I don't remember the time frame. Q. And how would the trademark demo work? A. Similar to the stamp demo, just different database. The features were a little bit different features. Most trademarks are black-and-white. Q. What sort of features would be extracted? A. I don't remember which features we were Page 54 using on trademarks. Q. Could you search for -- could you search an example trademark into a database of other trademarks? A. I believe so. Q. And would that work in a similar fashion to this stamps demo? A. I believe so. Q. And would that identify the stamp -- or the trademark that you were searching in any way? A. What do you mean by identify? Q. Would the -- would the trademarks returned ever match the trademark that you searched, or were they always different trademarks? A. Well, if your query was in a database, it would match. Flickner Depo at 62-63 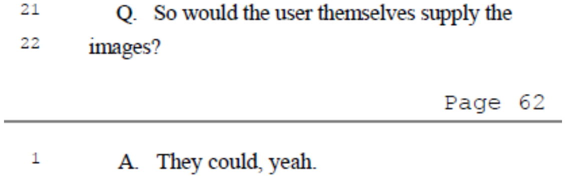 Q. So would the user themselves supply the images? Page 62 A. They could, yeah.

Element	'278 Claim Recitation	Exemplary Citations in Reference
		Flickner 8/29/23 Depo at 53 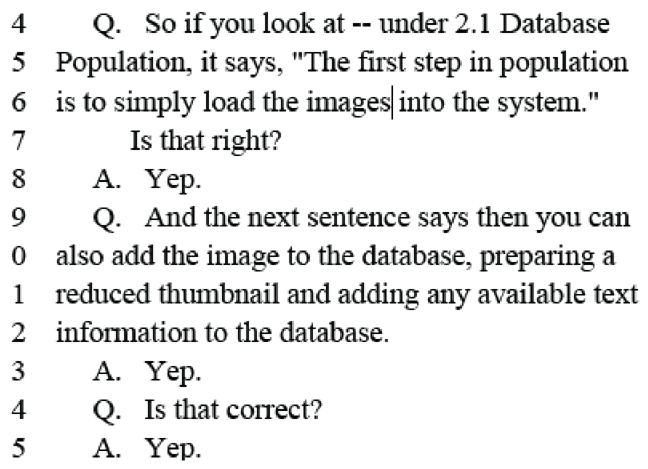 Q. So if you look at -- under 2.1 Database Population, it says, "The first step in population is to simply load the images into the system." Is that right? A. Yep. Q. And the next sentence says then you can also add the image to the database, preparing a reduced thumbnail and adding any available text information to the database. A. Yep. Q. Is that correct? A. Yep. *Id.* at 60–61

Element	'278 Claim Recitation	Exemplary Citations in Reference
		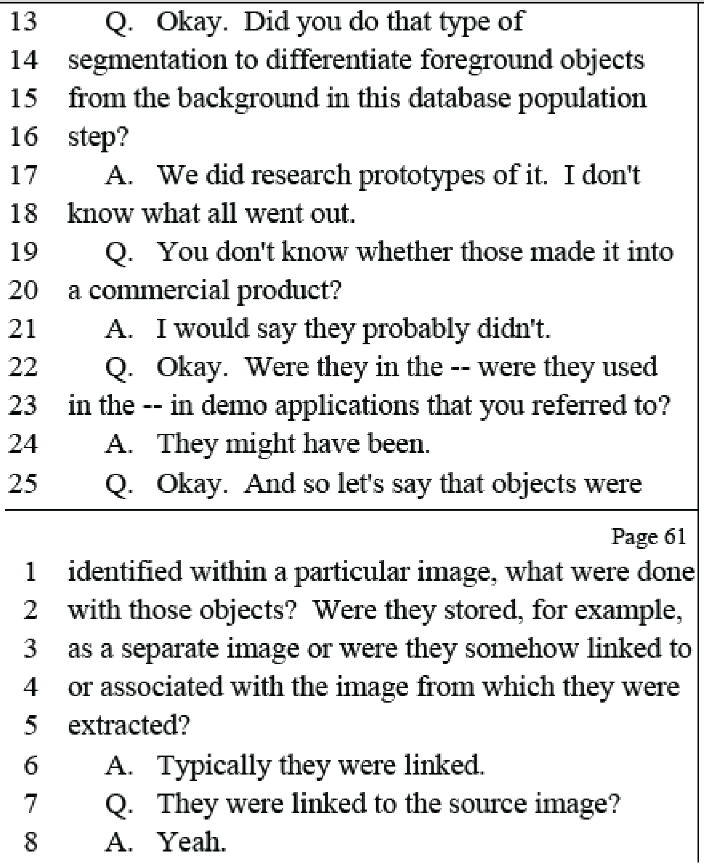 Q. Okay. Did you do that type of segmentation to differentiate foreground objects from the background in this database population step? A. We did research prototypes of it. I don't know what all went out. Q. You don't know whether those made it into a commercial product? A. I would say they probably didn't. Q. Okay. Were they in the -- were they used in the -- in demo applications that you referred to? A. They might have been. Q. Okay. And so let's say that objects were Page 61 identified within a particular image, what were done with those objects? Were they stored, for example, as a separate image or were they somehow linked to or associated with the image from which they were extracted? A. Typically they were linked. Q. They were linked to the source image? A. Yeah. *Id.* at 32–33

Element	'278 Claim Recitation	Exemplary Citations in Reference
		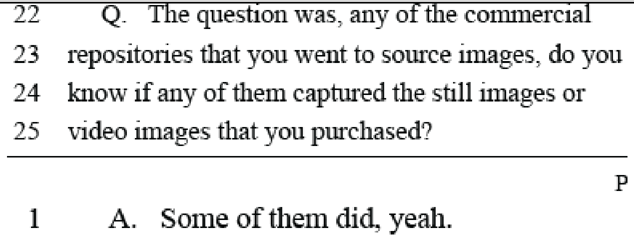 Q. The question was, any of the commercial repositories that you went to source images, do you know if any of them captured the still images or video images that you purchased? P A. Some of them did, yeah. *Id.* at 74–75

			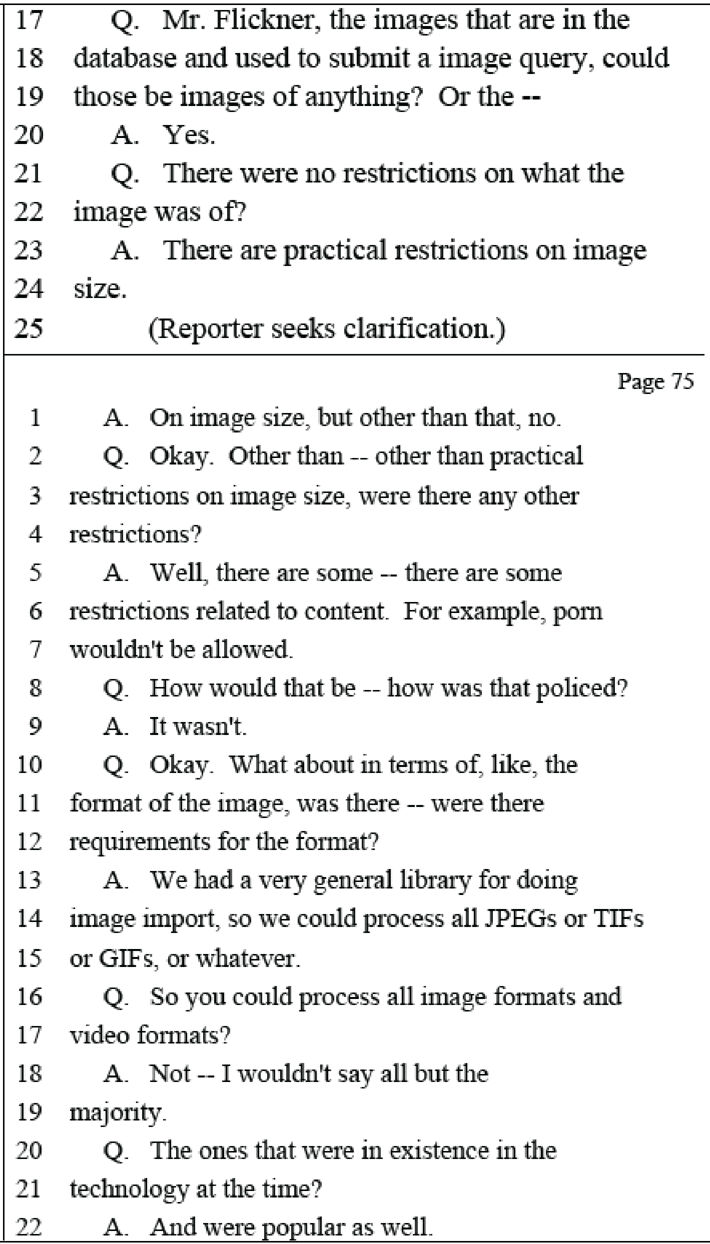 Q. Mr. Flickner, the images that are in the database and used to submit a image query, could those be images of anything? Or the -- A. Yes. Q. There were no restrictions on what the image was of? A. There are practical restrictions on image size. (Reporter seeks clarification.) Page 75 A. On image size, but other than that, no. Q. Okay. Other than -- other than practical restrictions on image size, were there any other restrictions? A. Well, there are some -- there are some restrictions related to content. For example, porn wouldn't be allowed. Q. How would that be -- how was that policed? A. It wasn't. Q. Okay. What about in terms of, like, the format of the image, was there -- were there requirements for the format? A. We had a very general library for doing image import, so we could process all JPEGs or TIFs or GIFs, or whatever. Q. So you could process all image formats and video formats? A. Not -- I wouldn't say all but the majority. Q. The ones that were in existence in the technology at the time? A. And were popular as well.

Element	'278 Claim Recitation	Exemplary Citations in Reference
		Id. at 114 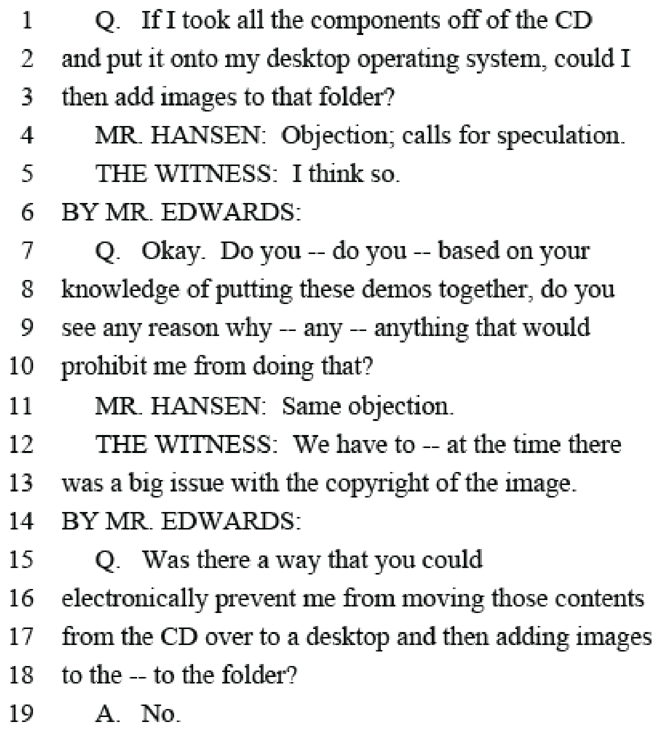 Q. If I took all the components off of the CD and put it onto my desktop operating system, could I then add images to that folder? MR. HANSEN: Objection; calls for speculation. THE WITNESS: I think so. BY MR. EDWARDS: Q. Okay. Do you -- do you -- based on your knowledge of putting these demos together, do you see any reason why -- any -- anything that would prohibit me from doing that? MR. HANSEN: Same objection. THE WITNESS: We have to -- at the time there was a big issue with the copyright of the image. BY MR. EDWARDS: Q. Was there a way that you could electronically prevent me from moving those contents from the CD over to a desktop and then adding images to the -- to the folder? A. No. *Id.* at 202

Element	'278 Claim Recitation	Exemplary Citations in Reference
		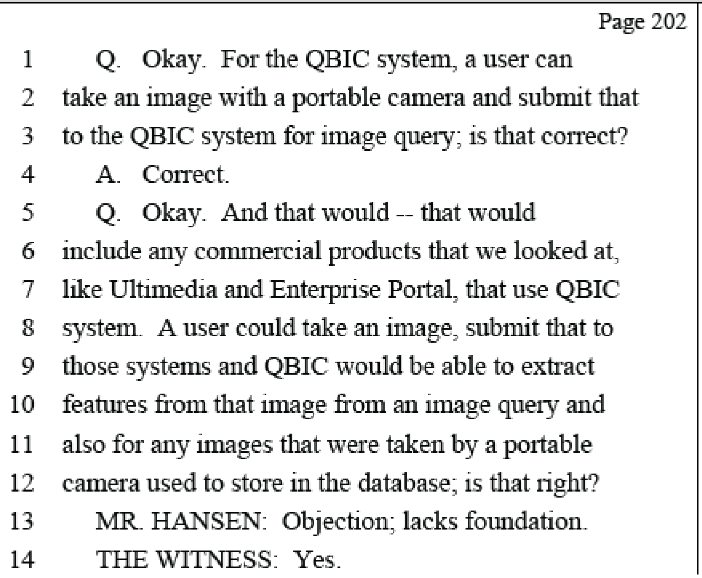 Page 202 Q. Okay. For the QBIC system, a user can take an image with a portable camera and submit that to the QBIC system for image query; is that correct? A. Correct. Q. Okay. And that would -- that would include any commercial products that we looked at, like Ultimedia and Enterprise Portal, that use QBIC system. A user could take an image, submit that to those systems and QBIC would be able to extract features from that image from an image query and also for any images that were taken by a portable camera used to store in the database; is that right? MR. HANSEN: Objection; lacks foundation. THE WITNESS: Yes. *Id.* at 300 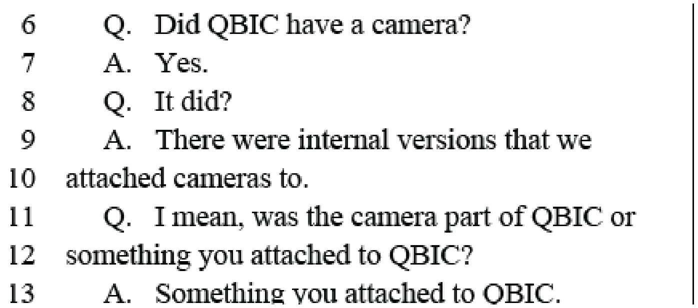 Q. Did QBIC have a camera? A. Yes. Q. It did? A. There were internal versions that we attached cameras to. Q. I mean, was the camera part of QBIC or something you attached to QBIC? A. Something you attached to QBIC. *Id.* at 312

Element	'278 Claim Recitation	Exemplary Citations in Reference
		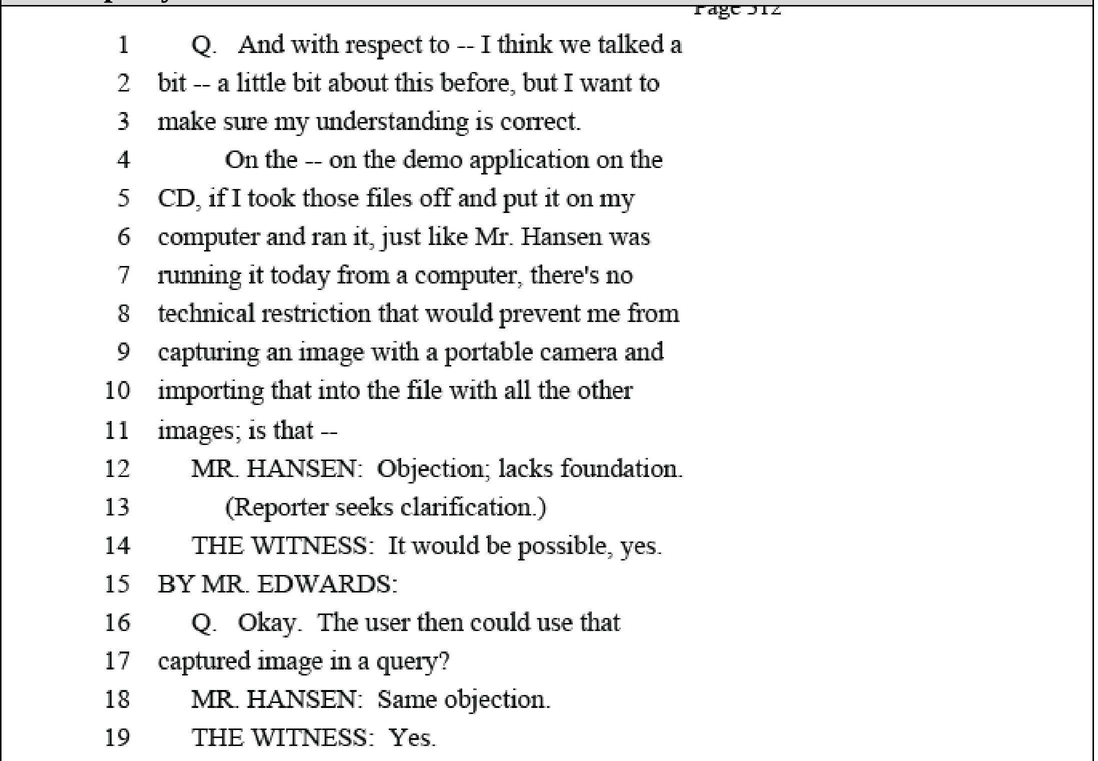 Page 312 Q. And with respect to -- I think we talked a bit -- a little bit about this before, but I want to make sure my understanding is correct. On the -- on the demo application on the CD, if I took those files off and put it on my computer and ran it, just like Mr. Hansen was running it today from a computer, there's no technical restriction that would prevent me from capturing an image with a portable camera and importing that into the file with all the other images; is that -- MR. HANSEN: Objection; lacks foundation. (Reporter seeks clarification.) THE WITNESS: It would be possible, yes. BY MR. EDWARDS: Q. Okay. The user then could use that captured image in a query? MR. HANSEN: Same objection. THE WITNESS: Yes. *See also* Flickner Depo at, e.g., 24, 41-42, 62-63 (user could take a picture or video and submit it to the QBIC system). To the extent it is found that QBIC does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that QBIC does not anticipate this claim, QBIC renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine QBIC with any of the other prior art references identified in the cover

Element	'278 Claim Recitation	Exemplary Citations in Reference
		pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1B]	determining that symbolic content is on the at least the portion of the document based on the image;	QBIC discloses determining that symbolic content is on the at least the portion of the document based on the image. Query by Image and Video Content: The QBIC System (IBM 0000893–902): "We have developed the QBIC (Query by Image Content) system to explore content-based retrieval methods. QBIC allows queries on large image and video databases based on • example images, • user-constructed sketches and drawings, • selected color and texture patterns, • camera and object motion, and • other graphical information." *Id.* at pp. 1–3. "**DATABASE POPULATION** In still image database population, the images are reduced to a standard-sized icon called a thumbnail and annotated with any available text information. Object identification is an optional but key part of this step. It lets users manually, semiautomatically, or fully automatically identify interesting regions—which we call objects—in the images. Internally, each object is represented as a binary mask. There may be an arbitrary number of objects per image. Objects can overlap and can consist of multiple disconnected components like the set of dots on a polka-dot dress. Text, like 'baby on beach,' can be associated with an outlined object or with the scene as a whole. **Object-outlining tools** Ideally, object identification would be automatic, but this is generally difficult. The alternative—manual identification—is tedious and can inhibit query-by-content applications. As a result, we have devoted considerable effort to

Element	'278 Claim Recitation	Exemplary Citations in Reference
		developing tools to aid in this step. In recent work, we have successfully used fully automatic unsupervised segmentation methods along with a foreground/background model to identify objects in a restricted class of images. The images, typical of museums and retail catalogs, have a small number of foreground objects on a generally separable background. Figure 4 shows example results. Even in this domain, robust algorithms are required because of the textured and variegated backgrounds. We also provide semiautomatic tools for identifying objects. One is an enhanced flood-fill technique. Flood-fill methods, found in most photo-editing programs, start from a single object pixel and repeatedly add adjacent pixels whose values are within some given threshold of the original pixel. Selecting the threshold, which must change from image to image and object to object, is tedious. We automatically calculate a dynamic threshold by having the user click on background as well as object points. For reasonably uniform objects that are distinct from the background, this operation allows fast object identification without manually adjusting a threshold. The example in Figure 3 shows an object, a fox, identified by using only a few clicks. We designed another outlining tool to help users track object edges. This tool takes a user-drawn curve and automatically aligns it with nearby image edges. Based on the 'snakes' concept developed in recent computer vision research, the tool finds the curve that maximizes the image gradient magnitude along the curve. The spline snake formulation we use allows for smooth solutions to the resulting nonlinear minimization problem. The computation is done at interactive speeds so that, as the user draws a curve, it is 'rubber-banded' to lie along object boundaries. . . . **Video data** For video data, database population has three major components: • shot detection, • representative frame creation for each shot, and • derivation of a layered representation of coherently moving structures/ objects. . . . Once the shot boundaries have been detected, each shot is represented using an

Element	'278 Claim Recitation	Exemplary Citations in Reference
		r-frame. R-frames are used for several purposes. First, during database population, r-frames are treated as still images in which objects can be identified by using the previously described methods." *Id*. at pp. 3–6. "INTEGRATION WITH TEXT AND PARAMETRIC ANNOTATION. Query by visual content complements and extends existing query methods. Systems must be able to integrate queries combining date, subject matter, price, and availability with content properties such as color, texture, and shape." *Id*. at p. 8. Image, Audio, and Video Extenders Administration and Programming (IBM 000001–617): "DB2 (DB2) Universal Database (UDB) is a powerful, object-relational database manager. It stores and protects traditional numeric and character data, as well as large, complex objects (LOBs). DB2 extenders help you exploit DB2 's object-relational features. The extenders define distinct data types and special functions for image, audio, video, and text objects. … Because of that, the extenders give your applications a single point of access to any or all of these types of data, along with traditional numeric and character data." *Id*. at p. 3. "A QBIC catalog is a set of files that hold data about the visual features of images. The Image Extender uses this data to search for images by content. You create a QBIC catalog for each column of images in a user table that you want to make available for searching by content. When you create a QBIC catalog you identify the features for which you want the Image Extender to analyze, store, and later query data. You can also add or drop features from a QBIC catalog after the catalog is created. A QBIC catalog can hold data for the following image features: **Average color** The sum of the color values for all pixels in an image divided by the number of

Element	'278 Claim Recitation	Exemplary Citations in Reference
		pixels in the image. (A pixel is the smallest element of an image that can be assigned color and intensity.) For example, if 50% of an image consists of blue pixels and the other 50% red pixels, the image has an average color value of purple. Average color is used to search for images that have a predominant color. If an image has a predominant color, the average color will be similar to the predominant color. **Histogram color** Measures the distribution of colors in an image against a spectrum of 64 colors. For each of the 64 colors, histogram color identifies the percentage of pixels in an image that have that color. For example, the histogram color of an image might be 40% white pixels, 50% blue, and 10% red; none of the pixels in the image have any of the remaining colors in the histogram spectrum. Histogram color is used to search for images that have a variety of colors. **Positional color** The average color value for the pixels in a specified area in an image. For example, the upper right-hand corner of an image might show a bright yellow sun; the positional color of this area of the image is bright yellow. Positional color is used to search for images that have a predominant color in a particular area. **Texture** Measures the coarseness, contrast, and directionality of an image. Coarseness indicates the size of repeating items in an image (for example, pebbles versus boulders). Contrast identifies the brightness variations in an image (light versus dark). Directionality indicates whether a direction predominates in an image (as in the vertical direction of a picket fence) or does not predominate (as in an image of sand). Texture is used to search for images that have a particular pattern. To make an image available for searching by content, you catalog the image. When you catalog an image, the Image Extender analyzes the image, by computing the feature values for the image, and stores the values in a QBIC catalog.

Element	'278 Claim Recitation	Exemplary Citations in Reference
		When you search for an image by content, your query identifies one or more features for the search (such as average color), a source for each feature (such as an example image), and a target set of cataloged images. The Image Extender computes the feature value of the source and compares it to the cataloged feature values for the target images. It then computes a score that indicates how similar the feature values of the target images are to the source. You can have the Image Extender return the images whose features are most similar to the source. The Image Extender will return the handle of each image and the image score. You can also have the Image Extender return only the score of a single image." *Id.* at pp. 20–21. Managing Enterprise Information Portal (IBM 000777–880): "Many paper-intensive enterprises, such as insurance companies and financial institutions, administer large volumes of business-related content. The need for an enterprise solution for managing and accessing business information spans many industries. A content server is a software system that stores multimedia, business forms, documents, and related data, along with metadata that allows employees to process and work with the content." *Id.* at p. 3. "The image search server uses IBM's QBIC (query by image content) technology to help you search for objects by certain visual properties, such as color and texture. The image search server analyzes images and stores the image information in a database. Then users can run image queries, which use the visual properties of images, to match colors, textures, and their positions without describing them in words. You can combine content-based queries with text and keyword searches for more powerful retrieval of image and multimedia data." *Id.* at p. 63. IBM0002390 at 2391-92:

Element	'278 Claim Recitation	Exemplary Citations in Reference
		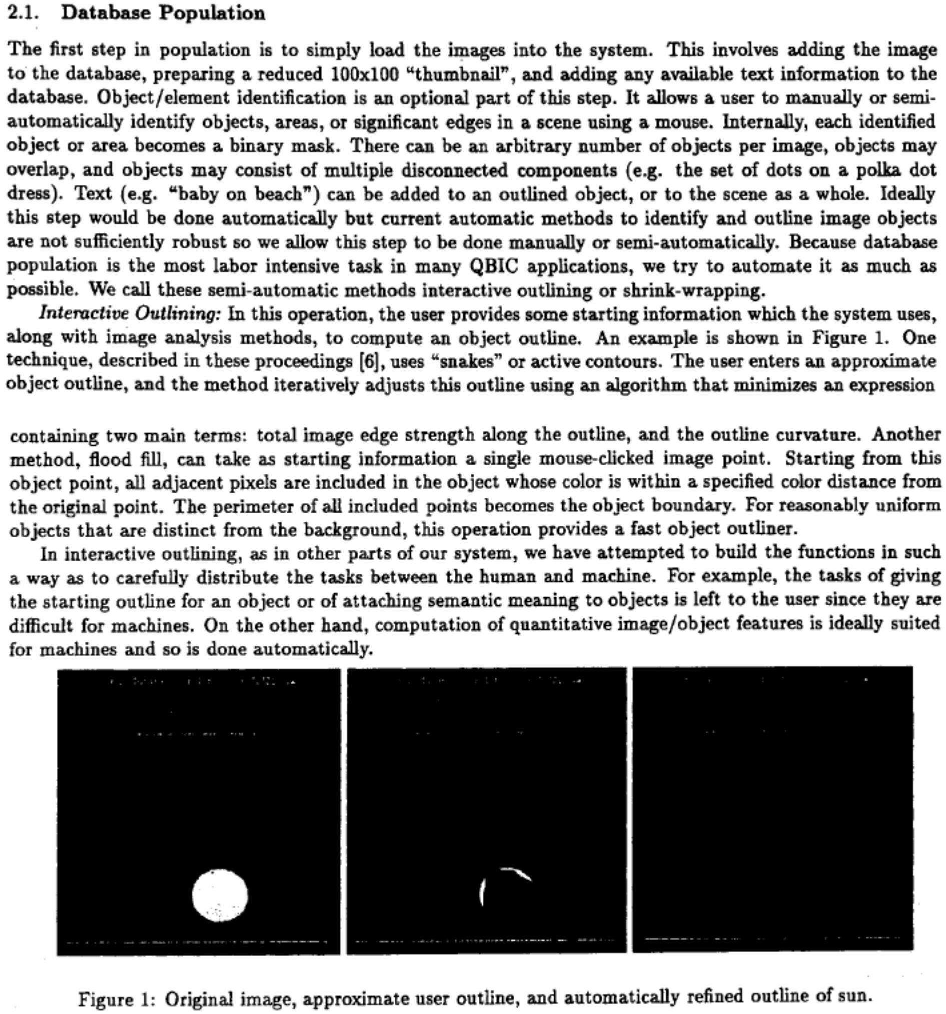 **2.1. Database Population** The first step in population is to simply load the images into the system. This involves adding the image to the database, preparing a reduced 100x100 "thumbnail", and adding any available text information to the database. Object/element identification is an optional part of this step. It allows a user to manually or semi-automatically identify objects, areas, or significant edges in a scene using a mouse. Internally, each identified object or area becomes a binary mask. There can be an arbitrary number of objects per image, objects may overlap, and objects may consist of multiple disconnected components (e.g. the set of dots on a polka dot dress). Text (e.g. "baby on beach") can be added to an outlined object, or to the scene as a whole. Ideally this step would be done automatically but current automatic methods to identify and outline image objects are not sufficiently robust so we allow this step to be done manually or semi-automatically. Because database population is the most labor intensive task in many QBIC applications, we try to automate it as much as possible. We call these semi-automatic methods interactive outlining or shrink-wrapping. *Interactive Outlining:* In this operation, the user provides some starting information which the system uses, along with image analysis methods, to compute an object outline. An example is shown in Figure 1. One technique, described in these proceedings [6], uses "snakes" or active contours. The user enters an approximate object outline, and the method iteratively adjusts this outline using an algorithm that minimizes an expression containing two main terms: total image edge strength along the outline, and the outline curvature. Another method, flood fill, can take as starting information a single mouse-clicked image point. Starting from this object point, all adjacent pixels are included in the object whose color is within a specified color distance from the original point. The perimeter of all included points becomes the object boundary. For reasonably uniform objects that are distinct from the background, this operation provides a fast object outliner. In interactive outlining, as in other parts of our system, we have attempted to build the functions in such a way as to carefully distribute the tasks between the human and machine. For example, the tasks of giving the starting outline for an object or of attaching semantic meaning to objects is left to the user since they are difficult for machines. On the other hand, computation of quantitative image/object features is ideally suited for machines and so is done automatically. Figure 1: Original image, approximate user outline, and automatically refined outline of sun.

Element	**'278 Claim Recitation**	**Exemplary Citations in Reference**
		QBIC Demo (IBM000747)  Flickner Depo at 162 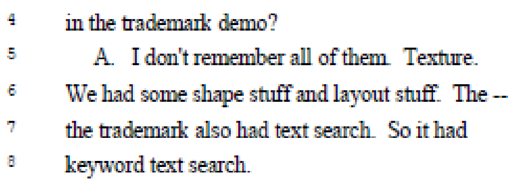 What -- what -- what features did you use in the trademark demo? A. I don't remember all of them. Texture. We had some shape stuff and layout stuff. The -- the trademark also had text search. So it had keyword text search. Flickner Depo at 50

Element	'278 Claim Recitation	Exemplary Citations in Reference
		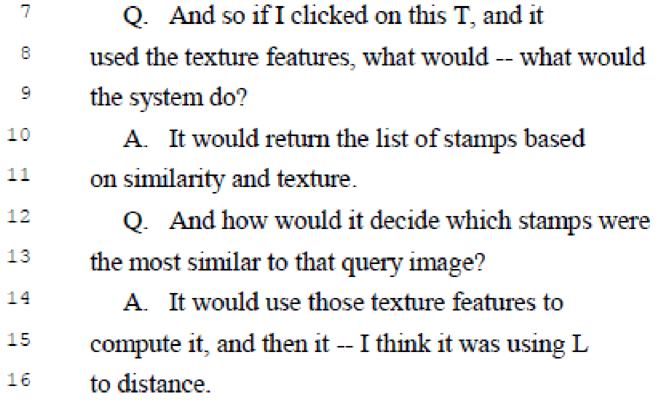 Q. And so if I clicked on this T, and it used the texture features, what would -- what would the system do? A. It would return the list of stamps based on similarity and texture. Q. And how would it decide which stamps were the most similar to that query image? A. It would use those texture features to compute it, and then it -- I think it was using L to distance. Flickner Depo at 22-25

Element	'278 Claim Recitation	Exemplary Citations in Reference
		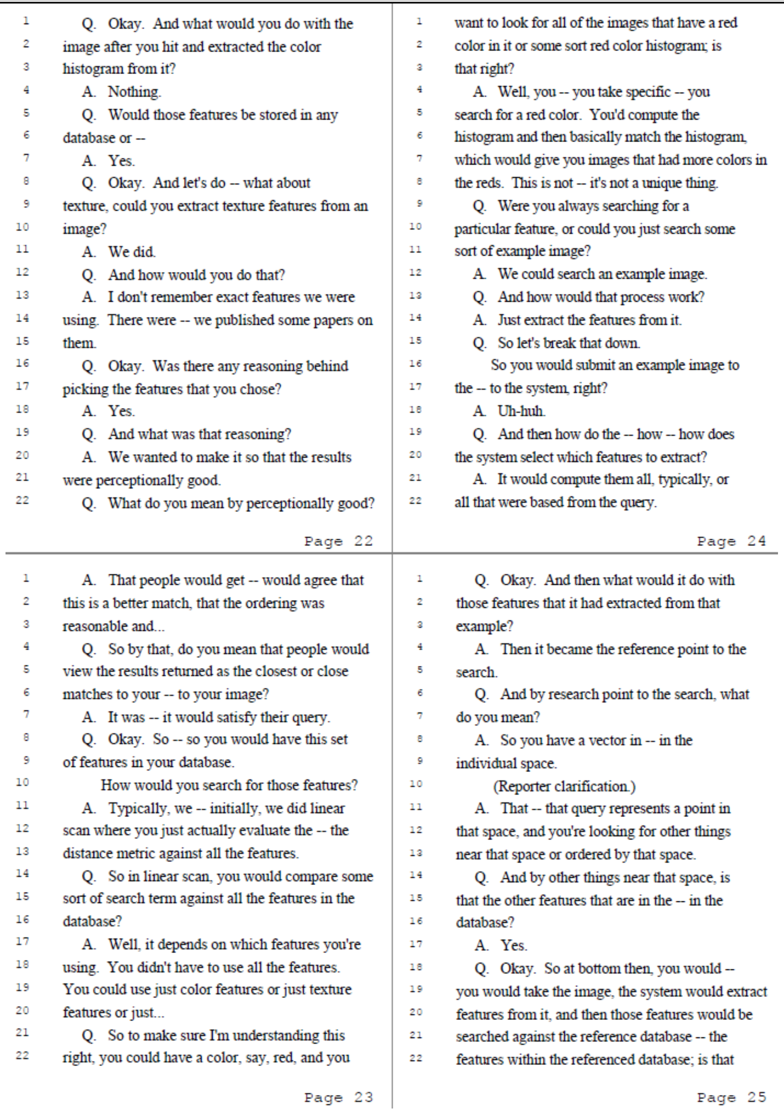

Q. Okay. And what would you do with the image after you hit and extracted the color histogram from it?
A. Nothing.
Q. Would those features be stored in any database or --
A. Yes.
Q. Okay. And let's do -- what about texture, could you extract texture features from an image?
A. We did.
Q. And how would you do that?
A. I don't remember exact features we were using. There were -- we published some papers on them.
Q. Okay. Was there any reasoning behind picking the features that you chose?
A. Yes.
Q. And what was that reasoning?
A. We wanted to make it so that the results were perceptionally good.
Q. What do you mean by perceptionally good?

Page 22

A. That people would get -- would agree that this is a better match, that the ordering was reasonable and...
Q. So by that, do you mean that people would view the results returned as the closest or close matches to your -- to your image?
A. It was -- it would satisfy their query.
Q. Okay. So -- so you would have this set of features in your database.
How would you search for those features?
A. Typically, we -- initially, we did linear scan where you just actually evaluate the -- the distance metric against all the features.
Q. So in linear scan, you would compare some sort of search term against all the features in the database?
A. Well, it depends on which features you're using. You didn't have to use all the features. You could use just color features or just texture features or just...
Q. So to make sure I'm understanding this right, you could have a color, say, red, and you

Page 23

want to look for all of the images that have a red color in it or some sort red color histogram; is that right?
A. Well, you -- you take specific -- you search for a red color. You'd compute the histogram and then basically match the histogram, which would give you images that had more colors in the reds. This is not -- it's not a unique thing.
Q. Were you always searching for a particular feature, or could you just search some sort of example image?
A. We could search an example image.
Q. And how would that process work?
A. Just extract the features from it.
Q. So let's break that down.
So you would submit an example image to the -- to the system, right?
A. Uh-huh.
Q. And then how do the -- how -- how does the system select which features to extract?
A. It would compute them all, typically, or all that were based from the query.

Page 24

Q. Okay. And then what would it do with those features that it had extracted from that example?
A. Then it became the reference point to the search.
Q. And by research point to the search, what do you mean?
A. So you have a vector in -- in the individual space.
(Reporter clarification.)
A. That -- that query represents a point in that space, and you're looking for other things near that space or ordered by that space.
Q. And by other things near that space, is that the other features that are in the -- in the database?
A. Yes.
Q. Okay. So at bottom then, you would -- you would take the image, the system would extract features from it, and then those features would be searched against the reference database -- the features within the referenced database; is that

Page 25

Element	'278 Claim Recitation	Exemplary Citations in Reference
		Flickner Depo at 36 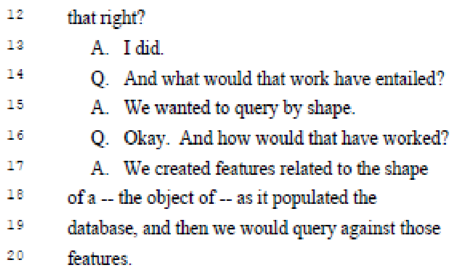 Q. So underlying this paper, you mentioned you did more work on the shape side of things; is that right? A. I did. Q. And what would that work have entailed? A. We wanted to query by shape. Q. Okay. And how would that have worked? A. We created features related to the shape of a -- the object of -- as it populated the database, and then we would query against those features. Flickner Depo at 214-215

Element	'278 Claim Recitation	Exemplary Citations in Reference
		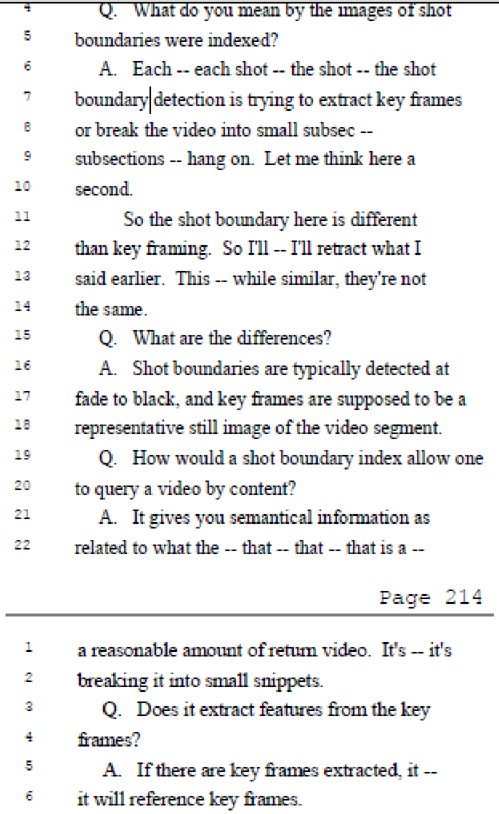 Q. What do you mean by the images of shot boundaries were indexed? A. Each -- each shot -- the shot -- the shot boundary detection is trying to extract key frames or break the video into small subsec -- subsections -- hang on. Let me think here a second. So the shot boundary here is different than key framing. So I'll -- I'll retract what I said earlier. This -- while similar, they're not the same. Q. What are the differences? A. Shot boundaries are typically detected at fade to black, and key frames are supposed to be a representative still image of the video segment. Q. How would a shot boundary index allow one to query a video by content? A. It gives you semantical information as related to what the -- that -- that -- that is a -- Page 214 a reasonable amount of return video. It's -- it's breaking it into small snippets. Q. Does it extract features from the key frames? A. If there are key frames extracted, it -- it will reference key frames. Flickner Depo at 40-42

		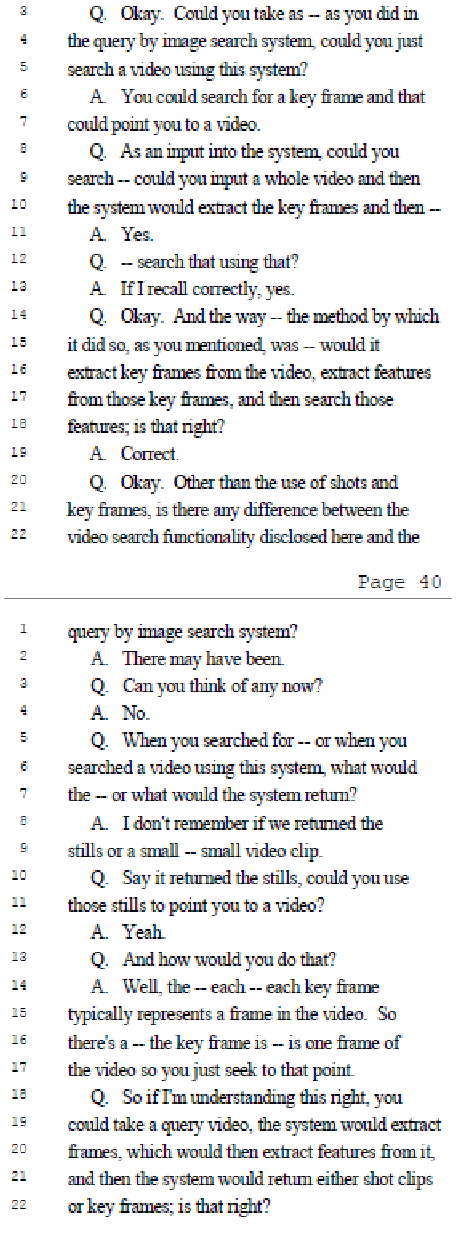 Q. Okay. Could you take as -- as you did in the query by image search system, could you just search a video using this system? A. You could search for a key frame and that could point you to a video. Q. As an input into the system, could you search -- could you input a whole video and then the system would extract the key frames and then -- A. Yes. Q. -- search that using that? A. If I recall correctly, yes. Q. Okay. And the way -- the method by which it did so, as you mentioned, was -- would it extract key frames from the video, extract features from those key frames, and then search those features; is that right? A. Correct. Q. Okay. Other than the use of shots and key frames, is there any difference between the video search functionality disclosed here and the Page 40 query by image search system? A. There may have been. Q. Can you think of any now? A. No. Q. When you searched for -- or when you searched a video using this system, what would the -- or what would the system return? A. I don't remember if we returned the stills or a small -- small video clip. Q. Say it returned the stills, could you use those stills to point you to a video? A. Yeah. Q. And how would you do that? A. Well, the -- each -- each key frame typically represents a frame in the video. So there's a -- the key frame is -- is one frame of the video so you just seek to that point. Q. So if I'm understanding this right, you could take a query video, the system would extract frames, which would then extract features from it, and then the system would return either shot clips or key frames; is that right?

Element	'278 Claim Recitation	Exemplary Citations in Reference
		 A. Right. Flickner Depo at 87. 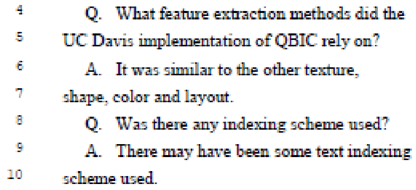 Q. What feature extraction methods did the UC Davis implementation of QBIC rely on? A. It was similar to the other texture, shape, color and layout. Q. Was there any indexing scheme used? A. There may have been some text indexing scheme used. Flickner 8/29/23 Depo at 42 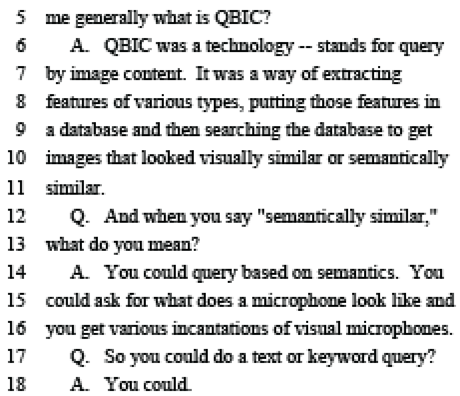 Q. So let's just start at a high level. Tell me generally what is QBIC? A. QBIC was a technology -- stands for query by image content. It was a way of extracting features of various types, putting those features in a database and then searching the database to get images that looked visually similar or semantically similar. Q. And when you say "semantically similar," what do you mean? A. You could query based on semantics. You could ask for what does a microphone look like and you get various incantations of visual microphones. Q. So you could do a text or keyword query? A. You could. *Id.* at 44 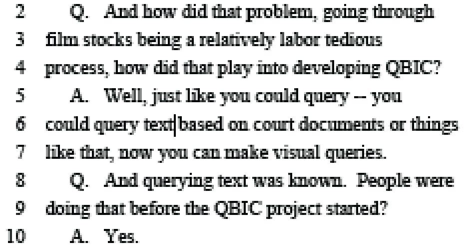 Q. And how did that problem, going through film stocks being a relatively labor tedious process, how did that play into developing QBIC? A. Well, just like you could query -- you could query text based on court documents or things like that, now you can make visual queries. Q. And querying text was known. People were doing that before the QBIC project started? A. Yes.

Element	'278 Claim Recitation	Exemplary Citations in Reference
		Id. at 45–47 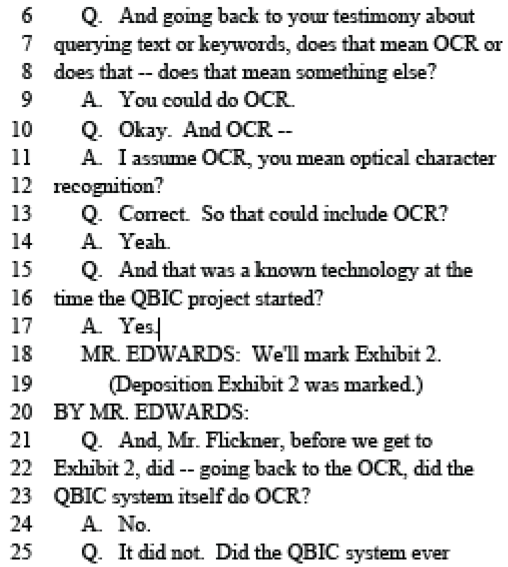 Q. And going back to your testimony about querying text or keywords, does that mean OCR or does that -- does that mean something else? A. You could do OCR. Q. Okay. And OCR -- A. I assume OCR, you mean optical character recognition? Q. Correct. So that could include OCR? A. Yeah. Q. And that was a known technology at the time the QBIC project started? A. Yes. MR. EDWARDS: We'll mark Exhibit 2. (Deposition Exhibit 2 was marked.) BY MR. EDWARDS: Q. And, Mr. Flickner, before we get to Exhibit 2, did -- going back to the OCR, did the QBIC system itself do OCR? A. No. Q. It did not. Did the QBIC system ever

Element	’278 Claim Recitation	Exemplary Citations in Reference
		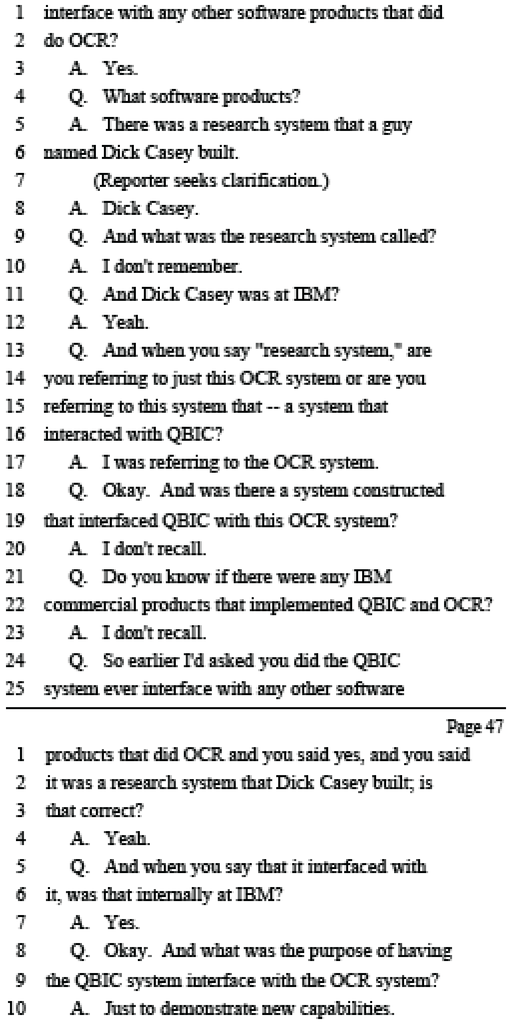 interface with any other software products that did do OCR? A. Yes. Q. What software products? A. There was a research system that a guy named Dick Casey built. (Reporter seeks clarification.) A. Dick Casey. Q. And what was the research system called? A. I don't remember. Q. And Dick Casey was at IBM? A. Yeah. Q. And when you say "research system," are you referring to just this OCR system or are you referring to this system that -- a system that interacted with QBIC? A. I was referring to the OCR system. Q. Okay. And was there a system constructed that interfaced QBIC with this OCR system? A. I don't recall. Q. Do you know if there were any IBM commercial products that implemented QBIC and OCR? A. I don't recall. Q. So earlier I'd asked you did the QBIC system ever interface with any other software Page 47 products that did OCR and you said yes, and you said it was a research system that Dick Casey built; is that correct? A. Yeah. Q. And when you say that it interfaced with it, was that internally at IBM? A. Yes. Q. Okay. And what was the purpose of having the QBIC system interface with the OCR system? A. Just to demonstrate new capabilities. *Id.* at 76

Element	'278 Claim Recitation	Exemplary Citations in Reference
		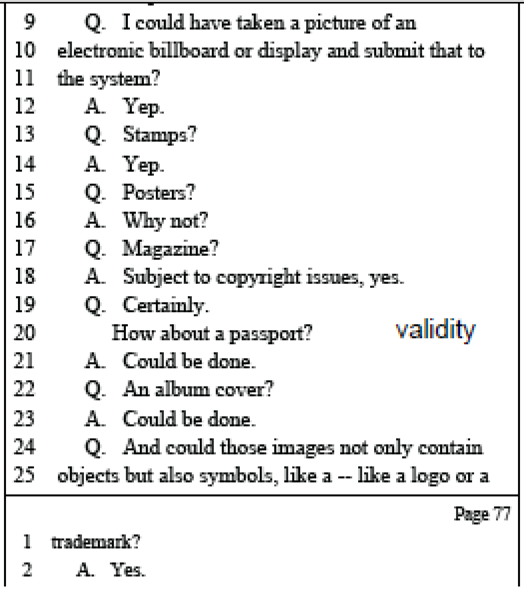 Q. I could have taken a picture of an electronic billboard or display and submit that to the system? A. Yep. Q. Stamps? A. Yep. Q. Posters? A. Why not? Q. Magazine? A. Subject to copyright issues, yes. Q. Certainly. How about a passport? validity A. Could be done. Q. An album cover? A. Could be done. Q. And could those images not only contain objects but also symbols, like a -- like a logo or a Page 77 trademark? A. Yes. *Id.* at 187–188 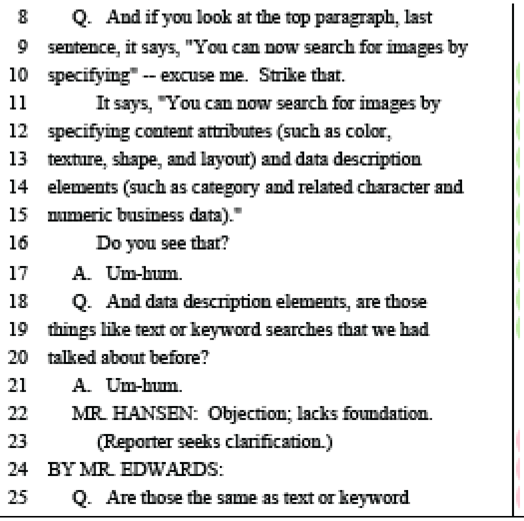 Q. And if you look at the top paragraph, last sentence, it says, "You can now search for images by specifying" -- excuse me. Strike that. It says, "You can now search for images by specifying content attributes (such as color, texture, shape, and layout) and data description elements (such as category and related character and numeric business data)." Do you see that? A. Um-hum. Q. And data description elements, are those things like text or keyword searches that we had talked about before? A. Um-hum. MR. HANSEN: Objection; lacks foundation. (Reporter seeks clarification.) BY MR. EDWARDS: Q. Are those the same as text or keyword

Element	'278 Claim Recitation	Exemplary Citations in Reference
		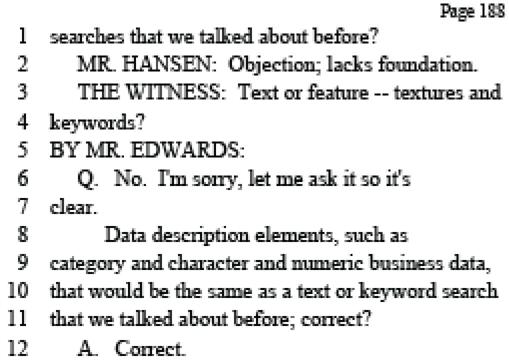 Page 188 searches that we talked about before? MR. HANSEN: Objection; lacks foundation. THE WITNESS: Text or feature -- textures and keywords? BY MR. EDWARDS: Q. No. I'm sorry, let me ask it so it's clear. Data description elements, such as category and character and numeric business data, that would be the same as a text or keyword search that we talked about before; correct? A. Correct. *Id.* at 190–191 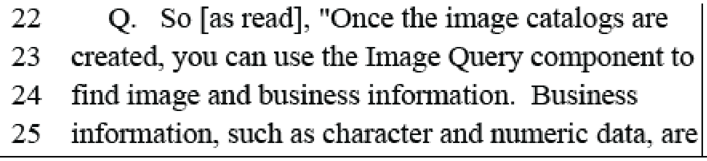 Q. So [as read], "Once the image catalogs are created, you can use the Image Query component to find image and business information. Business information, such as character and numeric data, are

		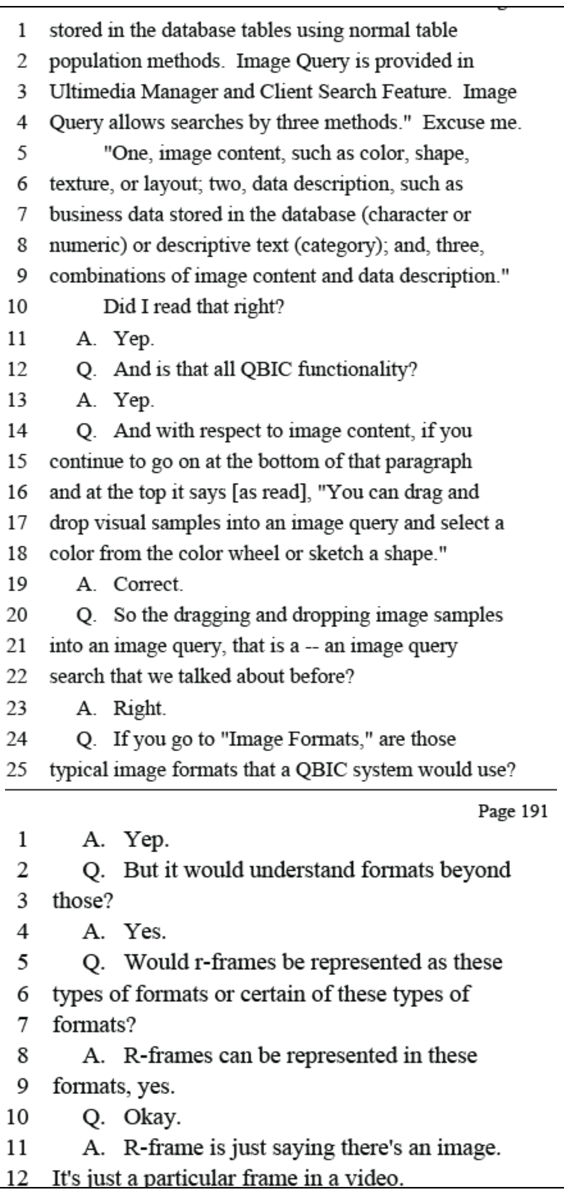 stored in the database tables using normal table population methods. Image Query is provided in Ultimedia Manager and Client Search Feature. Image Query allows searches by three methods." Excuse me. "One, image content, such as color, shape, texture, or layout; two, data description, such as business data stored in the database (character or numeric) or descriptive text (category); and, three, combinations of image content and data description." Did I read that right? A. Yep. Q. And is that all QBIC functionality? A. Yep. Q. And with respect to image content, if you continue to go on at the bottom of that paragraph and at the top it says [as read], "You can drag and drop visual samples into an image query and select a color from the color wheel or sketch a shape." A. Correct. Q. So the dragging and dropping image samples into an image query, that is a -- an image query search that we talked about before? A. Right. Q. If you go to "Image Formats," are those typical image formats that a QBIC system would use? Page 191 A. Yep. Q. But it would understand formats beyond those? A. Yes. Q. Would r-frames be represented as these types of formats or certain of these types of formats? A. R-frames can be represented in these formats, yes. Q. Okay. A. R-frame is just saying there's an image. It's just a particular frame in a video.

Element	'278 Claim Recitation	Exemplary Citations in Reference
		IBM0002263 at 2264

Description

Ultimedia Manager 1.1 opens new dimensions in visual queries, helping you manage your image database more efficiently. With its easy-to-use interface, even entry-level users can perform sophisticated searches. You can search by image content, data description, or a combination of both. Once the images are found, you can view the images in full resolution, use the images in another application, or read the associated textual comments.

You can use Ultimedia Manager 1.1 either as a stand-alone product or in a LAN environment. As a stand-alone product, Ultimedia Manager 1.1 employs full image classification and image query capabilities. Within the LAN environment, the Ultimedia Manager image catalogs can be sorted on a DB2 for OS/2 or DB2 for AIX database server. The image catalogs can be accessed by Windows or OS/2 clients using the Client Search Feature or the Visualizer Ultimedia Query.

Ultimedia Manager runs as an OS/2 application to a DB2 for OS/2 or DB2 for AIX database server in a LAN, or as a stand-alone application with a local DB2 for OS/2 database. The Client Search Feature runs on DOS or Windows, and accesses files in dBASE flat file format. A database manager is not required to access catalogs in dBASE file format.

Image Classification

With the Image Classifier, you can create visual indexes based on color, shape, category, and layout elements, which are areas of the image containing the patterns you want to use for searches. You can store these elements in catalogs in DB2 for OS/2 databases. The catalogs can be used locally or, through the use of utility programs, be transferred to DB2 for OS/2 or AIX servers. In addition, you can convert the catalog data into dBASE flat file format for access by Windows client applications.

The Image Classifier and Analyzer allows automatic or interactive classification of images for later image queries, flexibility in storage of catalog information, and access to multiple users. You can organize large image catalogs so that later searches for images can be performed easily. You can classify automatically using a batch program, or interactively by specifying areas of interest in an image and indicating the visual cue (color, shape, or texture) to be saved for later queries.

Image Query

Once the image catalogs are created, you can use the Image Query component to find image and business information. Business information, such as character and numeric data, are stored in the database tables using normal table population methods. Image Query is provided in Ultimedia Manager and Client Search Feature.

Image Query allows searches by three methods:

- Image content, such as color, shape, texture, or layout
- Data description, such as business data stored in the database (character or numeric) or descriptive text (category)
- Combinations of image content and data description.

Image Content: To search by image content, specify the attributes (color, shape, texture, or layout). You can drag and drop visual samples into an image query, select a color from the color wheel, or sketch a shape.

Data Description: To search by data description, enter character and numeric data values in the database tables using the simple fill-in-the-blanks query panel. You can enter data such as name, number, date, or a range of dates. You can integrate character or numeric data as part of the image query, enabling a more effective search for better results. The simple fill-in-the-blanks method for specifying desired conditions is so easy even non-programmers can use database table information as part of their searches.

Combinations: Image queries using either image content or data description, or combinations of both, allow you to search through thousands of images more quickly and effectively than the traditional text and keyword search methods.

Image Formats

Ultimedia Manager 1.1 understands these image formats:

- Audio Visual Connection (AVC) images
- OS/2 bitmap
- Windows 3.x bitmap
- PCX
- Intel and Motorola TIFF
- TARGA TGA
- TIFF (FAX)
- GIF
- Kodak Photo-CD

Product Positioning

The advanced image search technology in Ultimedia Manager, query by image content, was developed in conjunction with the IBM Almaden Research Center. It provides function not available with conventional, text-only image management systems.

Ultimedia Manager fulfills image search and retrieval needs of business users and creative professionals through use of online image libraries or collections. It is intended for users who have large image libraries and need fast, flexible access to their image data, users who may also want to add character or numeric data to their image queries, and users who feel limited by text-based systems.

Ultimedia Manager 1.1 works with the DB2 Family of products and the Visualizer Ultimedia Query product. Catalogs created by Utlimedia Manager can be stored in DB2 for OS/2 database tables. Image catalogs also can be stored in DB2 for OS/2 or DB2 for AIX database servers in client/server environments.

OS/2 and Windows clients in a LAN environment using the Client Search Feature and the Visualizer Ultimedia Query can access catalogs created with Ultimedia Manager 1.1. Visualizer Ultimedia Query provides users with the capability to:

- Add multimedia columns (such as image, video, audio, and formatted text) to DB2 for OS/2 database tables
- Query and display images
- View or listen to video and audio information
- Display formatted text

Users of Ultimedia Manager 1.1 and Client Search Feature are focused on image management and query applications. Visualizer Ultimedia Query users are

Element	'278 Claim Recitation	Exemplary Citations in Reference
		To the extent it is found that QBIC does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that QBIC does not anticipate this claim, QBIC renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine QBIC with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results. For example, the '278 patent acknowledges that acquiring information about objects through symbols, such as barcodes or some other means to encode detectable information in the image or on the object, were traditional methods known in the art and not part the present invention. '278 Patent, 3:39-48. Indeed, the '278 patent acknowledges that algorithms, software, and hardware components were commercially available at the time of the invention that could detect and decode symbols, such as barcodes or text, in an input image. *Id*. at 14:54-59. This limitation was thus known in the art by patentee's own admission, and by asserting this claim, Plaintiff is attempting to reclaim disavowed claim scope and claim subject matter which patentee has acknowledged is not part of the invention.
[1C]	extracting symbol information based on the symbolic content according to symbol type;	QBIC discloses extracting symbol information based on the symbolic content according to symbol type. Query by Image and Video Content: The QBIC System (IBM 0000893–902): "We have developed the QBIC (Query by Image Content) system to explore

Element	'278 Claim Recitation	Exemplary Citations in Reference
		content-based retrieval methods. QBIC allows queries on large image and video databases based on • example images, • user-constructed sketches and drawings, • selected color and texture patterns, • camera and object motion, and • other graphical information." *Id.* at pp. 1–3. "INTEGRATION WITH TEXT AND PARAMETRIC ANNOTATION. Query by visual content complements and extends existing query methods. Systems must be able to integrate queries combining date, subject matter, price, and availability with content properties such as color, texture, and shape." *Id.* at p. 8. Image, Audio, and Video Extenders Administration and Programming (IBM 000001–617): "DB2 (DB2) Universal Database (UDB) is a powerful, object-relational database manager. It stores and protects traditional numeric and character data, as well as large, complex objects (LOBs). DB2 extenders help you exploit DB2 's object-relational features. The extenders define distinct data types and special functions for image, audio, video, and text objects. … Because of that, the extenders give your applications a single point of access to any or all of these types of data, along with traditional numeric and character data." *Id.* at p. 3. Managing Enterprise Information Portal (IBM 000777–880): "Many paper-intensive enterprises, such as insurance companies and financial institutions, administer large volumes of business-related content. The need for an enterprise solution for managing and accessing business information spans many industries. A content server is a software system that stores multimedia, business forms, documents, and related data, along with metadata that allows employees to process and work with the content." *Id.* at p. 3. "The image search server uses IBM's QBIC (query by image content) technology to help you search for objects by certain visual properties, such as

Element	'278 Claim Recitation	Exemplary Citations in Reference
		color and texture. The image search server analyzes images and stores the image information in a database. Then users can run image queries, which use the visual properties of images, to match colors, textures, and their positions without describing them in words. You can combine content-based queries with text and keyword searches for more powerful retrieval of image and multimedia data." *Id.* at p. 63. QBIC Demo (IBM000747)  Location: http://172.16.85.176:3456/cgi-bin/ibm.tcl Keywords: france Flickner Depo at 185

Element	'278 Claim Recitation	Exemplary Citations in Reference
		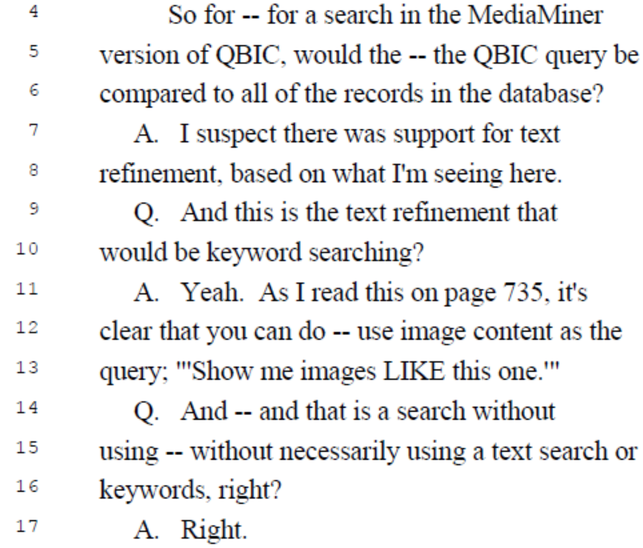 So for -- for a search in the MediaMiner version of QBIC, would the -- the QBIC query be compared to all of the records in the database? A. I suspect there was support for text refinement, based on what I'm seeing here. Q. And this is the text refinement that would be keyword searching? A. Yeah. As I read this on page 735, it's clear that you can do -- use image content as the query; "'Show me images LIKE this one.'" Q. And -- and that is a search without using -- without necessarily using a text search or keywords, right? A. Right. Flickner Depo at 217 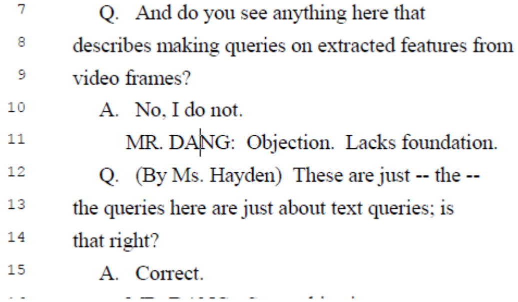 Q. And do you see anything here that describes making queries on extracted features from video frames? A. No, I do not. MR. DANG: Objection. Lacks foundation. Q. (By Ms. Hayden) These are just -- the -- the queries here are just about text queries; is that right? A. Correct. Flickner Depo at 25-26

Element	'278 Claim Recitation	Exemplary Citations in Reference
		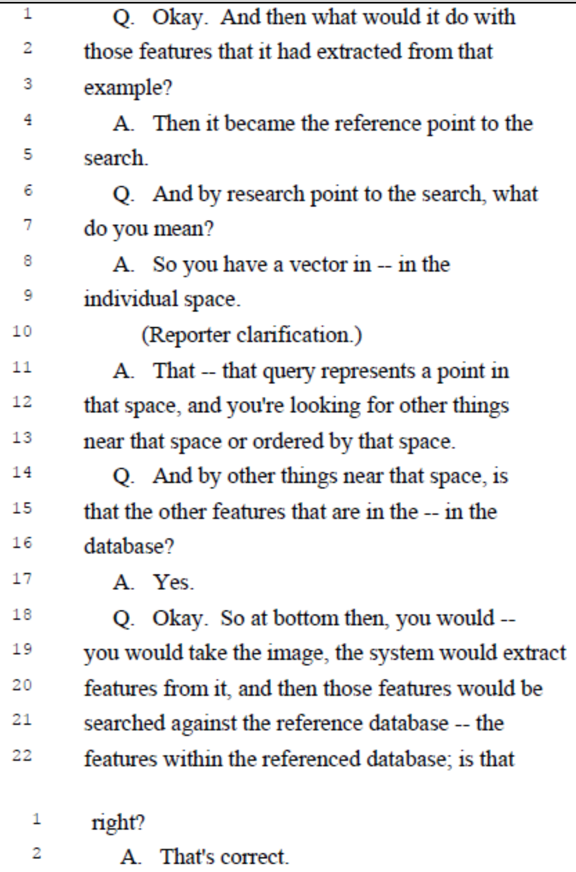 Q. Okay. And then what would it do with those features that it had extracted from that example? A. Then it became the reference point to the search. Q. And by research point to the search, what do you mean? A. So you have a vector in -- in the individual space. (Reporter clarification.) A. That -- that query represents a point in that space, and you're looking for other things near that space or ordered by that space. Q. And by other things near that space, is that the other features that are in the -- in the database? A. Yes. Q. Okay. So at bottom then, you would -- you would take the image, the system would extract features from it, and then those features would be searched against the reference database -- the features within the referenced database; is that right? A. That's correct. Flickner 8/29/31 Depo at 44

Element	'278 Claim Recitation	Exemplary Citations in Reference
		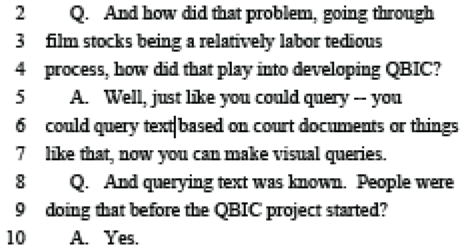 Q. And how did that problem, going through film stocks being a relatively labor tedious process, how did that play into developing QBIC? A. Well, just like you could query -- you could query text based on court documents or things like that, now you can make visual queries. Q. And querying text was known. People were doing that before the QBIC project started? A. Yes. *Id.* at 45–47 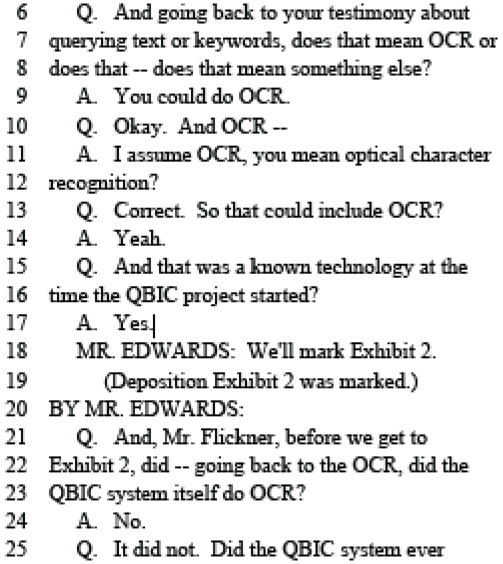 Q. And going back to your testimony about querying text or keywords, does that mean OCR or does that -- does that mean something else? A. You could do OCR. Q. Okay. And OCR -- A. I assume OCR, you mean optical character recognition? Q. Correct. So that could include OCR? A. Yeah. Q. And that was a known technology at the time the QBIC project started? A. Yes. MR. EDWARDS: We'll mark Exhibit 2. (Deposition Exhibit 2 was marked.) BY MR. EDWARDS: Q. And, Mr. Flickner, before we get to Exhibit 2, did -- going back to the OCR, did the QBIC system itself do OCR? A. No. Q. It did not. Did the QBIC system ever

Element	'278 Claim Recitation	Exemplary Citations in Reference
		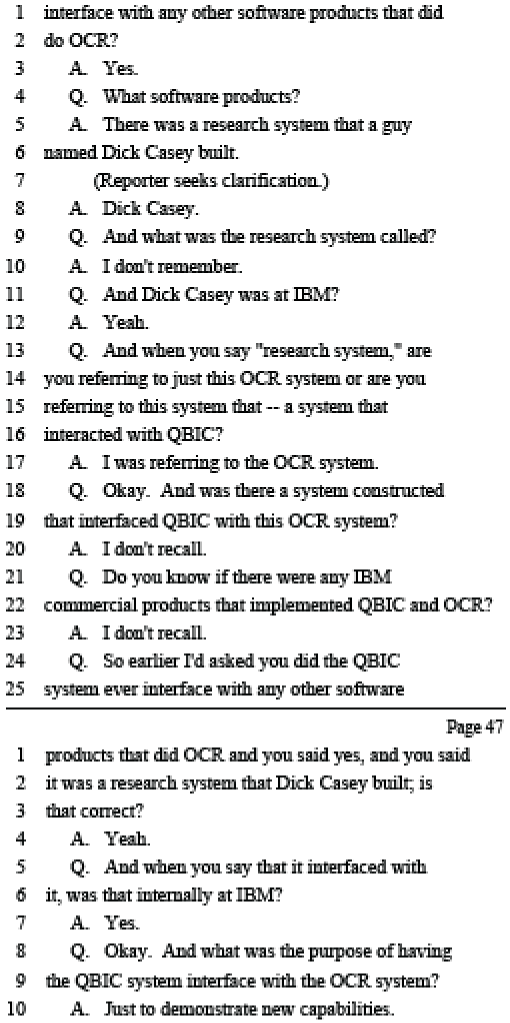 interface with any other software products that did do OCR? A. Yes. Q. What software products? A. There was a research system that a guy named Dick Casey built. (Reporter seeks clarification.) A. Dick Casey. Q. And what was the research system called? A. I don't remember. Q. And Dick Casey was at IBM? A. Yeah. Q. And when you say "research system," are you referring to just this OCR system or are you referring to this system that -- a system that interacted with QBIC? A. I was referring to the OCR system. Q. Okay. And was there a system constructed that interfaced QBIC with this OCR system? A. I don't recall. Q. Do you know if there were any IBM commercial products that implemented QBIC and OCR? A. I don't recall. Q. So earlier I'd asked you did the QBIC system ever interface with any other software Page 47 products that did OCR and you said yes, and you said it was a research system that Dick Casey built; is that correct? A. Yeah. Q. And when you say that it interfaced with it, was that internally at IBM? A. Yes. Q. Okay. And what was the purpose of having the QBIC system interface with the OCR system? A. Just to demonstrate new capabilities. *Id.* at 76

Element	'278 Claim Recitation	Exemplary Citations in Reference
		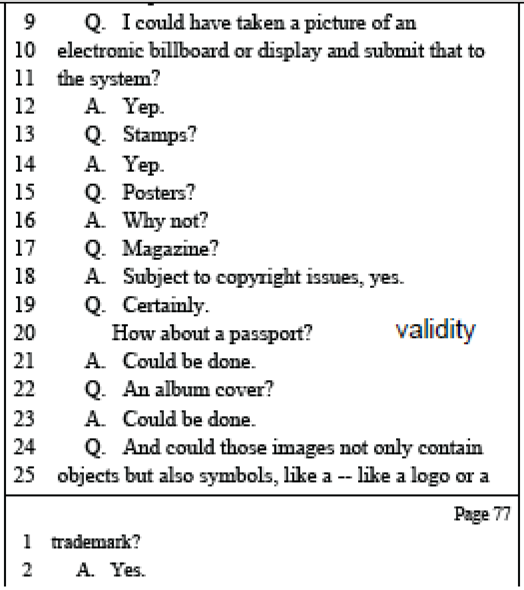 Q. I could have taken a picture of an electronic billboard or display and submit that to the system? A. Yep. Q. Stamps? A. Yep. Q. Posters? A. Why not? Q. Magazine? A. Subject to copyright issues, yes. Q. Certainly. How about a passport? validity A. Could be done. Q. An album cover? A. Could be done. Q. And could those images not only contain objects but also symbols, like a -- like a logo or a Page 77 trademark? A. Yes. *Id.* at 187–188 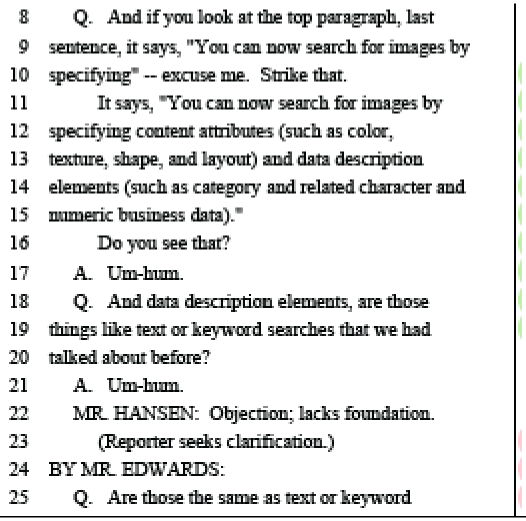 Q. And if you look at the top paragraph, last sentence, it says, "You can now search for images by specifying" -- excuse me. Strike that. It says, "You can now search for images by specifying content attributes (such as color, texture, shape, and layout) and data description elements (such as category and related character and numeric business data)." Do you see that? A. Um-hum. Q. And data description elements, are those things like text or keyword searches that we had talked about before? A. Um-hum. MR. HANSEN: Objection; lacks foundation. (Reporter seeks clarification.) BY MR. EDWARDS: Q. Are those the same as text or keyword

Element	'278 Claim Recitation	Exemplary Citations in Reference
		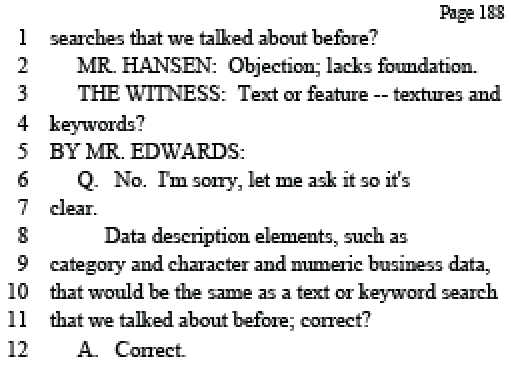Page 188 searches that we talked about before? MR. HANSEN: Objection; lacks foundation. THE WITNESS: Text or feature -- textures and keywords? BY MR. EDWARDS: Q. No. I'm sorry, let me ask it so it's clear. Data description elements, such as category and character and numeric business data, that would be the same as a text or keyword search that we talked about before; correct? A. Correct. IBM0002263 at 2264

Description

Ultimedia Manager 1.1 opens new dimensions in visual queries, helping you manage your image database more efficiently. With its easy-to-use interface, even entry-level users can perform sophisticated searches. You can search by image content, data description, or a combination of both. Once the images are found, you can view the images in full resolution, use the images in another application, or read the associated textual comments.

You can use Ultimedia Manager 1.1 either as a stand-alone product or in a LAN environment. As a stand-alone product, Ultimedia Manager 1.1 employs full image classification and image query capabilities. Within the LAN environment, the Ultimedia Manager image catalogs can be sorted on a DB2 for OS/2 or DB2 for AIX database server. The image catalogs can be accessed by Windows or OS/2 clients using the Client Search Feature or the Visualizer Ultimedia Query.

Ultimedia Manager runs as an OS/2 application to a DB2 for OS/2 or DB2 for AIX database server in a LAN, or as a stand-alone application with a local DB2 for OS/2 database. The Client Search Feature runs on DOS or Windows, and accesses files in dBASE flat file format. A database manager is not required to access catalogs in dBASE file format.

Image Classification

With the Image Classifier, you can create visual indexes based on color, shape, category, and layout elements, which are areas of the image containing the patterns you want to use for searches. You can store these elements in catalogs in DB2 for OS/2 databases. The catalogs can be used locally or, through the use of utility programs, be transferred to DB2 for OS/2 or AIX servers. In addition, you can convert the catalog data into dBASE flat file format for access by Windows client applications.

The Image Classifier and Analyzer allows automatic or interactive classification of images for later image queries, flexibility in storage of catalog information, and access to multiple users. You can organize large image catalogs so that later searches for images can be performed easily. You can classify automatically using a batch program, or interactively by specifying areas of interest in an image and indicating the visual cue (color, shape, or texture) to be saved for later queries.

Image Query

Once the image catalogs are created, you can use the Image Query component to find image and business information. Business information, such as character and numeric data, are stored in the database tables using normal table population methods. Image Query is provided in Ultimedia Manager and Client Search Feature.

Image Query allows searches by three methods:

- Image content, such as color, shape, texture, or layout
- Data description, such as business data stored in the database (character or numeric) or descriptive text (category)
- Combinations of image content and data description.

Image Content: To search by image content, specify the attributes (color, shape, texture, or layout). You can drag and drop visual samples into an image query, select a color from the color wheel, or sketch a shape.

Data Description: To search by data description, enter character and numeric data values in the database tables using the simple fill-in-the-blanks query panel. You can enter data such as name, number, date, or a range of dates. You can integrate character or numeric data as part of the image query, enabling a more effective search for better results. The simple fill-in-the-blanks method for specifying desired conditions is so easy even non-programmers can use database table information as part of their searches.

Combinations: Image queries using either image content or data description, or combinations of both, allow you to search through thousands of images more quickly and effectively than the traditional text and keyword search methods.

Image Formats

Ultimedia Manager 1.1 understands these image formats:

- Audio Visual Connection (AVC) images
- OS/2 bitmap
- Windows 3.x bitmap
- PCX
- Intel and Motorola TIFF
- TARGA TGA
- TIFF (FAX)
- GIF
- Kodak Photo-CD

Product Positioning

The advanced image search technology in Ultimedia Manager, query by image content, was developed in conjunction with the IBM Almaden Research Center. It provides function not available with conventional, text-only image management systems.

Ultimedia Manager fulfills image search and retrieval needs of business users and creative professionals through use of online image libraries or collections. It is intended for users who have large image libraries and need fast, flexible access to their image data, users who may also want to add character or numeric data to their image queries, and users who feel limited by text-based systems.

Ultimedia Manager 1.1 works with the DB2 Family of products and the Visualizer Ultimedia Query product. Catalogs created by Utlimedia Manager can be stored in DB2 for OS/2 database tables. Image catalogs also can be stored in DB2 for OS/2 or DB2 for AIX database servers in client/server environments.

OS/2 and Windows clients in a LAN environment using the Client Search Feature and the Visualizer Ultimedia Query can access catalogs created with Ultimedia Manager 1.1. Visualizer Ultimedia Query provides users with the capability to:

- Add multimedia columns (such as image, video, audio, and formatted text) to DB2 for OS/2 database tables
- Query and display images
- View or listen to video and audio information
- Display formatted text

Users of Ultimedia Manager 1.1 and Client Search Feature are focused on image management and query applications. Visualizer Ultimedia Query users are

Element	'278 Claim Recitation	Exemplary Citations in Reference
		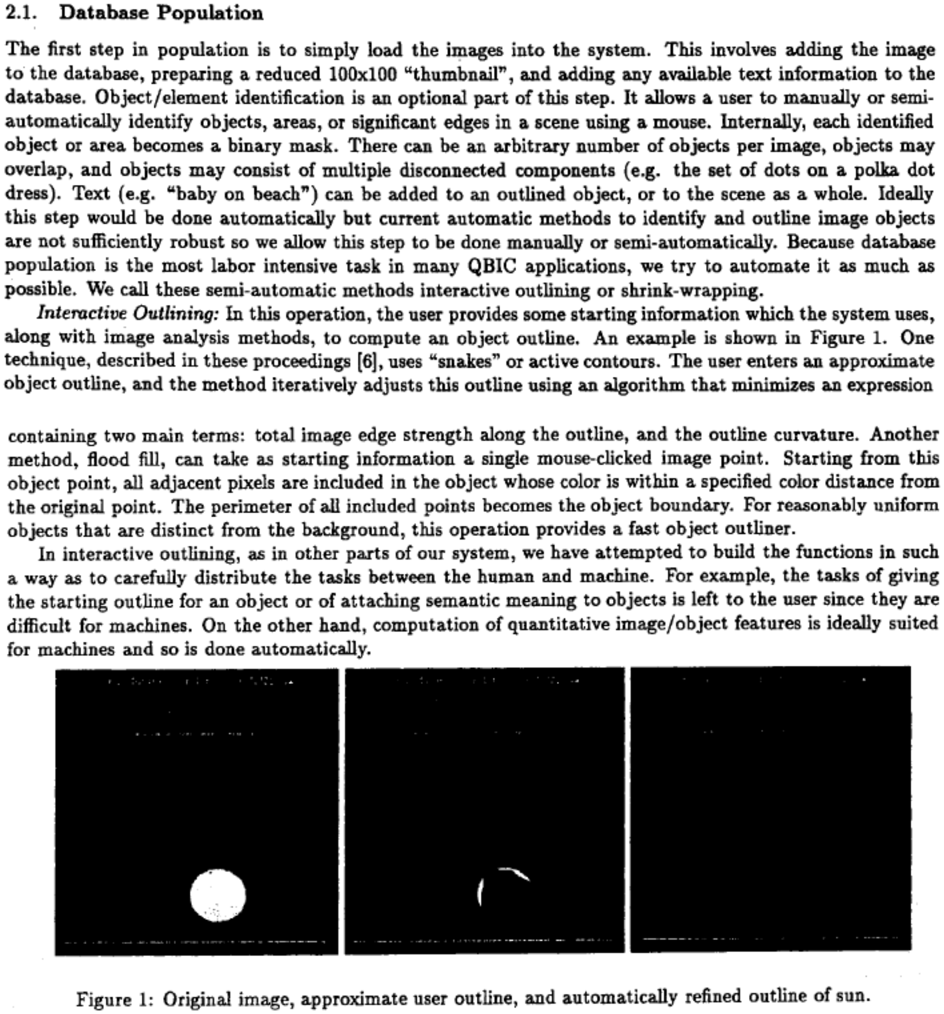 **2.1. Database Population** The first step in population is to simply load the images into the system. This involves adding the image to the database, preparing a reduced 100x100 "thumbnail", and adding any available text information to the database. Object/element identification is an optional part of this step. It allows a user to manually or semi-automatically identify objects, areas, or significant edges in a scene using a mouse. Internally, each identified object or area becomes a binary mask. There can be an arbitrary number of objects per image, objects may overlap, and objects may consist of multiple disconnected components (e.g. the set of dots on a polka dot dress). Text (e.g. "baby on beach") can be added to an outlined object, or to the scene as a whole. Ideally this step would be done automatically but current automatic methods to identify and outline image objects are not sufficiently robust so we allow this step to be done manually or semi-automatically. Because database population is the most labor intensive task in many QBIC applications, we try to automate it as much as possible. We call these semi-automatic methods interactive outlining or shrink-wrapping. *Interactive Outlining:* In this operation, the user provides some starting information which the system uses, along with image analysis methods, to compute an object outline. An example is shown in Figure 1. One technique, described in these proceedings [6], uses "snakes" or active contours. The user enters an approximate object outline, and the method iteratively adjusts this outline using an algorithm that minimizes an expression containing two main terms: total image edge strength along the outline, and the outline curvature. Another method, flood fill, can take as starting information a single mouse-clicked image point. Starting from this object point, all adjacent pixels are included in the object whose color is within a specified color distance from the original point. The perimeter of all included points becomes the object boundary. For reasonably uniform objects that are distinct from the background, this operation provides a fast object outliner. In interactive outlining, as in other parts of our system, we have attempted to build the functions in such a way as to carefully distribute the tasks between the human and machine. For example, the tasks of giving the starting outline for an object or of attaching semantic meaning to objects is left to the user since they are difficult for machines. On the other hand, computation of quantitative image/object features is ideally suited for machines and so is done automatically. Figure 1: Original image, approximate user outline, and automatically refined outline of sun. *See also* Flickner Depo at 22-25, 40-42, 54-55, 87, 98-90.

Element	'278 Claim Recitation	Exemplary Citations in Reference
		To the extent it is found that QBIC does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that QBIC does not anticipate this claim, QBIC renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine QBIC with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results. For example, the '278 patent acknowledges that acquiring information about objects through symbols, such as barcodes or some other means to encode detectable information in the image or on the object, were traditional methods known in the art and not part the present invention. '278 Patent, 3:39-48. Indeed, the '278 patent acknowledges that algorithms, software, and hardware components were commercially available at the time of the invention that could detect and decode symbols, such as barcodes or text, in an input image. *Id.* at 14:54-59. This limitation was thus known in the art by patentee's own admission, and by asserting this claim, Plaintiff is attempting to reclaim disavowed claim scope and claim subject matter which patentee has acknowledged is not part of the invention.
[1D]	determining a validity of the document based at least in part on the image and the symbol information; and	QBIC discloses determining a validity of the document based at least in part on the image and the symbol information. IBM0002390 at 2391-92:

Element	'278 Claim Recitation	Exemplary Citations in Reference
		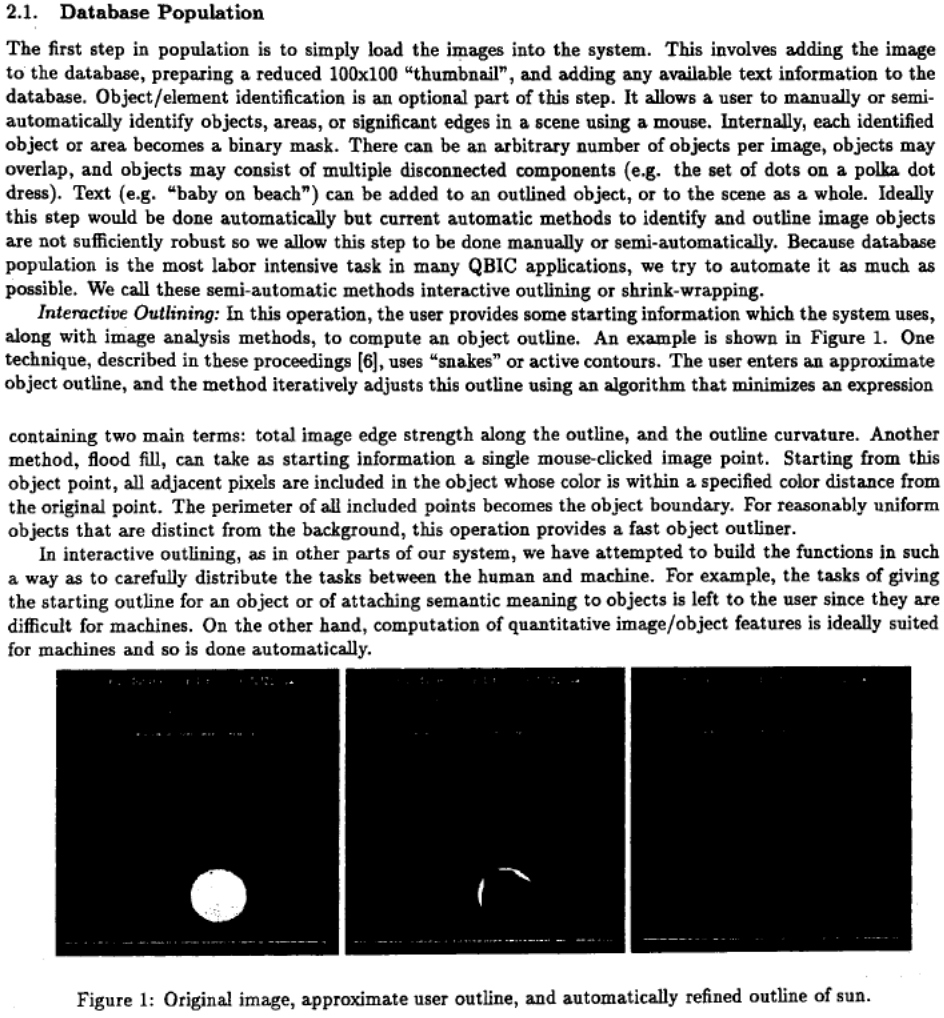 2.1. **Database Population** The first step in population is to simply load the images into the system. This involves adding the image to the database, preparing a reduced 100x100 "thumbnail", and adding any available text information to the database. Object/element identification is an optional part of this step. It allows a user to manually or semi-automatically identify objects, areas, or significant edges in a scene using a mouse. Internally, each identified object or area becomes a binary mask. There can be an arbitrary number of objects per image, objects may overlap, and objects may consist of multiple disconnected components (e.g. the set of dots on a polka dot dress). Text (e.g. "baby on beach") can be added to an outlined object, or to the scene as a whole. Ideally this step would be done automatically but current automatic methods to identify and outline image objects are not sufficiently robust so we allow this step to be done manually or semi-automatically. Because database population is the most labor intensive task in many QBIC applications, we try to automate it as much as possible. We call these semi-automatic methods interactive outlining or shrink-wrapping. *Interactive Outlining:* In this operation, the user provides some starting information which the system uses, along with image analysis methods, to compute an object outline. An example is shown in Figure 1. One technique, described in these proceedings [6], uses "snakes" or active contours. The user enters an approximate object outline, and the method iteratively adjusts this outline using an algorithm that minimizes an expression containing two main terms: total image edge strength along the outline, and the outline curvature. Another method, flood fill, can take as starting information a single mouse-clicked image point. Starting from this object point, all adjacent pixels are included in the object whose color is within a specified color distance from the original point. The perimeter of all included points becomes the object boundary. For reasonably uniform objects that are distinct from the background, this operation provides a fast object outliner. In interactive outlining, as in other parts of our system, we have attempted to build the functions in such a way as to carefully distribute the tasks between the human and machine. For example, the tasks of giving the starting outline for an object or of attaching semantic meaning to objects is left to the user since they are difficult for machines. On the other hand, computation of quantitative image/object features is ideally suited for machines and so is done automatically. Figure 1: Original image, approximate user outline, and automatically refined outline of sun. Query by Image and Video Content: The QBIC System (IBM 0000893–902):

Element	'278 Claim Recitation	Exemplary Citations in Reference
		"We have developed the QBIC (Query by Image Content) system to explore content-based retrieval methods. QBIC allows queries on large image and video databases based on • example images, • user-constructed sketches and drawings, • selected color and texture patterns, • camera and object motion, and • other graphical information." *Id.* at pp. 1–3. "**SAMPLE QUERIES** For each full-scene image, identified image object, r-frame, and identified video object resulting from the above processing, a set of features is computed to allow content-based queries. The features are computed and stored during database population. We present a brief description of the features and the associated queries. Mathematical details on the features and matching methods can be found in Ashley et al. and Niblack et al. Average color queries let users find images or objects that are similar to a selected color, say from a color wheel, or to the color of an object. The feature used in the query is a 3D vector of Munsell color coordinates. Histogram color queries return items with matching color distributions—say, a fabric pattern with approximately 40 percent red and 20 percent blue. For this case, the underlying feature is a 256-element histogram computed over a quantized version of the color space. Figure 7 shows a histogram query on still images and a color query on video r-frames. Note that in the query specification for the histogram query of Figure 7, the user has selected percentages of two colors (blue and white) by adjusting sliders. Using such a query, an advertising agent could, for example, search for a picture of a beach scene, one predominantly blue (for sky and water) and white (for sand and clouds); or find images with similar color spreads for a uniform ad campaign. The average color query demonstrates a query against a video shot database where the user is searching for red r-frames. Again, the query specification is on the left and the best hits are on the right. Figure 8 shows an example texture query. In this case, the query is specified by selecting from a sampler—a set of prestored example images. The underlying

Element	'278 Claim Recitation	Exemplary Citations in Reference
		texture features are mathematical representations of coarseness, contrast, and directionality features. Coarseness measures the scale of a texture (pebbles versus boulders), contrast describes its vividness, and directionality describes whether it has a favored direction (like grass) or not (like a smooth object). An object shape query is shown in Figure 8. In this case, the query specification is the drawn shape on the left. Area, circularity, eccentricity, major-axis direction, features derived from the object moments, and a set of tangent angles around the object perimeter are the features used to characterize and match shapes. Figure 9 illustrates query by sketch. In this case, the query specification is a freehand drawing of the dominant lines and edges in the image. The sketch feature is an automatically extracted reduced-resolution 'edge map.' Matching is done by using a template-matching technique. A multiobject query asking for images that contain both a red round object and a green textured object is shown in the bottom of Figure 9. The features are standard color and texture. The matching is done by combining the color and texture distances. Combining distances is applied to arbitrary sets of objects and features to implement logical And semantics. . . . . **ANNOTATION AND DATABASE POPULATION TOOLS**. Automatic methods (such as our Positional Color query) that don't rely on object identification, methods that identify objects automatically as in the museum image example, fast and easy-to-use semiautomatic outlining tools, and motion-based segmentation algorithms will enable additional application areas. **FEATURE EXTRACTION AND MATCHING METHODS**. New mathematical representations of video, image, and object attributes that capture 'interesting' features for retrieval are needed. Features that describe new image properties such as alternate texture measures or that are based on fractals or wavelet representations, for example, may offer advantages of representation, indexability, and ease of similarity matching. **INTEGRATION WITH TEXT AND PARAMETRIC ANNOTATION**.

Element	'278 Claim Recitation	Exemplary Citations in Reference
		Query by visual content complements and extends existing query methods. Systems must be able to integrate queries combining date, subject matter, price, and availability with content properties such as color, texture, and shape." *Id.* at pp. 7–8. Image, Audio, and Video Extenders Administration and Programming (IBM 000001–617): "DB2 (DB2) Universal Database (UDB) is a powerful, object-relational database manager. It stores and protects traditional numeric and character data, as well as large, complex objects (LOBs). DB2 extenders help you exploit DB2 's object-relational features. The extenders define distinct data types and special functions for image, audio, video, and text objects. … Because of that, the extenders give your applications a single point of access to any or all of these types of data, along with traditional numeric and character data." *Id.* at p. 3. "A trigger defines a set of actions that are activated by a change to a table. Triggers can be used to perform actions such as validating input data, automatically generating a value for a newly inserted row, reading from other tables for cross-referencing purposes, or writing to other tables for auditing purposes. Triggers are often used for integrity checking or to enforce business rules." *Id.* at 16–17. Managing Enterprise Information Portal (IBM 000777–880): "Many paper-intensive enterprises, such as insurance companies and financial institutions, administer large volumes of business-related content. The need for an enterprise solution for managing and accessing business information spans many industries. A content server is a software system that stores multimedia, business forms, documents, and related data, along with metadata that allows employees to process and work with the content." *Id.* at p. 3. "The image search server uses IBM's QBIC (query by image content)

Element	'278 Claim Recitation	Exemplary Citations in Reference
		technology to help you search for objects by certain visual properties, such as color and texture. The image search server analyzes images and stores the image information in a database. Then users can run image queries, which use the visual properties of images, to match colors, textures, and their positions without describing them in words. You can combine content-based queries with text and keyword searches for more powerful retrieval of image and multimedia data." *Id*. at p. 63. Flickner Depo at 22-23 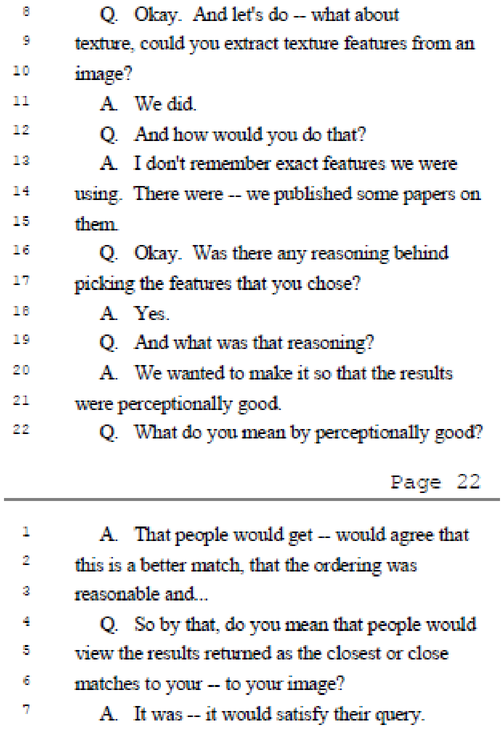 Q. Okay. And let's do -- what about texture, could you extract texture features from an image? A. We did. Q. And how would you do that? A. I don't remember exact features we were using. There were -- we published some papers on them. Q. Okay. Was there any reasoning behind picking the features that you chose? A. Yes. Q. And what was that reasoning? A. We wanted to make it so that the results were perceptionally good. Q. What do you mean by perceptionally good? Page 22 A. That people would get -- would agree that this is a better match, that the ordering was reasonable and... Q. So by that, do you mean that people would view the results returned as the closest or close matches to your -- to your image? A. It was -- it would satisfy their query. *Id.* at 33–37

Element	'278 Claim Recitation	Exemplary Citations in Reference
		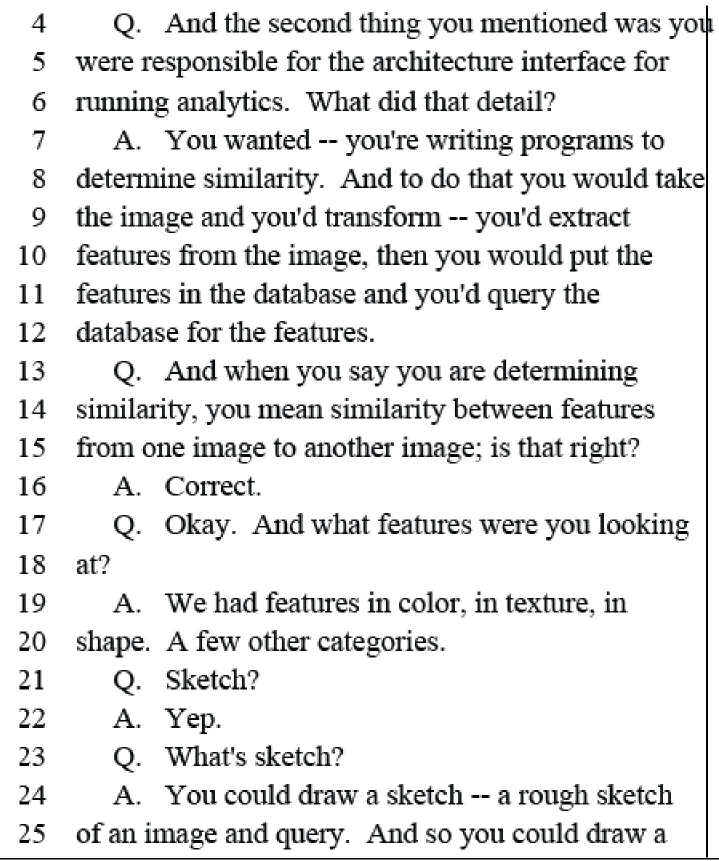 Q. And the second thing you mentioned was you were responsible for the architecture interface for running analytics. What did that detail? A. You wanted -- you're writing programs to determine similarity. And to do that you would take the image and you'd transform -- you'd extract features from the image, then you would put the features in the database and you'd query the database for the features. Q. And when you say you are determining similarity, you mean similarity between features from one image to another image; is that right? A. Correct. Q. Okay. And what features were you looking at? A. We had features in color, in texture, in shape. A few other categories. Q. Sketch? A. Yep. Q. What's sketch? A. You could draw a sketch -- a rough sketch of an image and query. And so you could draw a

		Page 34 circle for a sun, a line and maybe some trees and then you could send that sketch in and you'd get images of a sun in the trees. Q. And if I refer to images, can I -- let me -- let me back up. Earlier you talked about both still images and video were part of your responsibilities for these architecture interfaces; correct? A. Correct. Q. Okay. If I refer to "images," do you understand that is including both a still image or, for example, a frame of a video image? A. Sure. Q. Okay. So when I say "image," I'm including both in my questions from now on out. A. Okay. Q. Okay? A. Okay. Q. In terms of running analytics, that would include both extracting these characteristics from a query image and from an image that you're storing in a database; correct? A. Can you rephrase. Q. Yes. So the comparison that you're doing between two images; right, is it correct that that Page 35 comparison is for purposes of a query? You're querying image 1 to determine if there are images in the database that have similar features; is that right? A. Correct. Q. Okay. And the analytics that you're talking about, where you extract these characters from the image, that is -- those run on both the query image and the images that are stored in the database; is that right? A. Yeah. We call them features. Q. Features. Okay. So when you say "features," you're talking about the features that we talked about before: Color, texture, shape, sketch? A. For example. Q. Any other ones you can think of? A. Nothing comes to mind. Q. But when you refer to features they at least include those four; is that correct? A. Yes. Well, at least includes, yeah. Q. Is that correct? I'm sorry, I didn't -- you trailed off, so I -- that's why I reasked it. A. Yeah. Q. Let me reask it and then you can give me a Page 36 clean answer. When you refer to features, does it include at least color, texture, shape and sketch? A. Or a subset thereof. (Reporter seeks clarification.) A. Or a subset thereof. So you could do just color queries, you could do just shape queries, you could do just -- you didn't need to have all the feature sets in all -- in every query. Q. Understood. But would a user have the option to query at least those four features? A. Yes. Q. Okay. A. For their QBIC system, that's true. Q. And then I believe the last area you indicated you were responsible for was architecture interface for saving images; is that right? A. Correct. Q. And what did that entail? A. It's both saving images and saving features. That was to define the API, so you would give it a feature, and it would write the feature in a database or into a file. Q. Would that also include saving any Page 37 information associated with the images, for example, like any text description about the image? A. Yeah. Text is just another feature here. Q. Text is another feature? A. Yeah. Q. I understand when the images were saved, thumbnails were saved of those images; is that correct? A. Typically, yeah. Q. And was that a part of your responsibilities, to architect the interface for saving images? A. Saving . . . hmm. Can you rephrase. Q. What's -- what was confusing about my question? A. I just need you to repeat it. Q. Oh. One of the areas you mentioned that you were responsible for was to architect the interface for saving images. My question is, as part of that responsibility, did that include architecting the features related to saving thumbnails associated with those images? A. The thumbnail is an extracted feature. Q. Okay. So is that a "yes"?

Element	'278 Claim Recitation	Exemplary Citations in Reference
		Flickner Depo at 143 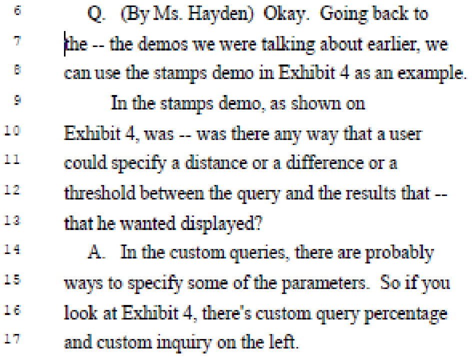 Q. (By Ms. Hayden) Okay. Going back to the -- the demos we were talking about earlier, we can use the stamps demo in Exhibit 4 as an example. In the stamps demo, as shown on Exhibit 4, was -- was there any way that a user could specify a distance or a difference or a threshold between the query and the results that -- that he wanted displayed? A. In the custom queries, there are probably ways to specify some of the parameters. So if you look at Exhibit 4, there's custom query percentage and custom inquiry on the left. *Id.* at 76–77 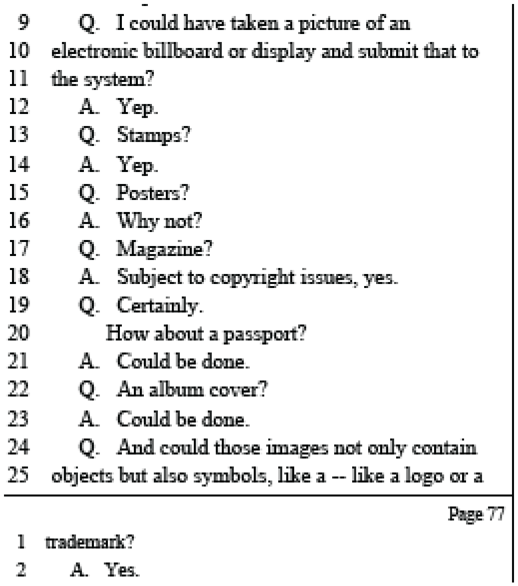 Q. I could have taken a picture of an electronic billboard or display and submit that to the system? A. Yep. Q. Stamps? A. Yep. Q. Posters? A. Why not? Q. Magazine? A. Subject to copyright issues, yes. Q. Certainly. How about a passport? A. Could be done. Q. An album cover? A. Could be done. Q. And could those images not only contain objects but also symbols, like a -- like a logo or a trademark? Page 77 A. Yes. *Id.* at 77–78

Element	'278 Claim Recitation	Exemplary Citations in Reference
		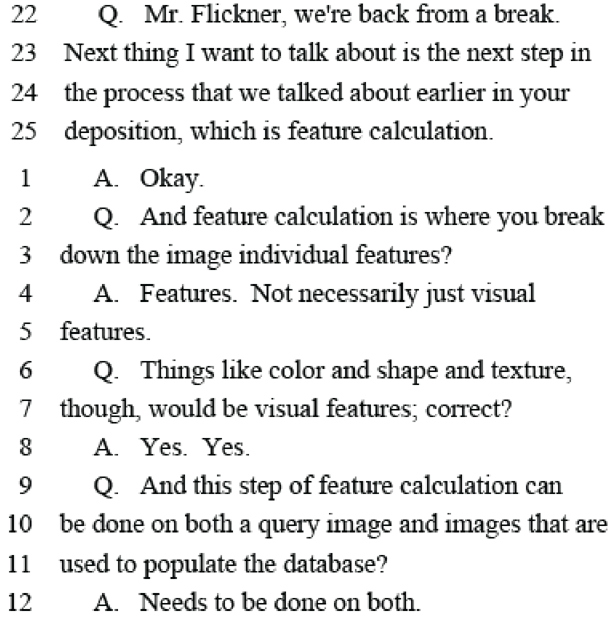 Q. Mr. Flickner, we're back from a break. Next thing I want to talk about is the next step in the process that we talked about earlier in your deposition, which is feature calculation. A. Okay. Q. And feature calculation is where you break down the image individual features? A. Features. Not necessarily just visual features. Q. Things like color and shape and texture, though, would be visual features; correct? A. Yes. Yes. Q. And this step of feature calculation can be done on both a query image and images that are used to populate the database? A. Needs to be done on both. *Id.* at 187–188

Element	'278 Claim Recitation	Exemplary Citations in Reference
		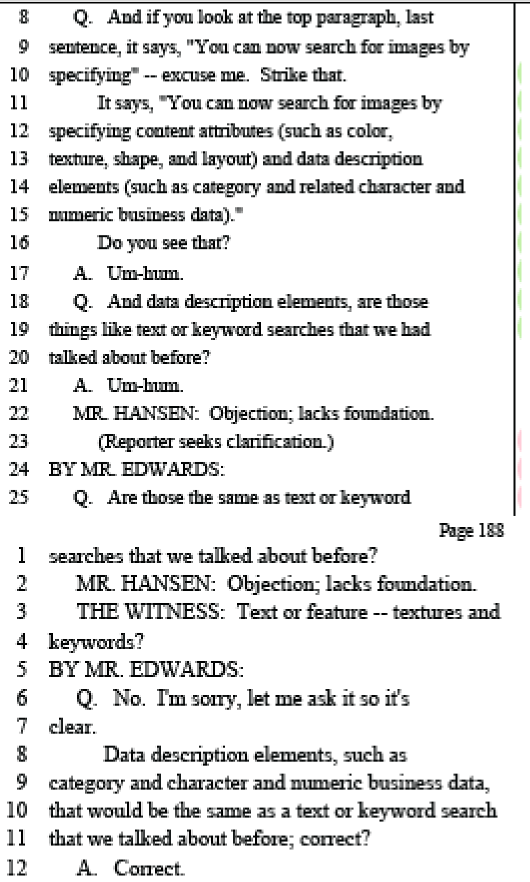 Q. And if you look at the top paragraph, last sentence, it says, "You can now search for images by specifying" -- excuse me. Strike that. It says, "You can now search for images by specifying content attributes (such as color, texture, shape, and layout) and data description elements (such as category and related character and numeric business data)." Do you see that? A. Um-hum. Q. And data description elements, are those things like text or keyword searches that we had talked about before? A. Um-hum. MR. HANSEN: Objection; lacks foundation. (Reporter seeks clarification.) BY MR. EDWARDS: Q. Are those the same as text or keyword Page 188 searches that we talked about before? MR. HANSEN: Objection; lacks foundation. THE WITNESS: Text or feature -- textures and keywords? BY MR. EDWARDS: Q. No. I'm sorry, let me ask it so it's clear. Data description elements, such as category and character and numeric business data, that would be the same as a text or keyword search that we talked about before; correct? A. Correct. *Id.* at 190–191 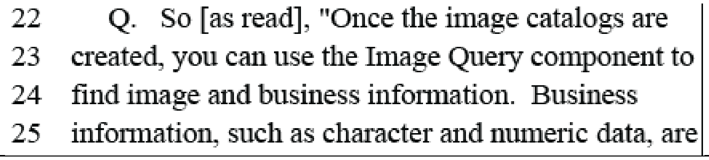 Q. So [as read], "Once the image catalogs are created, you can use the Image Query component to find image and business information. Business information, such as character and numeric data, are

		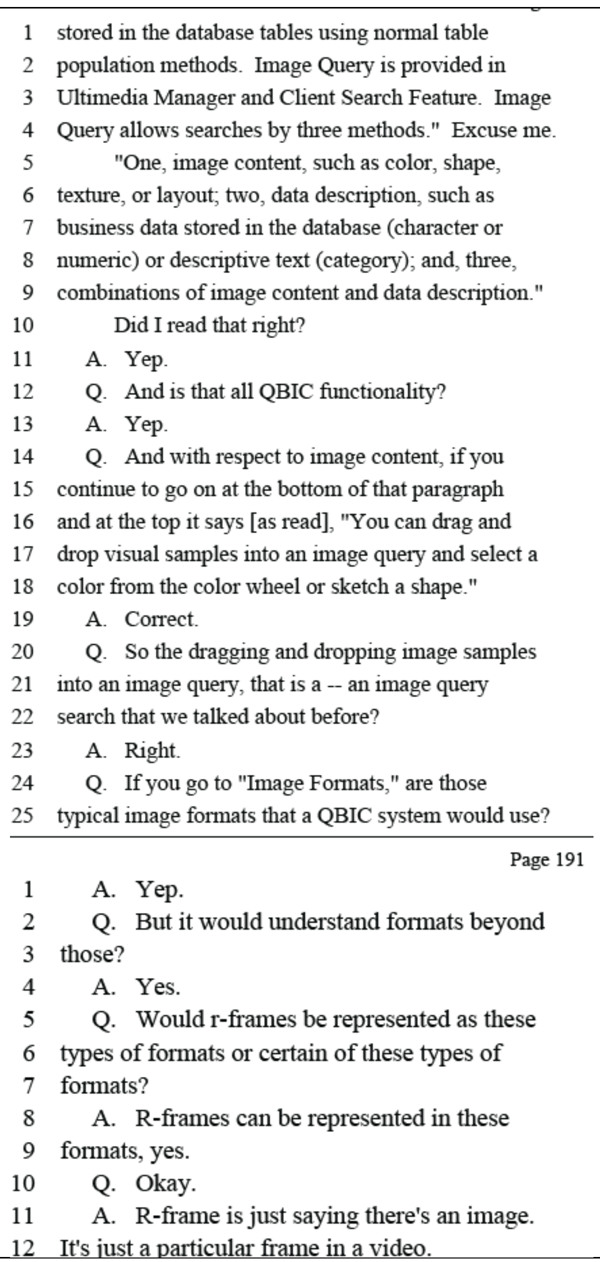stored in the database tables using normal table population methods. Image Query is provided in Ultimedia Manager and Client Search Feature. Image Query allows searches by three methods." Excuse me. "One, image content, such as color, shape, texture, or layout; two, data description, such as business data stored in the database (character or numeric) or descriptive text (category); and, three, combinations of image content and data description." Did I read that right? A. Yep. Q. And is that all QBIC functionality? A. Yep. Q. And with respect to image content, if you continue to go on at the bottom of that paragraph and at the top it says [as read], "You can drag and drop visual samples into an image query and select a color from the color wheel or sketch a shape." A. Correct. Q. So the dragging and dropping image samples into an image query, that is a -- an image query search that we talked about before? A. Right. Q. If you go to "Image Formats," are those typical image formats that a QBIC system would use? Page 191 A. Yep. Q. But it would understand formats beyond those? A. Yes. Q. Would r-frames be represented as these types of formats or certain of these types of formats? A. R-frames can be represented in these formats, yes. Q. Okay. A. R-frame is just saying there's an image. It's just a particular frame in a video.

Element	'278 Claim Recitation	Exemplary Citations in Reference
		See also Flickner Depo at 54-55, 89-90, 185, 217 To the extent it is found that QBIC does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that QBIC does not anticipate this claim, QBIC renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine QBIC with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1E]	recognizing the document as a first target object based at least in part on	QBIC discloses this limitation. The QBIC Project: Querying Images by Content Using Color, Texture, and Shape (IBM 0002390–404) at 2391 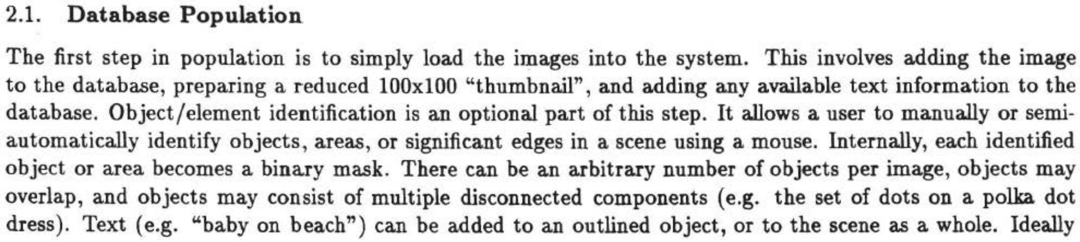 2.1. **Database Population** The first step in population is to simply load the images into the system. This involves adding the image to the database, preparing a reduced 100x100 "thumbnail", and adding any available text information to the database. Object/element identification is an optional part of this step. It allows a user to manually or semi-automatically identify objects, areas, or significant edges in a scene using a mouse. Internally, each identified object or area becomes a binary mask. There can be an arbitrary number of objects per image, objects may overlap, and objects may consist of multiple disconnected components (e.g. the set of dots on a polka dot dress). Text (e.g. "baby on beach") can be added to an outlined object, or to the scene as a whole. Ideally

Element	'278 Claim Recitation	Exemplary Citations in Reference
		The QBIC Project: Querying Images by Content Using Color, Texture, and Shape (IBM 0002390–404) at 2391-92 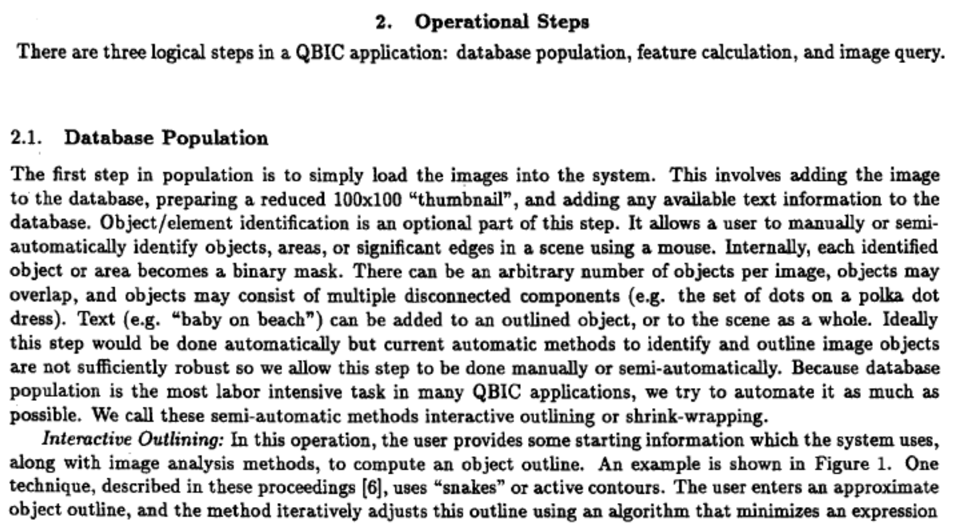 **2. Operational Steps** There are three logical steps in a QBIC application: database population, feature calculation, and image query. **2.1. Database Population** The first step in population is to simply load the images into the system. This involves adding the image to the database, preparing a reduced 100x100 "thumbnail", and adding any available text information to the database. Object/element identification is an optional part of this step. It allows a user to manually or semi-automatically identify objects, areas, or significant edges in a scene using a mouse. Internally, each identified object or area becomes a binary mask. There can be an arbitrary number of objects per image, objects may overlap, and objects may consist of multiple disconnected components (e.g. the set of dots on a polka dot dress). Text (e.g. "baby on beach") can be added to an outlined object, or to the scene as a whole. Ideally this step would be done automatically but current automatic methods to identify and outline image objects are not sufficiently robust so we allow this step to be done manually or semi-automatically. Because database population is the most labor intensive task in many QBIC applications, we try to automate it as much as possible. We call these semi-automatic methods interactive outlining or shrink-wrapping. *Interactive Outlining:* In this operation, the user provides some starting information which the system uses, along with image analysis methods, to compute an object outline. An example is shown in Figure 1. One technique, described in these proceedings [6], uses "snakes" or active contours. The user enters an approximate object outline, and the method iteratively adjusts this outline using an algorithm that minimizes an expression

Element	'278 Claim Recitation	Exemplary Citations in Reference
		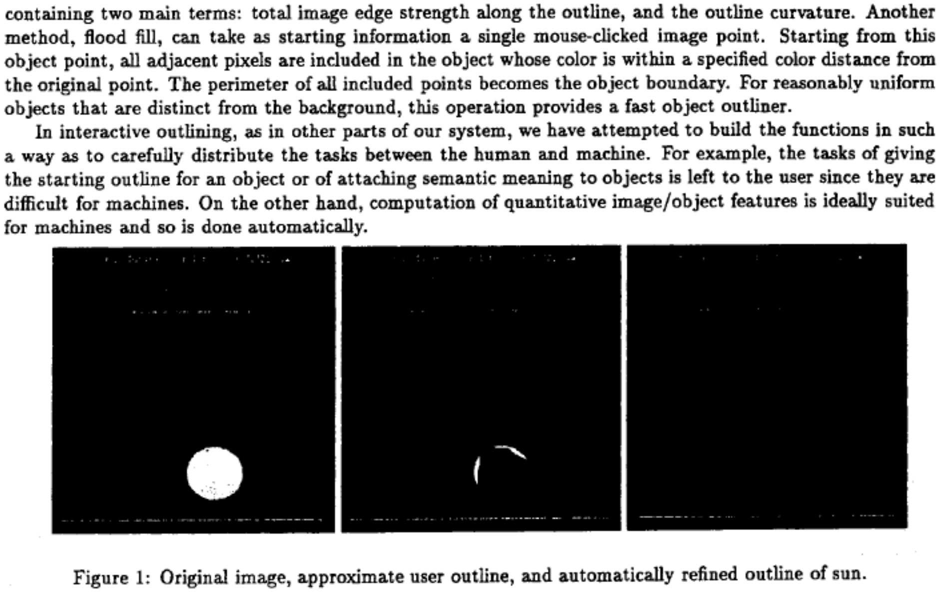 containing two main terms: total image edge strength along the outline, and the outline curvature. Another method, flood fill, can take as starting information a single mouse-clicked image point. Starting from this object point, all adjacent pixels are included in the object whose color is within a specified color distance from the original point. The perimeter of all included points becomes the object boundary. For reasonably uniform objects that are distinct from the background, this operation provides a fast object outliner. In interactive outlining, as in other parts of our system, we have attempted to build the functions in such a way as to carefully distribute the tasks between the human and machine. For example, the tasks of giving the starting outline for an object or of attaching semantic meaning to objects is left to the user since they are difficult for machines. On the other hand, computation of quantitative image/object features is ideally suited for machines and so is done automatically. Figure 1: Original image, approximate user outline, and automatically refined outline of sun. IBM0002390 at 2392-93:

Element	'278 Claim Recitation	Exemplary Citations in Reference
		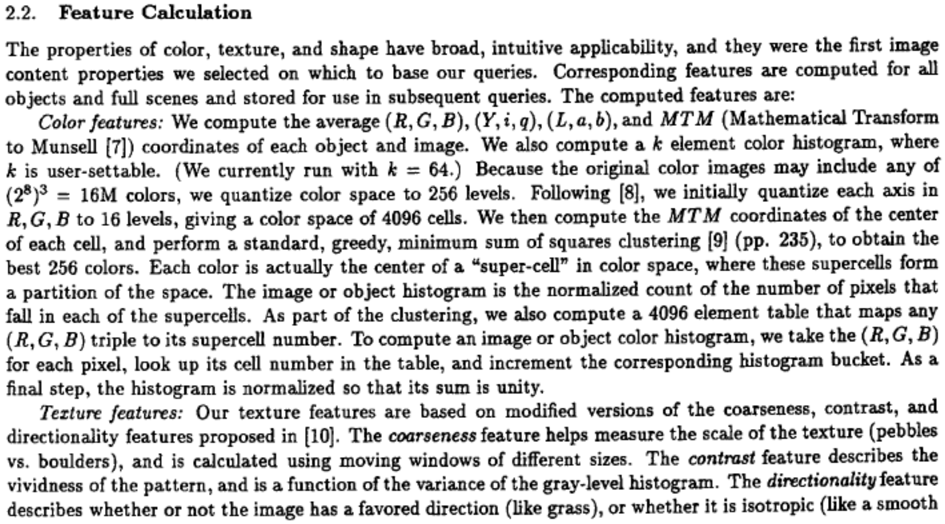 2.2. **Feature Calculation** The properties of color, texture, and shape have broad, intuitive applicability, and they were the first image content properties we selected on which to base our queries. Corresponding features are computed for all objects and full scenes and stored for use in subsequent queries. The computed features are: *Color features:* We compute the average (R, G, B), (Y, i, q), (L, a, b), and MTM (Mathematical Transform to Munsell [7]) coordinates of each object and image. We also compute a k element color histogram, where k is user-settable. (We currently run with $k = 64$.) Because the original color images may include any of $(2^8)^3 = 16$M colors, we quantize color space to 256 levels. Following [8], we initially quantize each axis in R, G, B to 16 levels, giving a color space of 4096 cells. We then compute the MTM coordinates of the center of each cell, and perform a standard, greedy, minimum sum of squares clustering [9] (pp. 235), to obtain the best 256 colors. Each color is actually the center of a "super-cell" in color space, where these supercells form a partition of the space. The image or object histogram is the normalized count of the number of pixels that fall in each of the supercells. As part of the clustering, we also compute a 4096 element table that maps any (R, G, B) triple to its supercell number. To compute an image or object color histogram, we take the (R, G, B) for each pixel, look up its cell number in the table, and increment the corresponding histogram bucket. As a final step, the histogram is normalized so that its sum is unity. *Texture features:* Our texture features are based on modified versions of the coarseness, contrast, and directionality features proposed in [10]. The *coarseness* feature helps measure the scale of the texture (pebbles vs. boulders), and is calculated using moving windows of different sizes. The *contrast* feature describes the vividness of the pattern, and is a function of the variance of the gray-level histogram. The *directionality* feature describes whether or not the image has a favored direction (like grass), or whether it is isotropic (like a smooth

Element	'278 Claim Recitation	Exemplary Citations in Reference
		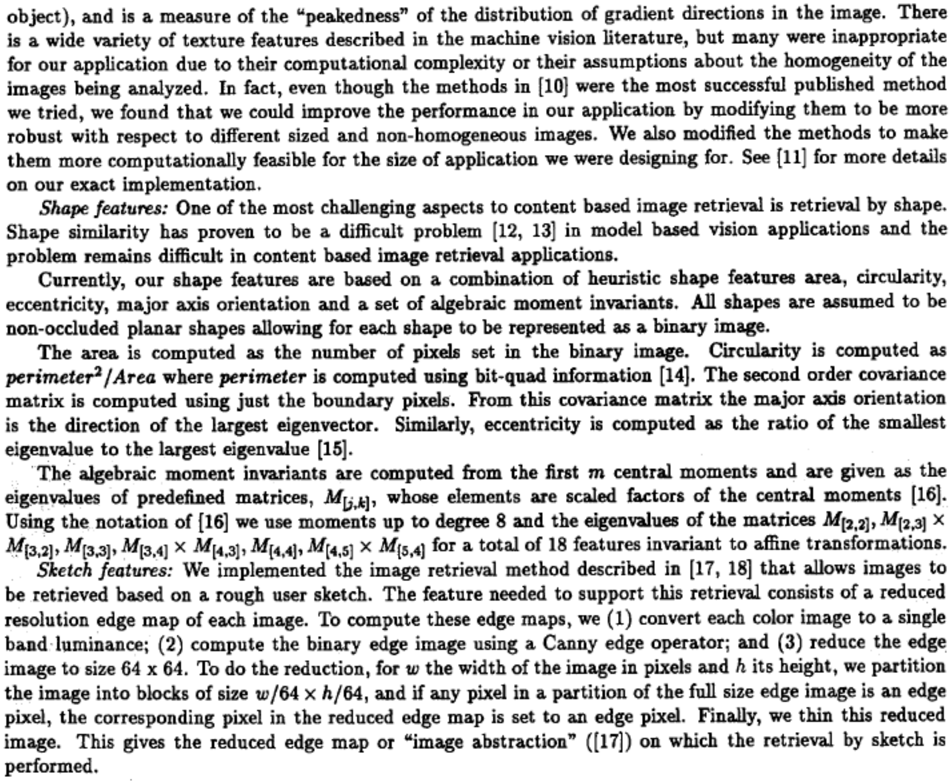 object), and is a measure of the "peakedness" of the distribution of gradient directions in the image. There is a wide variety of texture features described in the machine vision literature, but many were inappropriate for our application due to their computational complexity or their assumptions about the homogeneity of the images being analyzed. In fact, even though the methods in [10] were the most successful published method we tried, we found that we could improve the performance in our application by modifying them to be more robust with respect to different sized and non-homogeneous images. We also modified the methods to make them more computationally feasible for the size of application we were designing for. See [11] for more details on our exact implementation. *Shape features:* One of the most challenging aspects to content based image retrieval is retrieval by shape. Shape similarity has proven to be a difficult problem [12, 13] in model based vision applications and the problem remains difficult in content based image retrieval applications. Currently, our shape features are based on a combination of heuristic shape features area, circularity, eccentricity, major axis orientation and a set of algebraic moment invariants. All shapes are assumed to be non-occluded planar shapes allowing for each shape to be represented as a binary image. The area is computed as the number of pixels set in the binary image. Circularity is computed as $perimeter^2/Area$ where $perimeter$ is computed using bit-quad information [14]. The second order covariance matrix is computed using just the boundary pixels. From this covariance matrix the major axis orientation is the direction of the largest eigenvector. Similarly, eccentricity is computed as the ratio of the smallest eigenvalue to the largest eigenvalue [15]. The algebraic moment invariants are computed from the first m central moments and are given as the eigenvalues of predefined matrices, $M_{[j,k]}$, whose elements are scaled factors of the central moments [16]. Using the notation of [16] we use moments up to degree 8 and the eigenvalues of the matrices $M_{[2,2]}, M_{[2,3]} \times M_{[3,2]}, M_{[3,3]}, M_{[3,4]} \times M_{[4,3]}, M_{[4,4]}, M_{[4,5]} \times M_{[5,4]}$ for a total of 18 features invariant to affine transformations. *Sketch features:* We implemented the image retrieval method described in [17, 18] that allows images to be retrieved based on a rough user sketch. The feature needed to support this retrieval consists of a reduced resolution edge map of each image. To compute these edge maps, we (1) convert each color image to a single band luminance; (2) compute the binary edge image using a Canny edge operator; and (3) reduce the edge image to size 64 x 64. To do the reduction, for w the width of the image in pixels and h its height, we partition the image into blocks of size $w/64 \times h/64$, and if any pixel in a partition of the full size edge image is an edge pixel, the corresponding pixel in the reduced edge map is set to an edge pixel. Finally, we thin this reduced image. This gives the reduced edge map or "image abstraction" ([17]) on which the retrieval by sketch is performed. IBM0002390 at 2393-95:

Element	'278 Claim Recitation	Exemplary Citations in Reference
		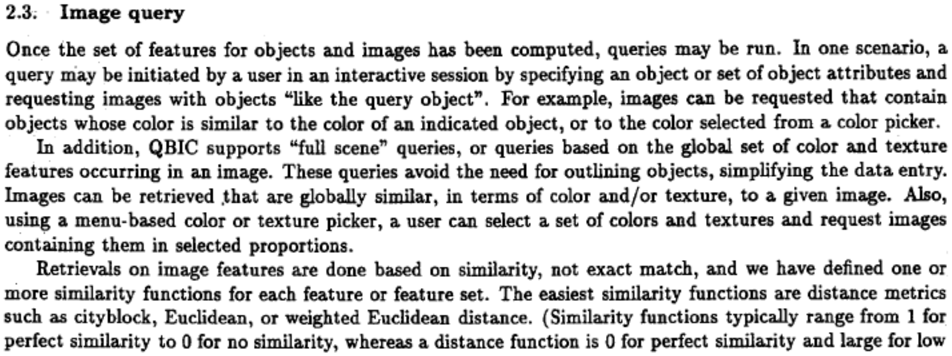 **2.3. Image query** Once the set of features for objects and images has been computed, queries may be run. In one scenario, a query may be initiated by a user in an interactive session by specifying an object or set of object attributes and requesting images with objects "like the query object". For example, images can be requested that contain objects whose color is similar to the color of an indicated object, or to the color selected from a color picker. In addition, QBIC supports "full scene" queries, or queries based on the global set of color and texture features occurring in an image. These queries avoid the need for outlining objects, simplifying the data entry. Images can be retrieved that are globally similar, in terms of color and/or texture, to a given image. Also, using a menu-based color or texture picker, a user can select a set of colors and textures and request images containing them in selected proportions. Retrievals on image features are done based on similarity, not exact match, and we have defined one or more similarity functions for each feature or feature set. The easiest similarity functions are distance metrics such as cityblock, Euclidean, or weighted Euclidean distance. (Similarity functions typically range from 1 for perfect similarity to 0 for no similarity, whereas a distance function is 0 for perfect similarity and large for low

Element	'278 Claim Recitation	Exemplary Citations in Reference
		similarity, so we must convert as necessary to obtain consistent measures.) Different similarity measures may be used, and they may vary from application to application. The matching methods and similarity functions for the different feature types are described below: *Color:* For average color, the distance between a query object and database object is weighted Euclidean distance. The weights are user adjustable so that we may, for example, downweight the L component of an (L, a, b) query. Best results, as judged by our subjective evaluation in on-going experiments, are obtained by representing the average color in the MTM color space, and by inversely weighting each component by its standard deviation over the samples in the database. Improved results are obtained when color matching is done not on average color, but on the distribution of colors that occur in an object or image. In this way, a red and blue object matches other red and blue objects better than to a purple object. To match color histograms, we use a method similar to the one described in [8]. Let X be the query histogram and Y the histogram of an item in the database, both normalized. We compute the element by element difference histogram Z. Then the similarity between X and Y is given by $\|Z\| = Z^TAZ$, where A is a symmetric color similarity matrix with $$a(i,j) \equiv 1 - d(c_i, c_j)/d_{max}$$ where c_i and c_j are the ith and jth colors in the color histograms, and $d(c_i, c_j)$ is the MTM color distance, and d_{max} is the maximum distance between any two colors ([8] used Luv distance.) This metric gives the (weighted) length of the difference vector between X and Y, weighted by A, which accounts for the perceptual distance between different pairs of colors. The result is a measure that accounts for both the difference between the amounts in any given color (e.g. a particular shade of red), as well as differences between similar colors (e.g. red and orange). Based on our experience, we also found it useful for the user to be able to request images with $x\%$ of color 1, $y\%$ of color 2, etc. where both the percentages and the colors are explicitly specified by the user. This allows searches such as for a beach scene (25% white (for the sand) and 50% blue (for the sky and water)), and for many scenes for which the user has seen before, and for which he/she can remember approximate color percentages. The color picker that allows this type of query is shown in Figure 2. 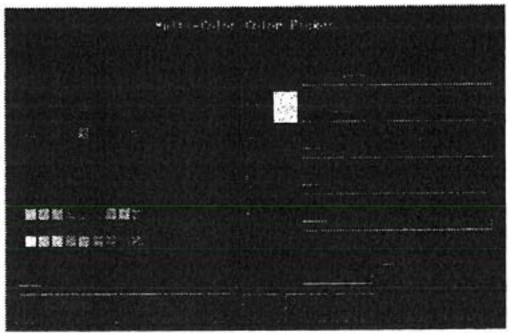 Figure 2: Multi-Color Color Picker with two selected colors. The palette on the left displays the set of selectable colors. Once a color is selected, it appears in a box on the right, and the amount of the color can be specified by the sliders. In this case, 25% light blue and 15% white have been selected. *Texture:* Texture distance is computed as weighted Euclidean distance in the three dimensional texture space. The most common weighting factors are the inverse variances for each component, computed over the

Element	'278 Claim Recitation	Exemplary Citations in Reference
		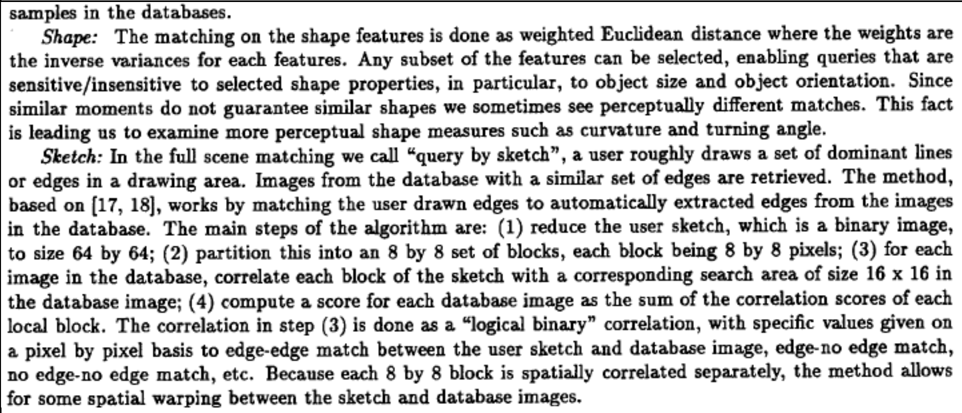 samples in the databases. *Shape:* The matching on the shape features is done as weighted Euclidean distance where the weights are the inverse variances for each features. Any subset of the features can be selected, enabling queries that are sensitive/insensitive to selected shape properties, in particular, to object size and object orientation. Since similar moments do not guarantee similar shapes we sometimes see perceptually different matches. This fact is leading us to examine more perceptual shape measures such as curvature and turning angle. *Sketch:* In the full scene matching we call "query by sketch", a user roughly draws a set of dominant lines or edges in a drawing area. Images from the database with a similar set of edges are retrieved. The method, based on [17, 18], works by matching the user drawn edges to automatically extracted edges from the images in the database. The main steps of the algorithm are: (1) reduce the user sketch, which is a binary image, to size 64 by 64; (2) partition this into an 8 by 8 set of blocks, each block being 8 by 8 pixels; (3) for each image in the database, correlate each block of the sketch with a corresponding search area of size 16 x 16 in the database image; (4) compute a score for each database image as the sum of the correlation scores of each local block. The correlation in step (3) is done as a "logical binary" correlation, with specific values given on a pixel by pixel basis to edge-edge match between the user sketch and database image, edge-no edge match, no edge-no edge match, etc. Because each 8 by 8 block is spatially correlated separately, the method allows for some spatial warping between the sketch and database images. IBM0002390 at 2395 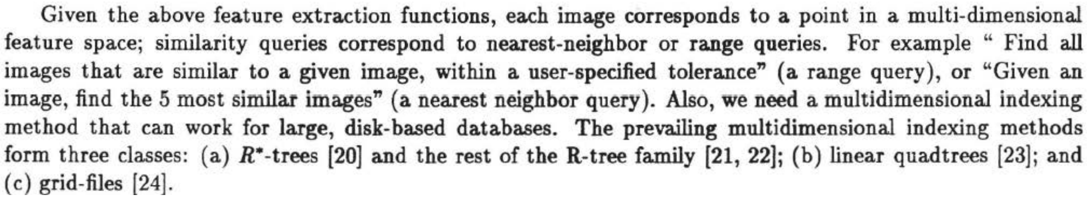 Given the above feature extraction functions, each image corresponds to a point in a multi-dimensional feature space; similarity queries correspond to nearest-neighbor or range queries. For example " Find all images that are similar to a given image, within a user-specified tolerance" (a range query), or "Given an image, find the 5 most similar images" (a nearest neighbor query). Also, we need a multidimensional indexing method that can work for large, disk-based databases. The prevailing multidimensional indexing methods form three classes: (a) R^*-trees [20] and the rest of the R-tree family [21, 22]; (b) linear quadtrees [23]; and (c) grid-files [24]. IBM0002390 at 2396:

Element	'278 Claim Recitation	Exemplary Citations in Reference
		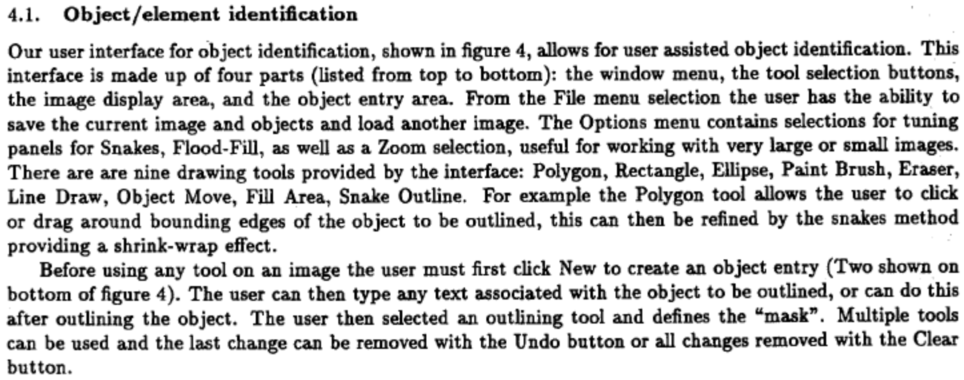 **4.1. Object/element identification** Our user interface for object identification, shown in figure 4, allows for user assisted object identification. This interface is made up of four parts (listed from top to bottom): the window menu, the tool selection buttons, the image display area, and the object entry area. From the File menu selection the user has the ability to save the current image and objects and load another image. The Options menu contains selections for tuning panels for Snakes, Flood-Fill, as well as a Zoom selection, useful for working with very large or small images. There are are nine drawing tools provided by the interface: Polygon, Rectangle, Ellipse, Paint Brush, Eraser, Line Draw, Object Move, Fill Area, Snake Outline. For example the Polygon tool allows the user to click or drag around bounding edges of the object to be outlined, this can then be refined by the snakes method providing a shrink-wrap effect. Before using any tool on an image the user must first click New to create an object entry (Two shown on bottom of figure 4). The user can then type any text associated with the object to be outlined, or can do this after outlining the object. The user then selected an outlining tool and defines the "mask". Multiple tools can be used and the last change can be removed with the Undo button or all changes removed with the Clear button. QBIC (Query by Image Content) at 1

Element	'278 Claim Recitation	Exemplary Citations in Reference
		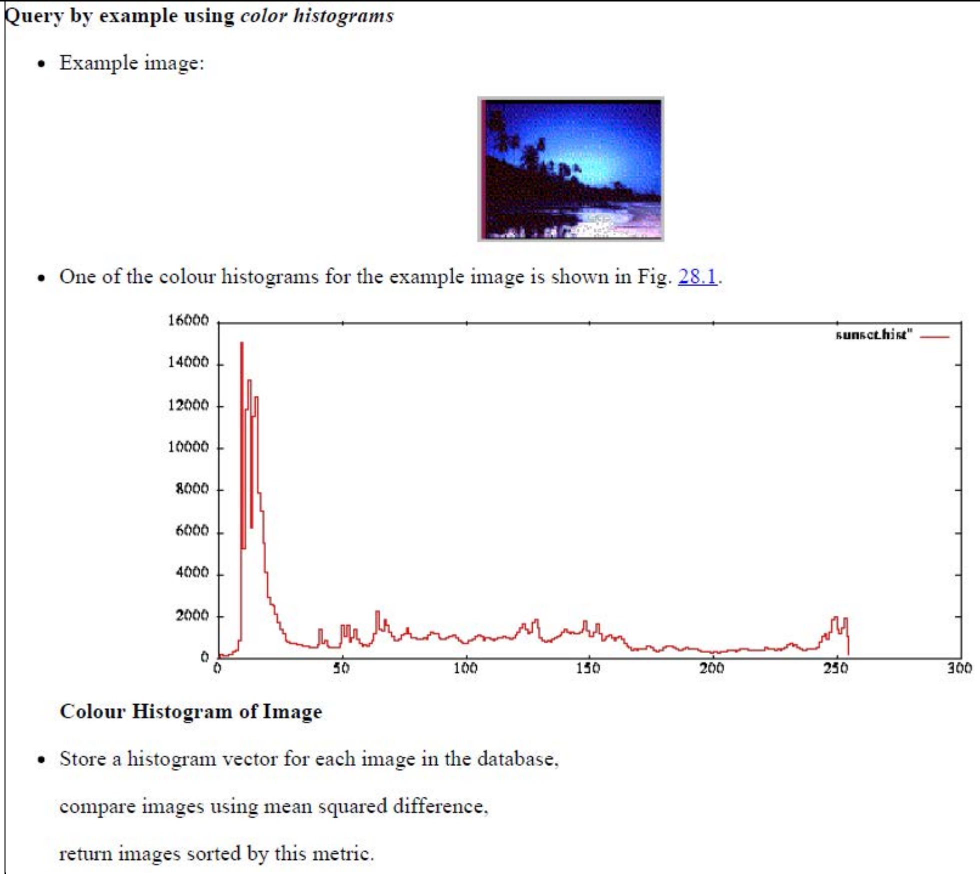 **Query by example using *color histograms*** - Example image: - One of the colour histograms for the example image is shown in Fig. 28.1. **Colour Histogram of Image** - Store a histogram vector for each image in the database, compare images using mean squared difference, return images sorted by this metric. Query by Image and Video Content: The QBIC System (IBM 0000893–902): "We have developed the QBIC (Query by Image Content) system to explore content-based retrieval methods. QBIC allows queries on large image and video databases based on • example images, • user-constructed sketches and drawings, • selected color and texture patterns, • camera and object motion, and • other graphical information." *Id.* at pp. 1–3.

Element	'278 Claim Recitation	Exemplary Citations in Reference
		"INTEGRATION WITH TEXT AND PARAMETRIC ANNOTATION. Query by visual content complements and extends existing query methods. Systems must be able to integrate queries combining date, subject matter, price, and availability with content properties such as color, texture, and shape." *Id.* at p. 8. Image, Audio, and Video Extenders Administration and Programming (IBM 000001–617): "DB2 (DB2) Universal Database (UDB) is a powerful, object-relational database manager. It stores and protects traditional numeric and character data, as well as large, complex objects (LOBs). DB2 extenders help you exploit DB2's object-relational features. The extenders define distinct data types and special functions for image, audio, video, and text objects. … Because of that, the extenders give your applications a single point of access to any or all of these types of data, along with traditional numeric and character data." *Id.* at p. 3. Managing Enterprise Information Portal (IBM 000777–880): "Many paper-intensive enterprises, such as insurance companies and financial institutions, administer large volumes of business-related content. The need for an enterprise solution for managing and accessing business information spans many industries. A content server is a software system that stores multimedia, business forms, documents, and related data, along with metadata that allows employees to process and work with the content." *Id.* at p. 3. "The image search server uses IBM's QBIC (query by image content) technology to help you search for objects by certain visual properties, such as color and texture. The image search server analyzes images and stores the image information in a database. Then users can run image queries, which use the visual properties of images, to match colors, textures, and their positions without describing them in words. You can combine content-based queries with

Element	'278 Claim Recitation	Exemplary Citations in Reference
		text and keyword searches for more powerful retrieval of image and multimedia data." *Id.* at p. 63. QBIC Demo (IBM000747) 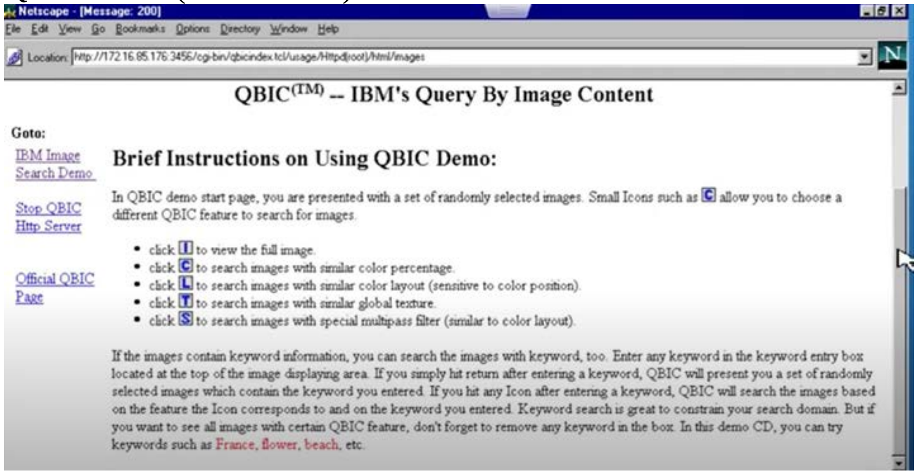 QBIC(TM) -- IBM's Query By Image Content **Brief Instructions on Using QBIC Demo:** In QBIC demo start page, you are presented with a set of randomly selected images. Small Icons such as C allow you to choose a different QBIC feature to search for images. - click I to view the full image. - click C to search images with similar color percentage. - click L to search images with similar color layout (sensitive to color position). - click T to search images with similar global texture. - click S to search images with special multipass filter (similar to color layout). If the images contain keyword information, you can search the images with keyword, too. Enter any keyword in the keyword entry box located at the top of the image displaying area. If you simply hit return after entering a keyword, QBIC will present you a set of randomly selected images which contain the keyword you entered. If you hit any Icon after entering a keyword, QBIC will search the images based on the feature the Icon corresponds to and on the keyword you entered. Keyword search is great to constrain your search domain. But if you want to see all images with certain QBIC feature, don't forget to remove any keyword in the box. In this demo CD, you can try keywords such as France, flower, beach, etc. QBIC Demo (IBM000747)

Element	’278 Claim Recitation	Exemplary Citations in Reference
		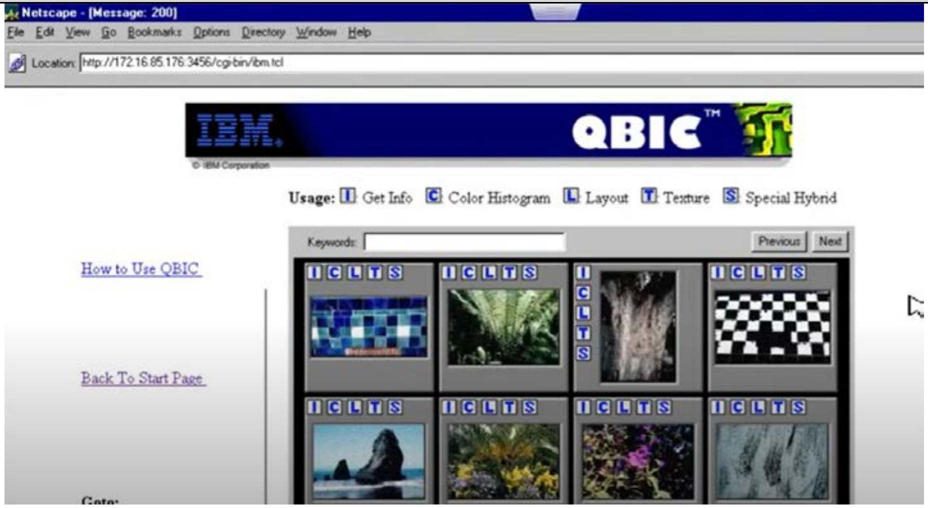 QBIC Demo (IBM000747) 

Element	'278 Claim Recitation	Exemplary Citations in Reference
		 http://172.16.85.176:3456/cgi-bin/ibm.tcl/search?fname=HybridColor&france10%2Ejpg.x=0 **Goto:** Custom Color % Query Custom Paint Query Random Images Official QBIC Page **Query was:** Example: =france10.jpg Query Type: QbHybridColorFeatureClass Flickner Depo at 78-79

Element	'278 Claim Recitation	Exemplary Citations in Reference
		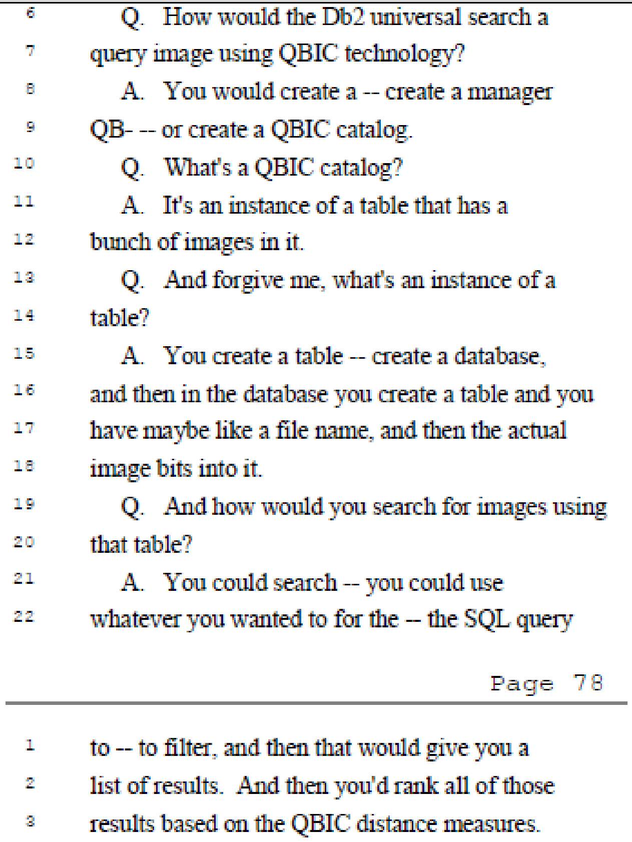 Q. How would the Db2 universal search a query image using QBIC technology? A. You would create a -- create a manager QB- -- or create a QBIC catalog. Q. What's a QBIC catalog? A. It's an instance of a table that has a bunch of images in it. Q. And forgive me, what's an instance of a table? A. You create a table -- create a database, and then in the database you create a table and you have maybe like a file name, and then the actual image bits into it. Q. And how would you search for images using that table? A. You could search -- you could use whatever you wanted to for the -- the SQL query Page 78 to -- to filter, and then that would give you a list of results. And then you'd rank all of those results based on the QBIC distance measures. Flickner Depo at 201

Element	’278 Claim Recitation	Exemplary Citations in Reference
		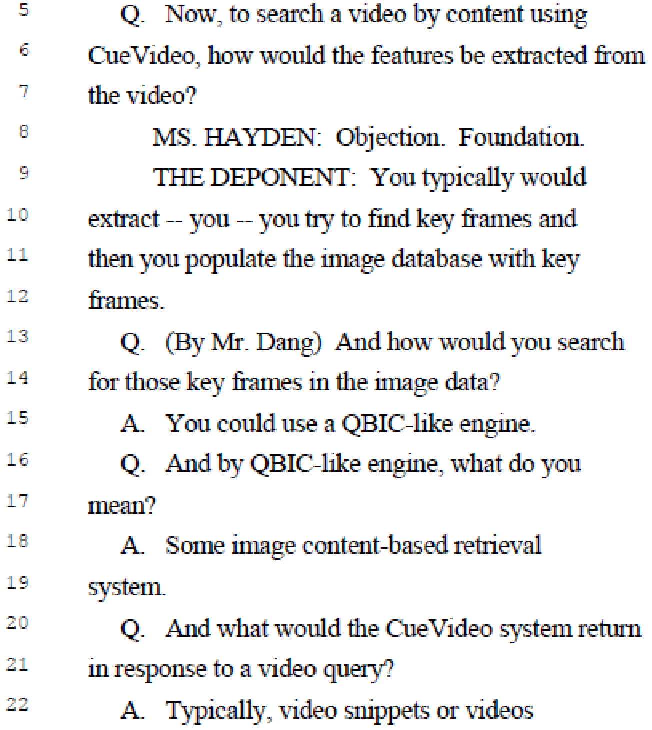 Q. Now, to search a video by content using CueVideo, how would the features be extracted from the video? MS. HAYDEN: Objection. Foundation. THE DEPONENT: You typically would extract -- you -- you try to find key frames and then you populate the image database with key frames. Q. (By Mr. Dang) And how would you search for those key frames in the image data? A. You could use a QBIC-like engine. Q. And by QBIC-like engine, what do you mean? A. Some image content-based retrieval system. Q. And what would the CueVideo system return in response to a video query? A. Typically, video snippets or videos Flickner Depo at 159 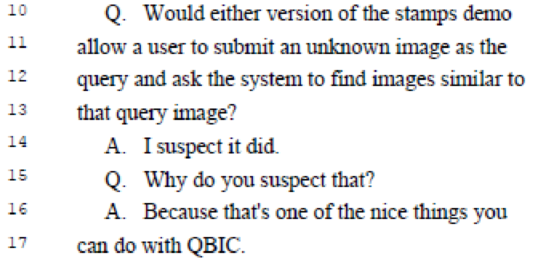 Q. Would either version of the stamps demo allow a user to submit an unknown image as the query and ask the system to find images similar to that query image? A. I suspect it did. Q. Why do you suspect that? A. Because that's one of the nice things you can do with QBIC. Flickner Depo at 209-210

Element	'278 Claim Recitation	Exemplary Citations in Reference
		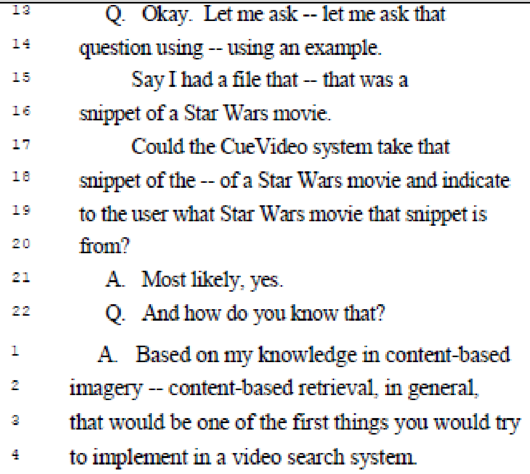 Q. Okay. Let me ask -- let me ask that question using -- using an example. Say I had a file that -- that was a snippet of a Star Wars movie. Could the CueVideo system take that snippet of the -- of a Star Wars movie and indicate to the user what Star Wars movie that snippet is from? A. Most likely, yes. Q. And how do you know that? A. Based on my knowledge in content-based imagery -- content-based retrieval, in general, that would be one of the first things you would try to implement in a video search system. Flickner Depo at 75-76 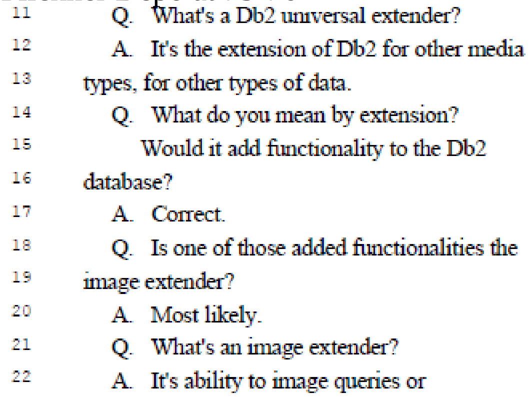 Q. What's a Db2 universal extender? A. It's the extension of Db2 for other media types, for other types of data. Q. What do you mean by extension? Would it add functionality to the Db2 database? A. Correct. Q. Is one of those added functionalities the image extender? A. Most likely. Q. What's an image extender? A. It's ability to image queries or

Element	’278 Claim Recitation	Exemplary Citations in Reference
		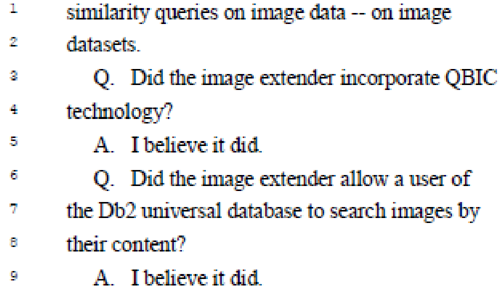 similarity queries on image data -- on image datasets. Q. Did the image extender incorporate QBIC technology? A. I believe it did. Q. Did the image extender allow a user of the Db2 universal database to search images by their content? A. I believe it did. Flickner Depo at 162-163 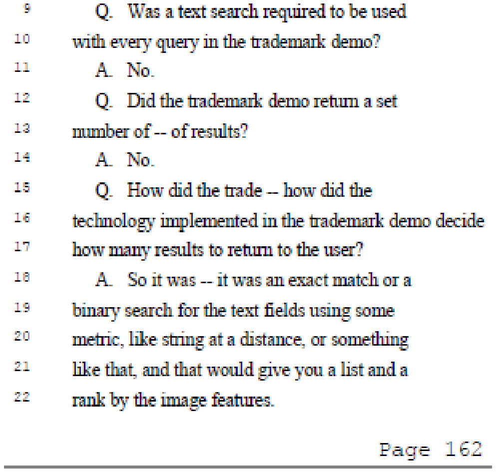 Q. Was a text search required to be used with every query in the trademark demo? A. No. Q. Did the trademark demo return a set number of -- of results? A. No. Q. How did the trade -- how did the technology implemented in the trademark demo decide how many results to return to the user? A. So it was -- it was an exact match or a binary search for the text fields using some metric, like string at a distance, or something like that, and that would give you a list and a rank by the image features. Page 162 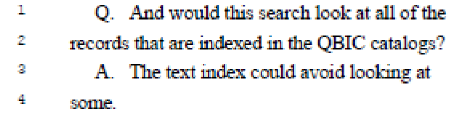 Q. And would this search look at all of the records that are indexed in the QBIC catalogs? A. The text index could avoid looking at some. Flickner 8/29/23 Depo at 33–37

Element	'278 Claim Recitation	Exemplary Citations in Reference
		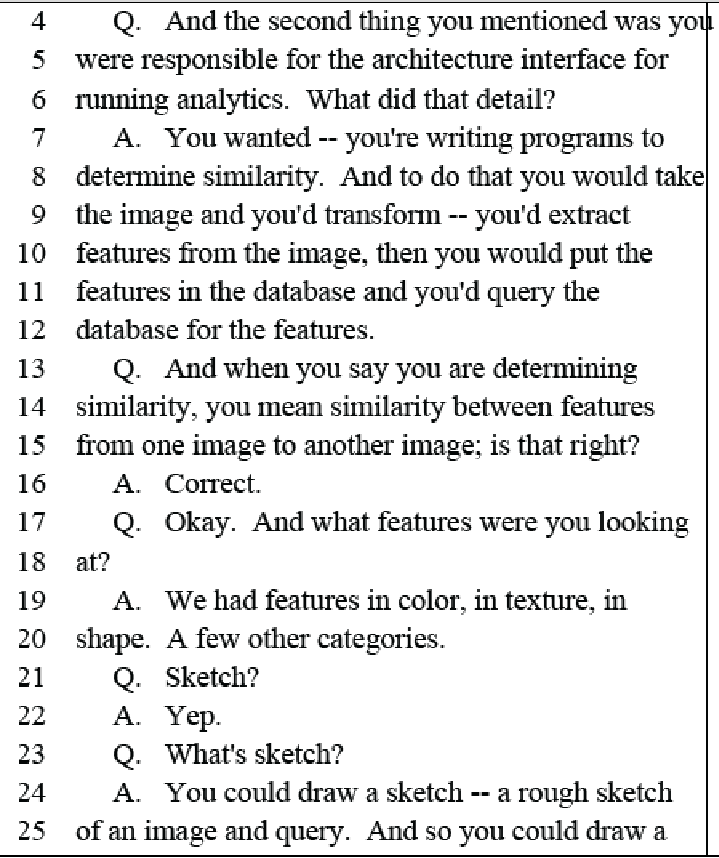 Q. And the second thing you mentioned was you were responsible for the architecture interface for running analytics. What did that detail? A. You wanted -- you're writing programs to determine similarity. And to do that you would take the image and you'd transform -- you'd extract features from the image, then you would put the features in the database and you'd query the database for the features. Q. And when you say you are determining similarity, you mean similarity between features from one image to another image; is that right? A. Correct. Q. Okay. And what features were you looking at? A. We had features in color, in texture, in shape. A few other categories. Q. Sketch? A. Yep. Q. What's sketch? A. You could draw a sketch -- a rough sketch of an image and query. And so you could draw a

Page 34

circle for a sun, a line and maybe some trees and then you could send that sketch in and you'd get images of a sun in the trees.
Q. And if I refer to images, can I -- let me -- let me back up.
Earlier you talked about both still images and video were part of your responsibilities for these architecture interfaces; correct?
A. Correct.
Q. Okay. If I refer to "images," do you understand that is including both a still image or, for example, a frame of a video image?
A. Sure.
Q. Okay. So when I say "image," I'm including both in my questions from now on out.
A. Okay.
Q. Okay?
A. Okay.
Q. In terms of running analytics, that would include both extracting these characteristics from a query image and from an image that you're storing in a database; correct?
A. Can you rephrase.
Q. Yes. So the comparison that you're doing between two images; right, is it correct that that

Page 35

comparison is for purposes of a query? You're querying image 1 to determine if there are images in the database that have similar features; is that right?
A. Correct.
Q. Okay. And the analytics that you're talking about, where you extract these characters from the image, that is -- those run on both the query image and the images that are stored in the database; is that right?
A. Yeah. We call them features.
Q. Features. Okay.
So when you say "features," you're talking about the features that we talked about before: Color, texture, shape, sketch?
A. For example.
Q. Any other ones you can think of?
A. Nothing comes to mind.
Q. But when you refer to features they at least include those four; is that correct?
A. Yes. Well, at least includes, yeah.
Q. Is that correct? I'm sorry, I didn't -- you trailed off, so I -- that's why I reasked it.
A. Yeah.
Q. Let me reask it and then you can give me a

Page 36

clean answer.
When you refer to features, does it include at least color, texture, shape and sketch?
A. Or a subset thereof.
(Reporter seeks clarification.)
A. Or a subset thereof.
So you could do just color queries, you could do just shape queries, you could do just -- you didn't need to have all the feature sets in all -- in every query.
Q. Understood. But would a user have the option to query at least those four features?
A. Yes.
Q. Okay.
A. For their QBIC system, that's true.
Q. And then I believe the last area you indicated you were responsible for was architecture interface for saving images; is that right?
A. Correct.
Q. And what did that entail?
A. It's both saving images and saving features. That was to define the API, so you would give it a feature, and it would write the feature in a database or into a file.
Q. Would that also include saving any

Page 37

information associated with the images, for example, like any text description about the image?
A. Yeah. Text is just another feature here.
Q. Text is another feature?
A. Yeah.
Q. I understand when the images were saved, thumbnails were saved of those images; is that correct?
A. Typically, yeah.
Q. And was that a part of your responsibilities, to architect the interface for saving images?
A. Saving . . . hmm. Can you rephrase.
Q. What's -- what was confusing about my question?
A. I just need you to repeat it.
Q. Oh. One of the areas you mentioned that you were responsible for was to architect the interface for saving images.
My question is, as part of that responsibility, did that include architecting the features related to saving thumbnails associated with those images?
A. The thumbnail is an extracted feature.
Q. Okay. So is that a "yes"?

Element	'278 Claim Recitation	Exemplary Citations in Reference
		Id. at 52–53 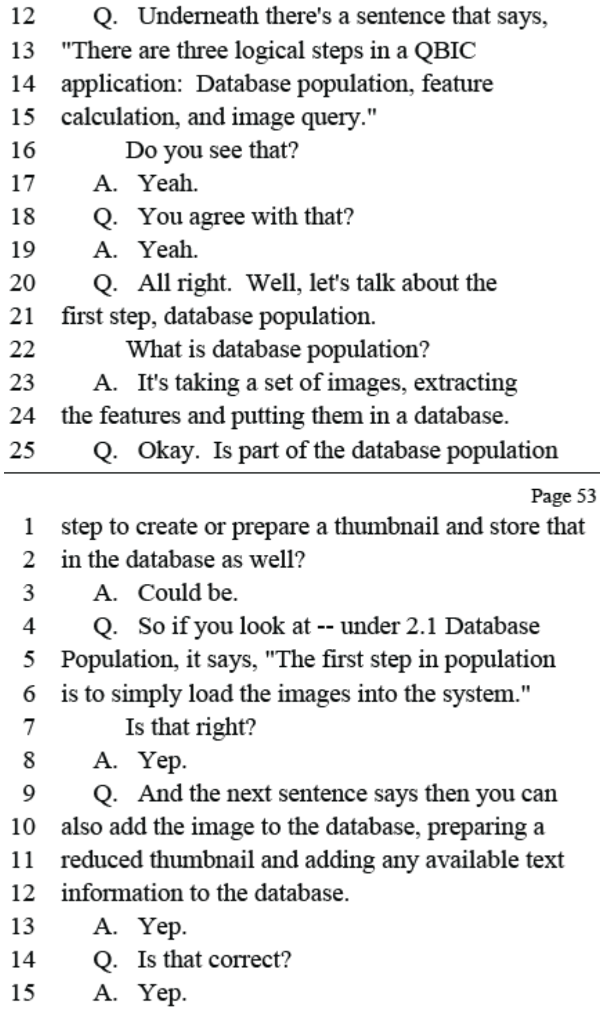 Q. Underneath there's a sentence that says, "There are three logical steps in a QBIC application: Database population, feature calculation, and image query." Do you see that? A. Yeah. Q. You agree with that? A. Yeah. Q. All right. Well, let's talk about the first step, database population. What is database population? A. It's taking a set of images, extracting the features and putting them in a database. Q. Okay. Is part of the database population step to create or prepare a thumbnail and store that in the database as well? Page 53 A. Could be. Q. So if you look at -- under 2.1 Database Population, it says, "The first step in population is to simply load the images into the system." Is that right? A. Yep. Q. And the next sentence says then you can also add the image to the database, preparing a reduced thumbnail and adding any available text information to the database. A. Yep. Q. Is that correct? A. Yep. *Id.* at 77–78

Element	'278 Claim Recitation	Exemplary Citations in Reference
		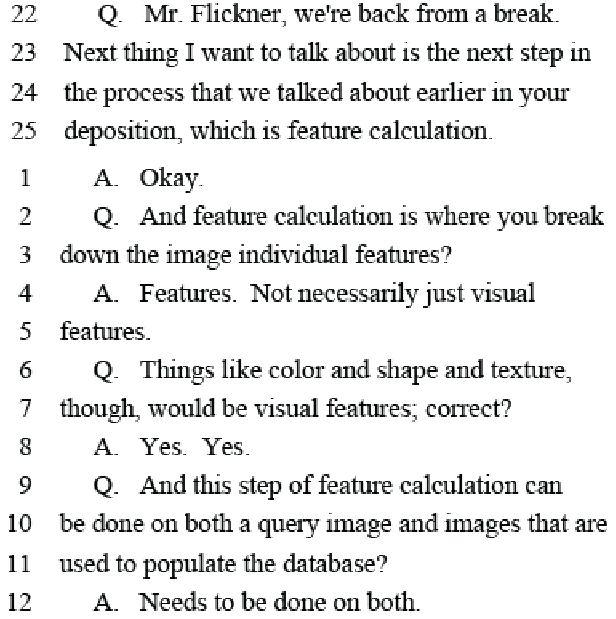 Q. Mr. Flickner, we're back from a break. Next thing I want to talk about is the next step in the process that we talked about earlier in your deposition, which is feature calculation. A. Okay. Q. And feature calculation is where you break down the image individual features? A. Features. Not necessarily just visual features. Q. Things like color and shape and texture, though, would be visual features; correct? A. Yes. Yes. Q. And this step of feature calculation can be done on both a query image and images that are used to populate the database? A. Needs to be done on both. *Id.* at 85–99 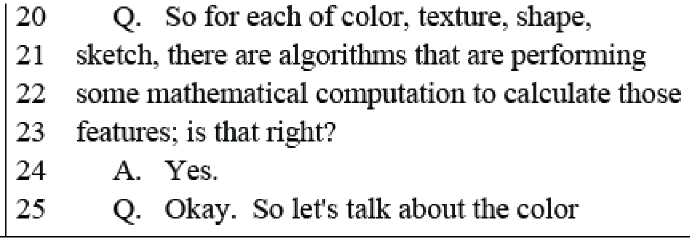 Q. So for each of color, texture, shape, sketch, there are algorithms that are performing some mathematical computation to calculate those features; is that right? A. Yes. Q. Okay. So let's talk about the color

		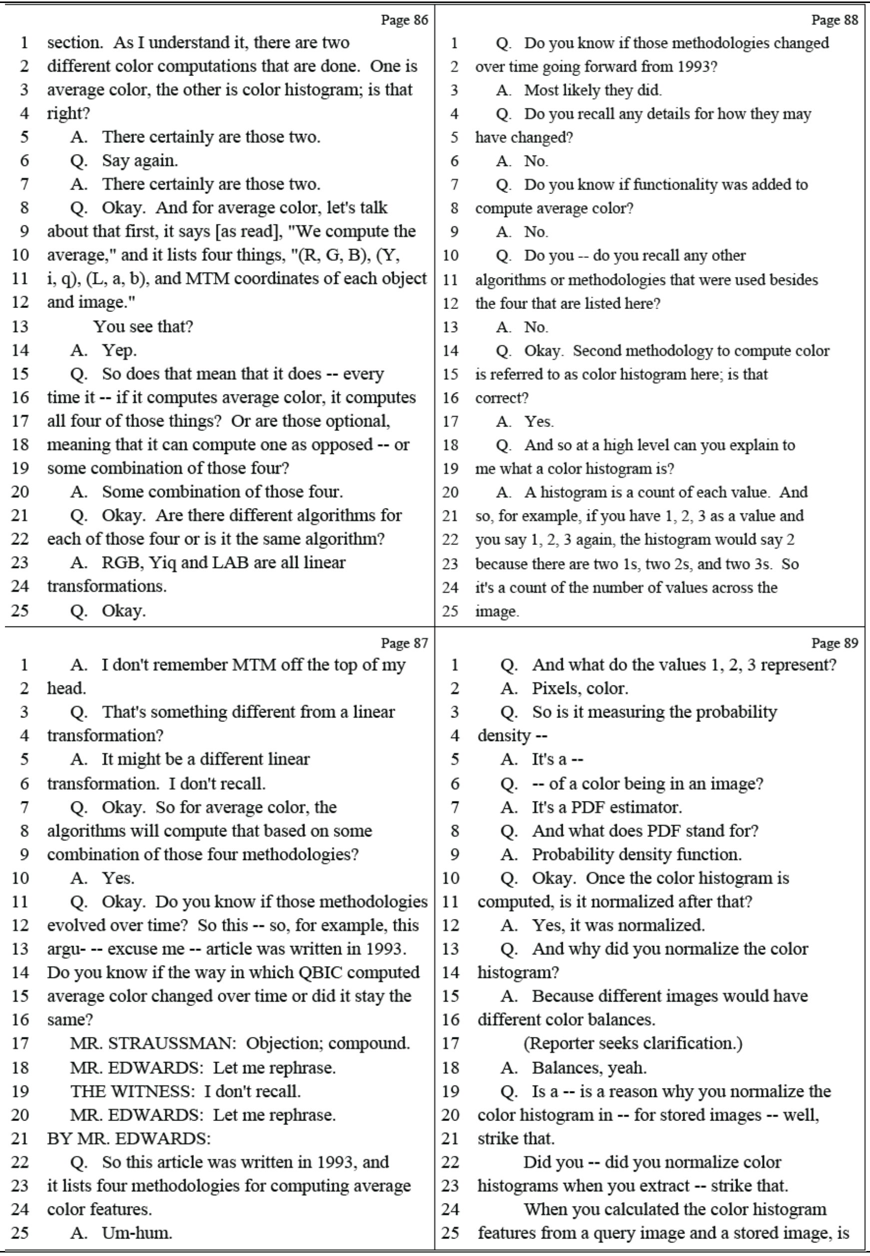

Page 86

section. As I understand it, there are two different color computations that are done. One is average color, the other is color histogram; is that right?
A. There certainly are those two.
Q. Say again.
A. There certainly are those two.
Q. Okay. And for average color, let's talk about that first, it says [as read], "We compute the average," and it lists four things, "(R, G, B), (Y, i, q), (L, a, b), and MTM coordinates of each object and image."
You see that?
A. Yep.
Q. So does that mean that it does -- every time it -- if it computes average color, it computes all four of those things? Or are those optional, meaning that it can compute one as opposed -- or some combination of those four?
A. Some combination of those four.
Q. Okay. Are there different algorithms for each of those four or is it the same algorithm?
A. RGB, Yiq and LAB are all linear transformations.
Q. Okay.

Page 87

A. I don't remember MTM off the top of my head.
Q. That's something different from a linear transformation?
A. It might be a different linear transformation. I don't recall.
Q. Okay. So for average color, the algorithms will compute that based on some combination of those four methodologies?
A. Yes.
Q. Okay. Do you know if those methodologies evolved over time? So this -- so, for example, this argu- -- excuse me -- article was written in 1993. Do you know if the way in which QBIC computed average color changed over time or did it stay the same?
MR. STRAUSSMAN: Objection; compound.
MR. EDWARDS: Let me rephrase.
THE WITNESS: I don't recall.
MR. EDWARDS: Let me rephrase.
BY MR. EDWARDS:
Q. So this article was written in 1993, and it lists four methodologies for computing average color features.
A. Um-hum.

Page 88

Q. Do you know if those methodologies changed over time going forward from 1993?
A. Most likely they did.
Q. Do you recall any details for how they may have changed?
A. No.
Q. Do you know if functionality was added to compute average color?
A. No.
Q. Do you -- do you recall any other algorithms or methodologies that were used besides the four that are listed here?
A. No.
Q. Okay. Second methodology to compute color is referred to as color histogram here; is that correct?
A. Yes.
Q. And so at a high level can you explain to me what a color histogram is?
A. A histogram is a count of each value. And so, for example, if you have 1, 2, 3 as a value and you say 1, 2, 3 again, the histogram would say 2 because there are two 1s, two 2s, and two 3s. So it's a count of the number of values across the image.

Page 89

Q. And what do the values 1, 2, 3 represent?
A. Pixels, color.
Q. So is it measuring the probability density --
A. It's a --
Q. -- of a color being in an image?
A. It's a PDF estimator.
Q. And what does PDF stand for?
A. Probability density function.
Q. Okay. Once the color histogram is computed, is it normalized after that?
A. Yes, it was normalized.
Q. And why did you normalize the color histogram?
A. Because different images would have different color balances.
(Reporter seeks clarification.)
A. Balances, yeah.
Q. Is a -- is a reason why you normalize the color histogram in -- for stored images -- well, strike that.
Did you -- did you normalize color histograms when you extract -- strike that.
When you calculated the color histogram features from a query image and a stored image, is

		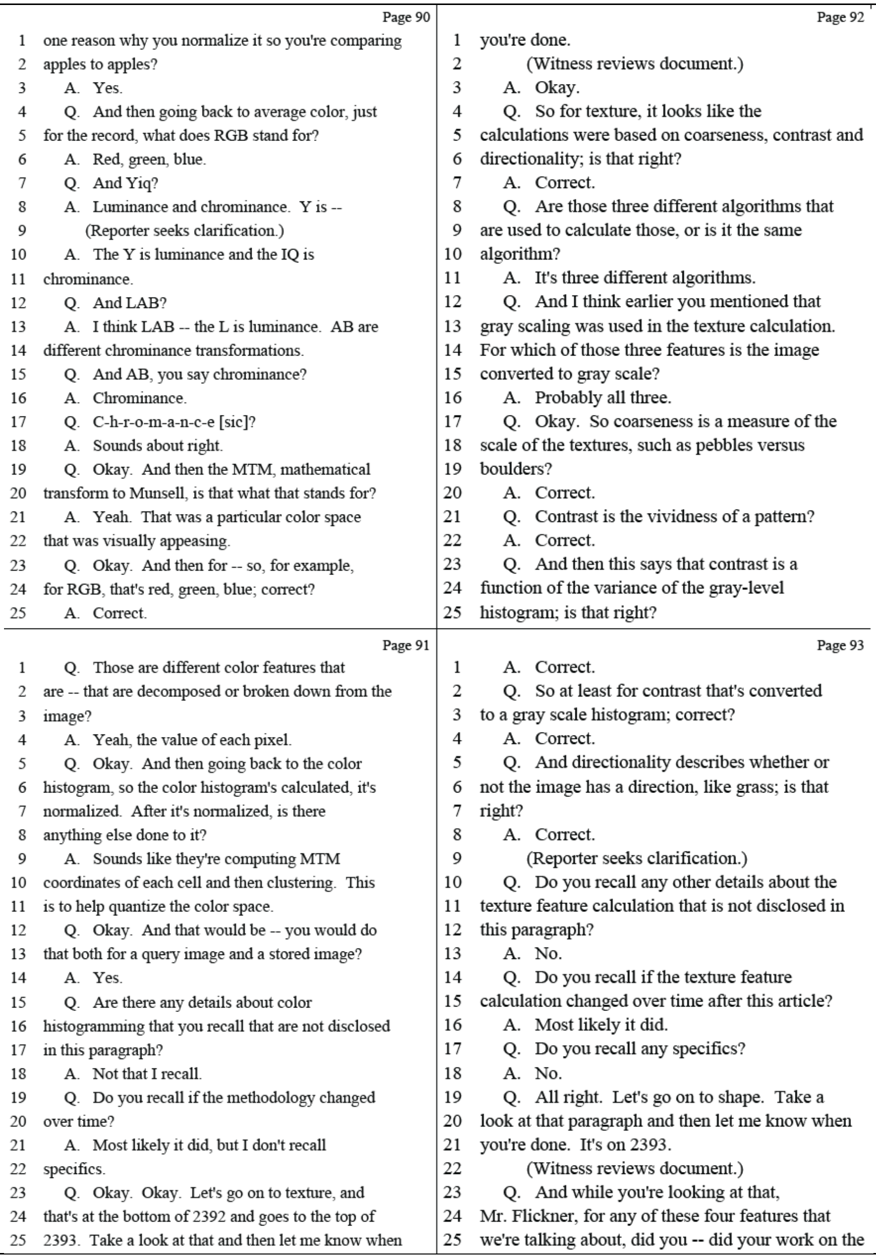

Page 90

one reason why you normalize it so you're comparing apples to apples?
A. Yes.
Q. And then going back to average color, just for the record, what does RGB stand for?
A. Red, green, blue.
Q. And Yiq?
A. Luminance and chrominance. Y is --
(Reporter seeks clarification.)
A. The Y is luminance and the IQ is chrominance.
Q. And LAB?
A. I think LAB -- the L is luminance. AB are different chrominance transformations.
Q. And AB, you say chrominance?
A. Chrominance.
Q. C-h-r-o-m-a-n-c-e [sic]?
A. Sounds about right.
Q. Okay. And then the MTM, mathematical transform to Munsell, is that what that stands for?
A. Yeah. That was a particular color space that was visually appeasing.
Q. Okay. And then for -- so, for example, for RGB, that's red, green, blue; correct?
A. Correct.

Page 91

Q. Those are different color features that are -- that are decomposed or broken down from the image?
A. Yeah, the value of each pixel.
Q. Okay. And then going back to the color histogram, so the color histogram's calculated, it's normalized. After it's normalized, is there anything else done to it?
A. Sounds like they're computing MTM coordinates of each cell and then clustering. This is to help quantize the color space.
Q. Okay. And that would be -- you would do that both for a query image and a stored image?
A. Yes.
Q. Are there any details about color histogramming that you recall that are not disclosed in this paragraph?
A. Not that I recall.
Q. Do you recall if the methodology changed over time?
A. Most likely it did, but I don't recall specifics.
Q. Okay. Okay. Let's go on to texture, and that's at the bottom of 2392 and goes to the top of 2393. Take a look at that and then let me know when

Page 92

you're done.
(Witness reviews document.)
A. Okay.
Q. So for texture, it looks like the calculations were based on coarseness, contrast and directionality; is that right?
A. Correct.
Q. Are those three different algorithms that are used to calculate those, or is it the same algorithm?
A. It's three different algorithms.
Q. And I think earlier you mentioned that gray scaling was used in the texture calculation. For which of those three features is the image converted to gray scale?
A. Probably all three.
Q. Okay. So coarseness is a measure of the scale of the textures, such as pebbles versus boulders?
A. Correct.
Q. Contrast is the vividness of a pattern?
A. Correct.
Q. And then this says that contrast is a function of the variance of the gray-level histogram; is that right?

Page 93

A. Correct.
Q. So at least for contrast that's converted to a gray scale histogram; correct?
A. Correct.
Q. And directionality describes whether or not the image has a direction, like grass; is that right?
A. Correct.
(Reporter seeks clarification.)
Q. Do you recall any other details about the texture feature calculation that is not disclosed in this paragraph?
A. No.
Q. Do you recall if the texture feature calculation changed over time after this article?
A. Most likely it did.
Q. Do you recall any specifics?
A. No.
Q. All right. Let's go on to shape. Take a look at that paragraph and then let me know when you're done. It's on 2393.
(Witness reviews document.)
Q. And while you're looking at that, Mr. Flickner, for any of these four features that we're talking about, did you -- did your work on the

Page 94

QBIC project -- were you more involved in one or more of these features than the other?
A. Yes.
Q. Which one?
A. Shape.
Q. Okay.
(Witness reviews document.)
A. So one thing I can state is originally I couldn't remember the name of the edge operator -- and the name is Canny, the Canny edge operator -- when we were talking about finding edges.
Q. Okay. So this -- the paragraph on IBM 2393 refreshed your memory that shape calculation used Canny edge detection?
A. No, that's sketch.
Q. Oh, I'm sorry.
A. I'm just stating that originally I was asked a question about standard edge operators. I couldn't remember the name. The name was Canny.
Q. Canny. And Canny is used in the sketch calculation?
A. Yeah, it's used in the sketch calculation.
Q. Understood. Okay. Let's go back to shape.
And for shape, is it a first step -- well,

Page 95

back up, strike that.
To calculate the shape feature in a query image or a stored database image, is the first step to convert that image into a binary image?
A. I don't think so.
Q. So if you look at the second paragraph -- and the reason why I ask is -- it's the second sentence in the second paragraph under "Shape features." It says [as read], "All shapes are assumed to be non-occluded planar shapes allowing for each shape to be represented as a binary image."
See that?
A. Yeah.
Q. And then the next sentence says [as read], "The area is computed as the number of pixels set in the binary image."
A. Right.
Q. So that's why I ask is it the first step in calculating shape to convert the image to binary?
A. Yes.
Q. And then once you convert it to binary image, then what do you do next in order to calculate the shape?
A. You're computing something called perround, which is perimeter squared over the area,

Page 96

of the object, of the blob.
Q. Can you say the name of the word again?
A. I think it's called perround.
Q. Per round?
A. Yeah.
Q. Those -- are those two words, p-e-r, r-o-u-n-d?
A. P-e-r -- one word.
Q. Okay. Perround, one word.
A. It's basically the perimeter squared over area.
(Reporter seeks clarification.)
A. Perimeter squared over area.
(Reporter seeks clarification.)
A. Length of the -- over area.
Q. Okay. So let's step back a little bit. So it says -- the second paragraph says [as read], "Our shape features are based on a combination of heuristic" -- h-e-u-r-i-s-t-i-c -- "shape features area, circularity, eccentricity, major access orientation and a set of algebraic moment invariants," i-n-v-a-r-i-a-n-t-s.
Did I read that correctly?
A. Um-hum.
Q. Okay. For circularity, that is computed

Page 97

as perimeter squared over area; correct?
A. Right.
Q. All right. Area is the number of pixels in the binary image?
A. In the binary blob.
Q. In the binary what?
A. Blob.
Q. Blob. And when you say "blob," what does that refer to?
A. The outline of the object that you're doing a shape query on.
Q. And so are each of these, area, circularity, eccentricity, major axis orientation and algebraic moment invariants, are those each different algorithms or is it all the same algorithm?
A. It's different algorithms.
Q. Okay. And since you worked on the shape feature the most, do you recall any other details that are not disclosed in this paragraph?
A. Well, this evolved and we ended up doing an implementation of something that's not in here.
Q. Okay. And what was that implementation?
A. Something called turning angle per round, turning angle, which is -- it looks at the tangent

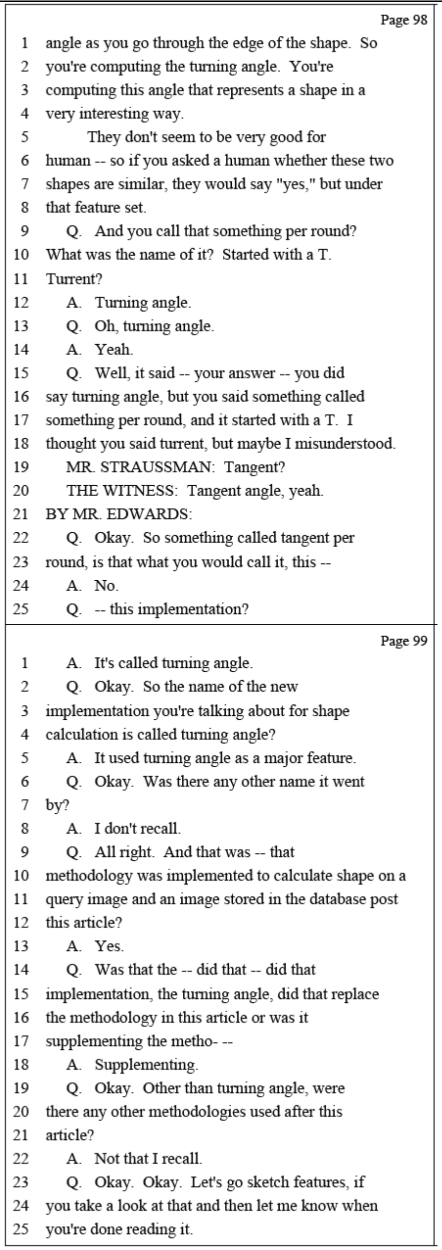

Page 98

angle as you go through the edge of the shape. So you're computing the turning angle. You're computing this angle that represents a shape in a very interesting way.

They don't seem to be very good for human -- so if you asked a human whether these two shapes are similar, they would say "yes," but under that feature set.

Q. And you call that something per round? What was the name of it? Started with a T. Turrent?

A. Turning angle.

Q. Oh, turning angle.

A. Yeah.

Q. Well, it said -- your answer -- you did say turning angle, but you said something called something per round, and it started with a T. I thought you said turrent, but maybe I misunderstood.

MR. STRAUSSMAN: Tangent?

THE WITNESS: Tangent angle, yeah.

BY MR. EDWARDS:

Q. Okay. So something called tangent per round, is that what you would call it, this --

A. No.

Q. -- this implementation?

Page 99

A. It's called turning angle.

Q. Okay. So the name of the new implementation you're talking about for shape calculation is called turning angle?

A. It used turning angle as a major feature.

Q. Okay. Was there any other name it went by?

A. I don't recall.

Q. All right. And that was -- that methodology was implemented to calculate shape on a query image and an image stored in the database post this article?

A. Yes.

Q. Was that the -- did that -- did that implementation, the turning angle, did that replace the methodology in this article or was it supplementing the metho- --

A. Supplementing.

Q. Okay. Other than turning angle, were there any other methodologies used after this article?

A. Not that I recall.

Q. Okay. Okay. Let's go sketch features, if you take a look at that and then let me know when you're done reading it.

Element	'278 Claim Recitation	Exemplary Citations in Reference
		Id. at 119–128 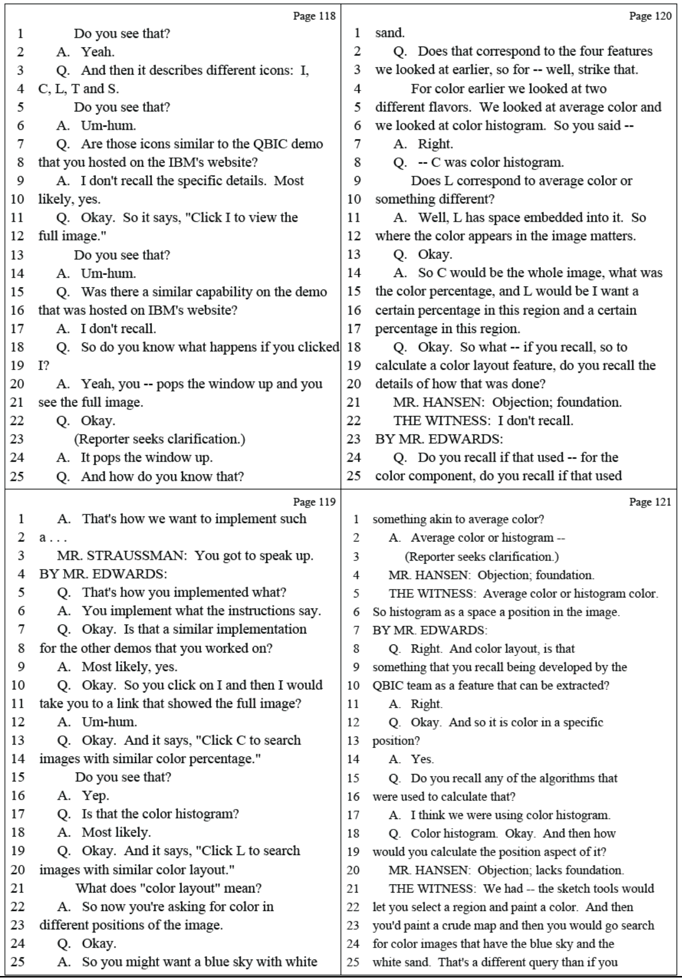 Page 118 Do you see that? A. Yeah. Q. And then it describes different icons: I, C, L, T and S. Do you see that? A. Um-hum. Q. Are those icons similar to the QBIC demo that you hosted on the IBM's website? A. I don't recall the specific details. Most likely, yes. Q. Okay. So it says, "Click I to view the full image." Do you see that? A. Um-hum. Q. Was there a similar capability on the demo that was hosted on IBM's website? A. I don't recall. Q. So do you know what happens if you clicked I? A. Yeah, you -- pops the window up and you see the full image. Q. Okay. (Reporter seeks clarification.) A. It pops the window up. Q. And how do you know that? Page 119 A. That's how we want to implement such a . . . MR. STRAUSSMAN: You got to speak up. BY MR. EDWARDS: Q. That's how you implemented what? A. You implement what the instructions say. Q. Okay. Is that a similar implementation for the other demos that you worked on? A. Most likely, yes. Q. Okay. So you click on I and then I would take you to a link that showed the full image? A. Um-hum. Q. Okay. And it says, "Click C to search images with similar color percentage." Do you see that? A. Yep. Q. Is that the color histogram? A. Most likely. Q. Okay. And it says, "Click L to search images with similar color layout." What does "color layout" mean? A. So now you're asking for color in different positions of the image. Q. Okay. A. So you might want a blue sky with white Page 120 sand. Q. Does that correspond to the four features we looked at earlier, so for -- well, strike that. For color earlier we looked at two different flavors. We looked at average color and we looked at color histogram. So you said -- A. Right. Q. -- C was color histogram. Does L correspond to average color or something different? A. Well, L has space embedded into it. So where the color appears in the image matters. Q. Okay. A. So C would be the whole image, what was the color percentage, and L would be I want a certain percentage in this region and a certain percentage in this region. Q. Okay. So what -- if you recall, so to calculate a color layout feature, do you recall the details of how that was done? MR. HANSEN: Objection; foundation. THE WITNESS: I don't recall. BY MR. EDWARDS: Q. Do you recall if that used -- for the color component, do you recall if that used Page 121 something akin to average color? A. Average color or histogram -- (Reporter seeks clarification.) MR. HANSEN: Objection; foundation. THE WITNESS: Average color or histogram color. So histogram as a space a position in the image. BY MR. EDWARDS: Q. Right. And color layout, is that something that you recall being developed by the QBIC team as a feature that can be extracted? A. Right. Q. Okay. And so it is color in a specific position? A. Yes. Q. Do you recall any of the algorithms that were used to calculate that? A. I think we were using color histogram. Q. Color histogram. Okay. And then how would you calculate the position aspect of it? MR. HANSEN: Objection; lacks foundation. THE WITNESS: We had -- the sketch tools would let you select a region and paint a color. And then you'd paint a crude map and then you would go search for color images that have the blue sky and the white sand. That's a different query than if you

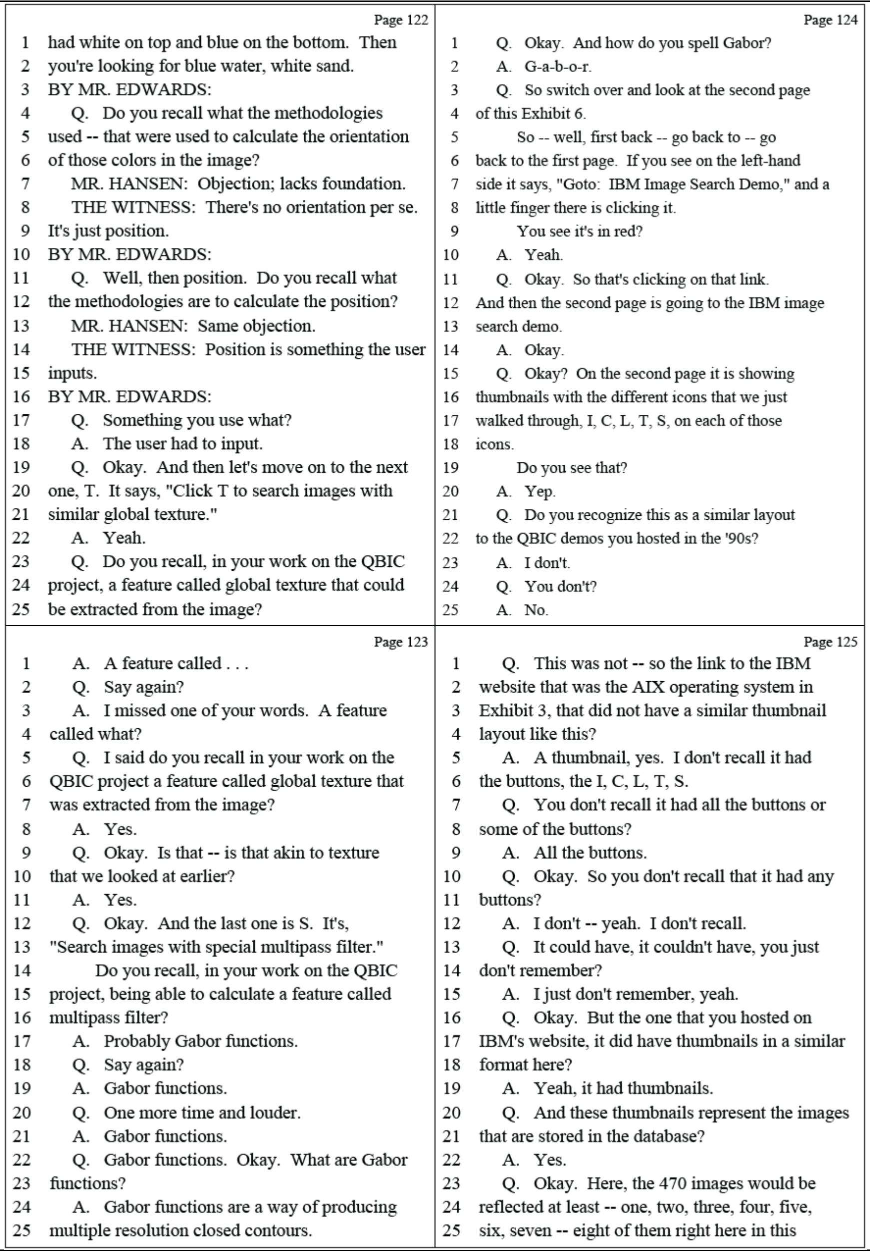

Page 122

had white on top and blue on the bottom. Then you're looking for blue water, white sand.
BY MR. EDWARDS:
Q. Do you recall what the methodologies used -- that were used to calculate the orientation of those colors in the image?
MR. HANSEN: Objection; lacks foundation.
THE WITNESS: There's no orientation per se. It's just position.
BY MR. EDWARDS:
Q. Well, then position. Do you recall what the methodologies are to calculate the position?
MR. HANSEN: Same objection.
THE WITNESS: Position is something the user inputs.
BY MR. EDWARDS:
Q. Something you use what?
A. The user had to input.
Q. Okay. And then let's move on to the next one, T. It says, "Click T to search images with similar global texture."
A. Yeah.
Q. Do you recall, in your work on the QBIC project, a feature called global texture that could be extracted from the image?

Page 123

A. A feature called . . .
Q. Say again?
A. I missed one of your words. A feature called what?
Q. I said do you recall in your work on the QBIC project a feature called global texture that was extracted from the image?
A. Yes.
Q. Okay. Is that -- is that akin to texture that we looked at earlier?
A. Yes.
Q. Okay. And the last one is S. It's, "Search images with special multipass filter."
Do you recall, in your work on the QBIC project, being able to calculate a feature called multipass filter?
A. Probably Gabor functions.
Q. Say again?
A. Gabor functions.
Q. One more time and louder.
A. Gabor functions.
Q. Gabor functions. Okay. What are Gabor functions?
A. Gabor functions are a way of producing multiple resolution closed contours.

Page 124

Q. Okay. And how do you spell Gabor?
A. G-a-b-o-r.
Q. So switch over and look at the second page of this Exhibit 6.
So -- well, first back -- go back to -- go back to the first page. If you see on the left-hand side it says, "Goto: IBM Image Search Demo," and a little finger there is clicking it.
You see it's in red?
A. Yeah.
Q. Okay. So that's clicking on that link. And then the second page is going to the IBM image search demo.
A. Okay.
Q. Okay? On the second page it is showing thumbnails with the different icons that we just walked through, I, C, L, T, S, on each of those icons.
Do you see that?
A. Yep.
Q. Do you recognize this as a similar layout to the QBIC demos you hosted in the '90s?
A. I don't.
Q. You don't?
A. No.

Page 125

Q. This was not -- so the link to the IBM website that was the AIX operating system in Exhibit 3, that did not have a similar thumbnail layout like this?
A. A thumbnail, yes. I don't recall it had the buttons, the I, C, L, T, S.
Q. You don't recall it had all the buttons or some of the buttons?
A. All the buttons.
Q. Okay. So you don't recall that it had any buttons?
A. I don't -- yeah. I don't recall.
Q. It could have, it couldn't have, you just don't remember?
A. I just don't remember, yeah.
Q. Okay. But the one that you hosted on IBM's website, it did have thumbnails in a similar format here?
A. Yeah, it had thumbnails.
Q. And these thumbnails represent the images that are stored in the database?
A. Yes.
Q. Okay. Here, the 470 images would be reflected at least -- one, two, three, four, five, six, seven -- eight of them right here in this

Element	'278 Claim Recitation	Exemplary Citations in Reference
		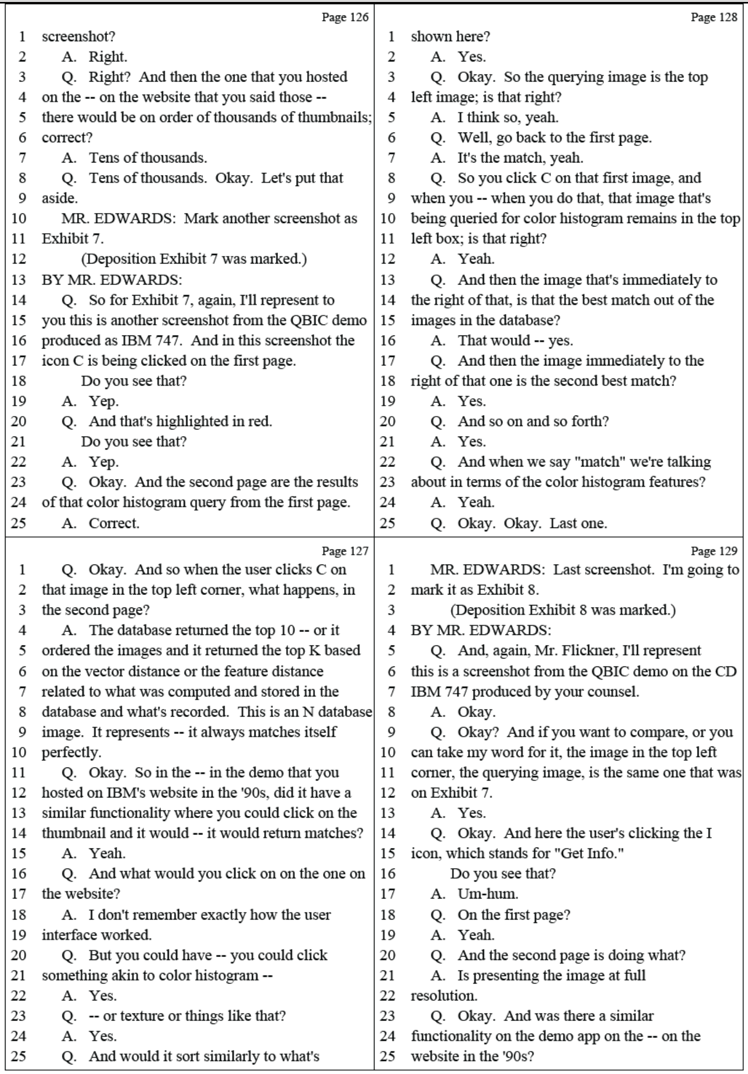 Page 126 screenshot? A. Right. Q. Right? And then the one that you hosted on the -- on the website that you said those -- there would be on order of thousands of thumbnails; correct? A. Tens of thousands. Q. Tens of thousands. Okay. Let's put that aside. MR. EDWARDS: Mark another screenshot as Exhibit 7. (Deposition Exhibit 7 was marked.) BY MR. EDWARDS: Q. So for Exhibit 7, again, I'll represent to you this is another screenshot from the QBIC demo produced as IBM 747. And in this screenshot the icon C is being clicked on the first page. Do you see that? A. Yep. Q. And that's highlighted in red. Do you see that? A. Yep. Q. Okay. And the second page are the results of that color histogram query from the first page. A. Correct. Page 127 Q. Okay. And so when the user clicks C on that image in the top left corner, what happens, in the second page? A. The database returned the top 10 -- or it ordered the images and it returned the top K based on the vector distance or the feature distance related to what was computed and stored in the database and what's recorded. This is an N database image. It represents -- it always matches itself perfectly. Q. Okay. So in the -- in the demo that you hosted on IBM's website in the '90s, did it have a similar functionality where you could click on the thumbnail and it would -- it would return matches? A. Yeah. Q. And what would you click on on the one on the website? A. I don't remember exactly how the user interface worked. Q. But you could have -- you could click something akin to color histogram -- A. Yes. Q. -- or texture or things like that? A. Yes. Q. And would it sort similarly to what's Page 128 shown here? A. Yes. Q. Okay. So the querying image is the top left image; is that right? A. I think so, yeah. Q. Well, go back to the first page. A. It's the match, yeah. Q. So you click C on that first image, and when you -- when you do that, that image that's being queried for color histogram remains in the top left box; is that right? A. Yeah. Q. And then the image that's immediately to the right of that, is that the best match out of the images in the database? A. That would -- yes. Q. And then the image immediately to the right of that one is the second best match? A. Yes. Q. And so on and so forth? A. Yes. Q. And when we say "match" we're talking about in terms of the color histogram features? A. Yeah. Q. Okay. Okay. Last one. Page 129 MR. EDWARDS: Last screenshot. I'm going to mark it as Exhibit 8. (Deposition Exhibit 8 was marked.) BY MR. EDWARDS: Q. And, again, Mr. Flickner, I'll represent this is a screenshot from the QBIC demo on the CD IBM 747 produced by your counsel. A. Okay. Q. Okay? And if you want to compare, or you can take my word for it, the image in the top left corner, the querying image, is the same one that was on Exhibit 7. A. Yes. Q. Okay. And here the user's clicking the I icon, which stands for "Get Info." Do you see that? A. Um-hum. Q. On the first page? A. Yeah. Q. And the second page is doing what? A. Is presenting the image at full resolution. Q. Okay. And was there a similar functionality on the demo app on the -- on the website in the '90s? *Id.* at 135 – 136

Element	'278 Claim Recitation	Exemplary Citations in Reference
		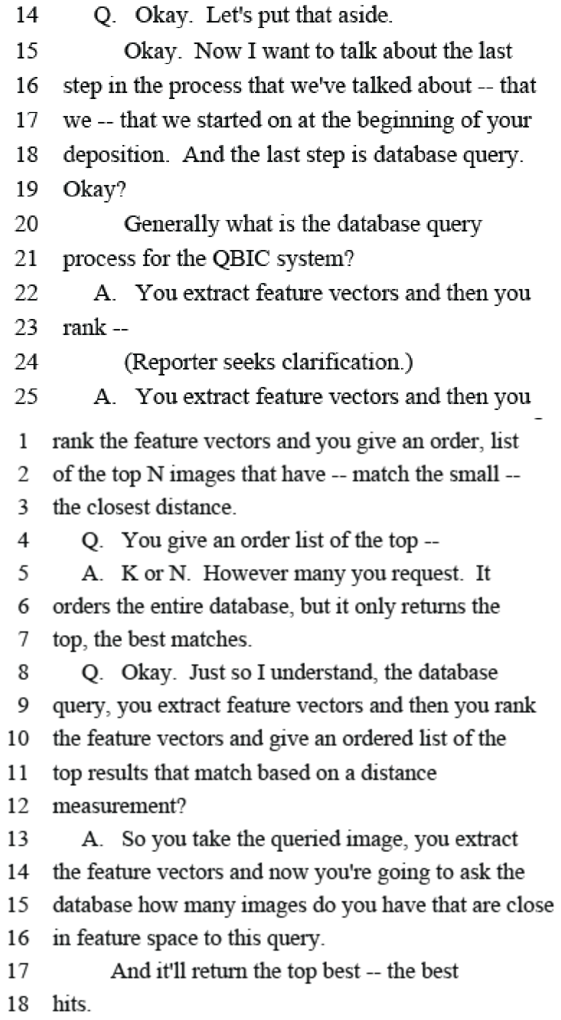 Q. Okay. Let's put that aside. Okay. Now I want to talk about the last step in the process that we've talked about -- that we -- that we started on at the beginning of your deposition. And the last step is database query. Okay? Generally what is the database query process for the QBIC system? A. You extract feature vectors and then you rank -- (Reporter seeks clarification.) A. You extract feature vectors and then you rank the feature vectors and you give an order, list of the top N images that have -- match the small -- the closest distance. Q. You give an order list of the top -- A. K or N. However many you request. It orders the entire database, but it only returns the top, the best matches. Q. Okay. Just so I understand, the database query, you extract feature vectors and then you rank the feature vectors and give an ordered list of the top results that match based on a distance measurement? A. So you take the queried image, you extract the feature vectors and now you're going to ask the database how many images do you have that are close in feature space to this query. And it'll return the top best -- the best hits. *Id.* at 139–165

Element	'278 Claim Recitation	Exemplary Citations in Reference
		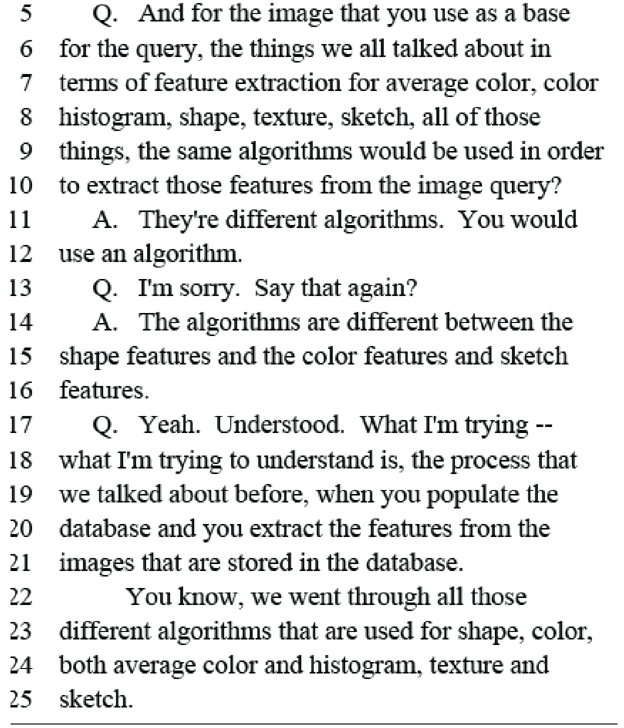 Q. And for the image that you use as a base for the query, the things we all talked about in terms of feature extraction for average color, color histogram, shape, texture, sketch, all of those things, the same algorithms would be used in order to extract those features from the image query? A. They're different algorithms. You would use an algorithm. Q. I'm sorry. Say that again? A. The algorithms are different between the shape features and the color features and sketch features. Q. Yeah. Understood. What I'm trying -- what I'm trying to understand is, the process that we talked about before, when you populate the database and you extract the features from the images that are stored in the database. You know, we went through all those different algorithms that are used for shape, color, both average color and histogram, texture and sketch.

Element	'278 Claim Recitation	Exemplary Citations in Reference
		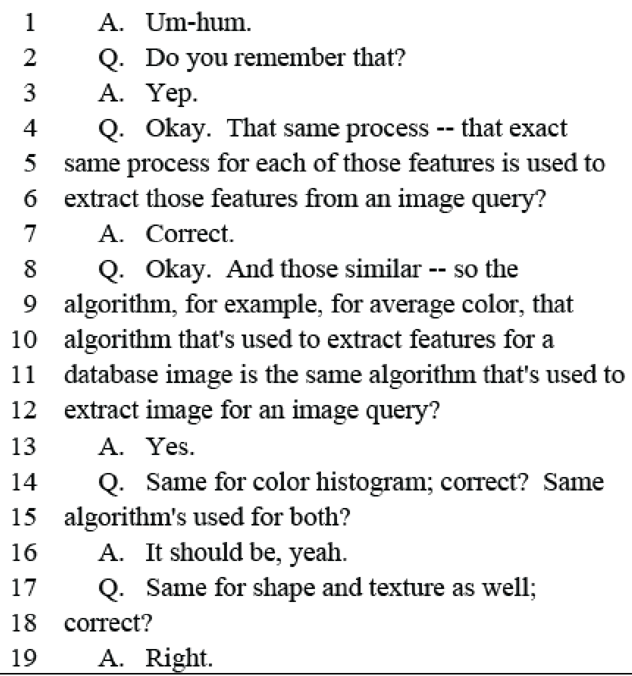 A. Um-hum. Q. Do you remember that? A. Yep. Q. Okay. That same process -- that exact same process for each of those features is used to extract those features from an image query? A. Correct. Q. Okay. And those similar -- so the algorithm, for example, for average color, that algorithm that's used to extract features for a database image is the same algorithm that's used to extract image for an image query? A. Yes. Q. Same for color histogram; correct? Same algorithm's used for both? A. It should be, yeah. Q. Same for shape and texture as well; correct? A. Right.

		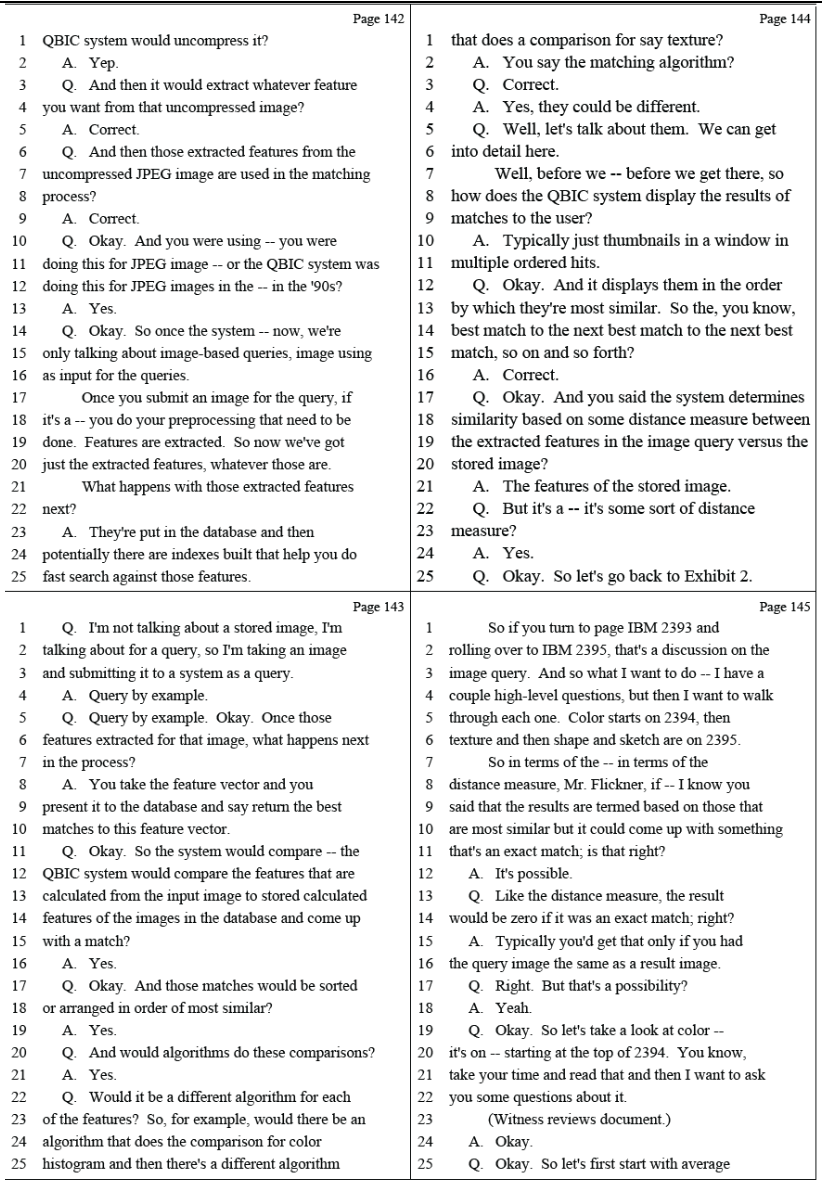 Page 142 QBIC system would uncompress it? A. Yep. Q. And then it would extract whatever feature you want from that uncompressed image? A. Correct. Q. And then those extracted features from the uncompressed JPEG image are used in the matching process? A. Correct. Q. Okay. And you were using -- you were doing this for JPEG image -- or the QBIC system was doing this for JPEG images in the -- in the '90s? A. Yes. Q. Okay. So once the system -- now, we're only talking about image-based queries, image using as input for the queries. Once you submit an image for the query, if it's a -- you do your preprocessing that need to be done. Features are extracted. So now we've got just the extracted features, whatever those are. What happens with those extracted features next? A. They're put in the database and then potentially there are indexes built that help you do fast search against those features. Page 143 Q. I'm not talking about a stored image, I'm talking about for a query, so I'm taking an image and submitting it to a system as a query. A. Query by example. Q. Query by example. Okay. Once those features extracted for that image, what happens next in the process? A. You take the feature vector and you present it to the database and say return the best matches to this feature vector. Q. Okay. So the system would compare -- the QBIC system would compare the features that are calculated from the input image to stored calculated features of the images in the database and come up with a match? A. Yes. Q. Okay. And those matches would be sorted or arranged in order of most similar? A. Yes. Q. And would algorithms do these comparisons? A. Yes. Q. Would it be a different algorithm for each of the features? So, for example, would there be an algorithm that does the comparison for color histogram and then there's a different algorithm Page 144 that does a comparison for say texture? A. You say the matching algorithm? Q. Correct. A. Yes, they could be different. Q. Well, let's talk about them. We can get into detail here. Well, before we -- before we get there, so how does the QBIC system display the results of matches to the user? A. Typically just thumbnails in a window in multiple ordered hits. Q. Okay. And it displays them in the order by which they're most similar. So the, you know, best match to the next best match to the next best match, so on and so forth? A. Correct. Q. Okay. And you said the system determines similarity based on some distance measure between the extracted features in the image query versus the stored image? A. The features of the stored image. Q. But it's a -- it's some sort of distance measure? A. Yes. Q. Okay. So let's go back to Exhibit 2. Page 145 So if you turn to page IBM 2393 and rolling over to IBM 2395, that's a discussion on the image query. And so what I want to do -- I have a couple high-level questions, but then I want to walk through each one. Color starts on 2394, then texture and then shape and sketch are on 2395. So in terms of the -- in terms of the distance measure, Mr. Flickner, if -- I know you said that the results are termed based on those that are most similar but it could come up with something that's an exact match; is that right? A. It's possible. Q. Like the distance measure, the result would be zero if it was an exact match; right? A. Typically you'd get that only if you had the query image the same as a result image. Q. Right. But that's a possibility? A. Yeah. Q. Okay. So let's take a look at color -- it's on -- starting at the top of 2394. You know, take your time and read that and then I want to ask you some questions about it. (Witness reviews document.) A. Okay. Q. Okay. So let's first start with average

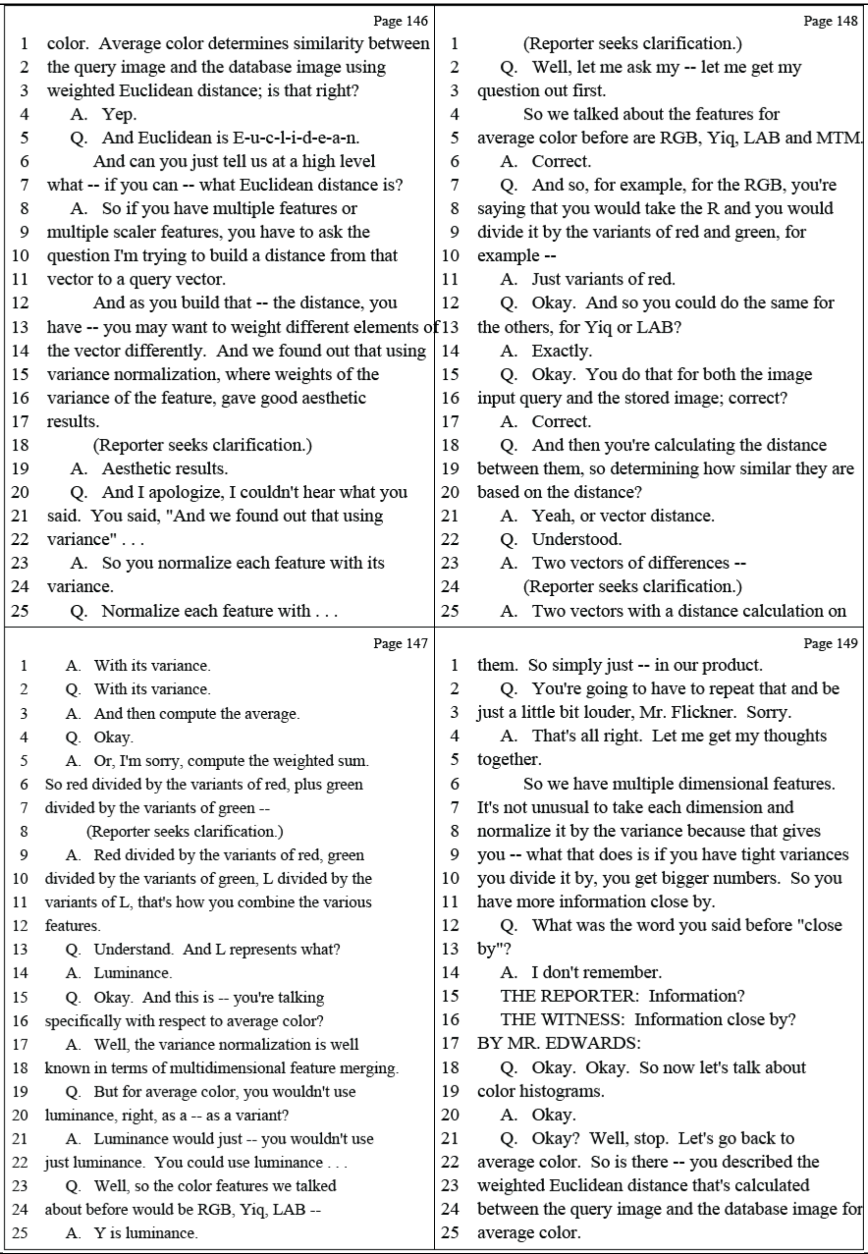

Page 146

color. Average color determines similarity between the query image and the database image using weighted Euclidean distance; is that right?
A. Yep.
Q. And Euclidean is E-u-c-l-i-d-e-a-n.
And can you just tell us at a high level what -- if you can -- what Euclidean distance is?
A. So if you have multiple features or multiple scaler features, you have to ask the question I'm trying to build a distance from that vector to a query vector.
And as you build that -- the distance, you have -- you may want to weight different elements of the vector differently. And we found out that using variance normalization, where weights of the variance of the feature, gave good aesthetic results.
(Reporter seeks clarification.)
A. Aesthetic results.
Q. And I apologize, I couldn't hear what you said. You said, "And we found out that using variance" . . .
A. So you normalize each feature with its variance.
Q. Normalize each feature with . . .

Page 147

A. With its variance.
Q. With its variance.
A. And then compute the average.
Q. Okay.
A. Or, I'm sorry, compute the weighted sum. So red divided by the variants of red, plus green divided by the variants of green --
(Reporter seeks clarification.)
A. Red divided by the variants of red, green divided by the variants of green, L divided by the variants of L, that's how you combine the various features.
Q. Understand. And L represents what?
A. Luminance.
Q. Okay. And this is -- you're talking specifically with respect to average color?
A. Well, the variance normalization is well known in terms of multidimensional feature merging.
Q. But for average color, you wouldn't use luminance, right, as a -- as a variant?
A. Luminance would just -- you wouldn't use just luminance. You could use luminance . . .
Q. Well, so the color features we talked about before would be RGB, Yiq, LAB --
A. Y is luminance.

Page 148

(Reporter seeks clarification.)
Q. Well, let me ask my -- let me get my question out first.
So we talked about the features for average color before are RGB, Yiq, LAB and MTM.
A. Correct.
Q. And so, for example, for the RGB, you're saying that you would take the R and you would divide it by the variants of red and green, for example --
A. Just variants of red.
Q. Okay. And so you could do the same for the others, for Yiq or LAB?
A. Exactly.
Q. Okay. You do that for both the image input query and the stored image; correct?
A. Correct.
Q. And then you're calculating the distance between them, so determining how similar they are based on the distance?
A. Yeah, or vector distance.
Q. Understood.
A. Two vectors of differences --
(Reporter seeks clarification.)
A. Two vectors with a distance calculation on

Page 149

them. So simply just -- in our product.
Q. You're going to have to repeat that and be just a little bit louder, Mr. Flickner. Sorry.
A. That's all right. Let me get my thoughts together.
So we have multiple dimensional features. It's not unusual to take each dimension and normalize it by the variance because that gives you -- what that does is if you have tight variances you divide it by, you get bigger numbers. So you have more information close by.
Q. What was the word you said before "close by"?
A. I don't remember.
THE REPORTER: Information?
THE WITNESS: Information close by?
BY MR. EDWARDS:
Q. Okay. Okay. So now let's talk about color histograms.
A. Okay.
Q. Okay? Well, stop. Let's go back to average color. So is there -- you described the weighted Euclidean distance that's calculated between the query image and the database image for average color.

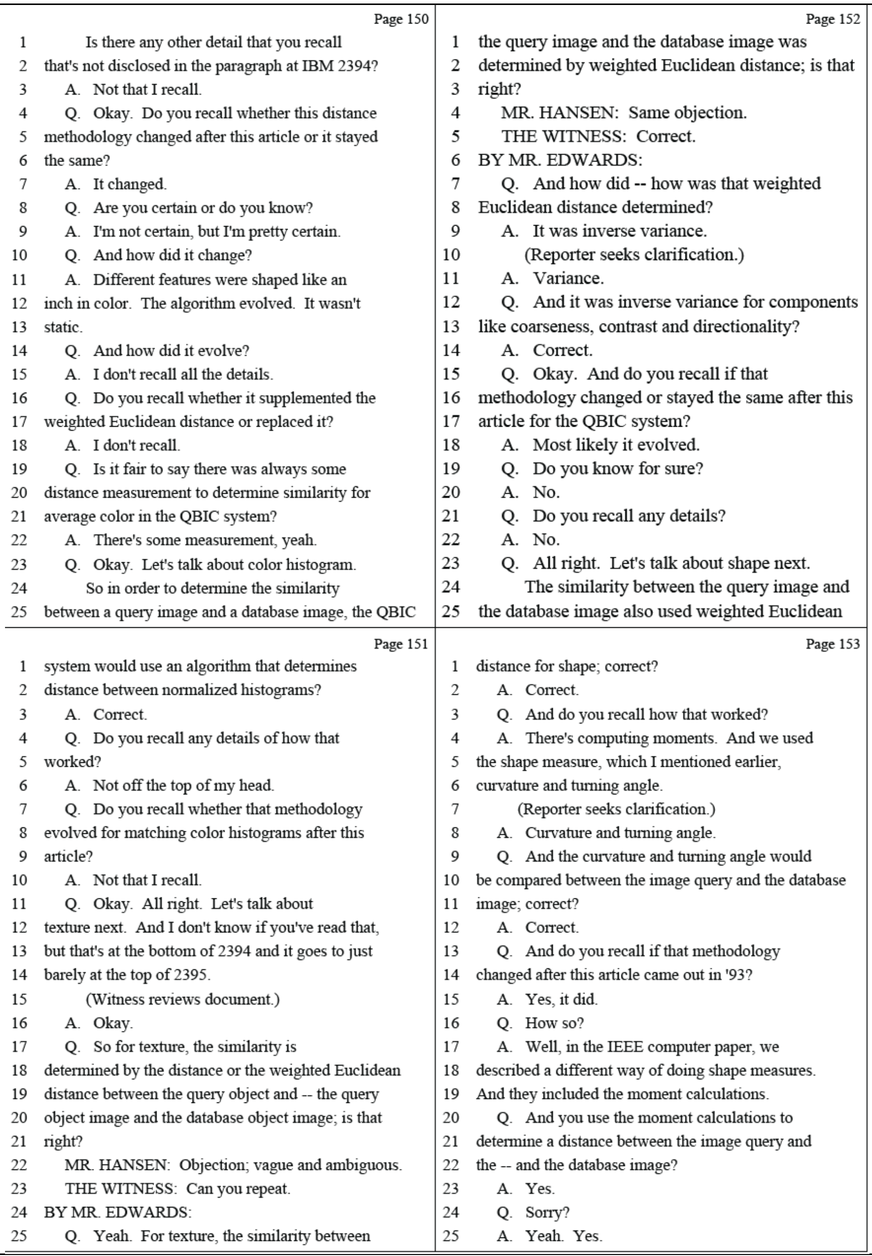

Page 150

Is there any other detail that you recall that's not disclosed in the paragraph at IBM 2394?
A. Not that I recall.
Q. Okay. Do you recall whether this distance methodology changed after this article or it stayed the same?
A. It changed.
Q. Are you certain or do you know?
A. I'm not certain, but I'm pretty certain.
Q. And how did it change?
A. Different features were shaped like an inch in color. The algorithm evolved. It wasn't static.
Q. And how did it evolve?
A. I don't recall all the details.
Q. Do you recall whether it supplemented the weighted Euclidean distance or replaced it?
A. I don't recall.
Q. Is it fair to say there was always some distance measurement to determine similarity for average color in the QBIC system?
A. There's some measurement, yeah.
Q. Okay. Let's talk about color histogram.
So in order to determine the similarity between a query image and a database image, the QBIC

Page 151

system would use an algorithm that determines distance between normalized histograms?
A. Correct.
Q. Do you recall any details of how that worked?
A. Not off the top of my head.
Q. Do you recall whether that methodology evolved for matching color histograms after this article?
A. Not that I recall.
Q. Okay. All right. Let's talk about texture next. And I don't know if you've read that, but that's at the bottom of 2394 and it goes to just barely at the top of 2395.
(Witness reviews document.)
A. Okay.
Q. So for texture, the similarity is determined by the distance or the weighted Euclidean distance between the query object and -- the query object image and the database object image; is that right?
MR. HANSEN: Objection; vague and ambiguous.
THE WITNESS: Can you repeat.
BY MR. EDWARDS:
Q. Yeah. For texture, the similarity between

Page 152

the query image and the database image was determined by weighted Euclidean distance; is that right?
MR. HANSEN: Same objection.
THE WITNESS: Correct.
BY MR. EDWARDS:
Q. And how did -- how was that weighted Euclidean distance determined?
A. It was inverse variance.
(Reporter seeks clarification.)
A. Variance.
Q. And it was inverse variance for components like coarseness, contrast and directionality?
A. Correct.
Q. Okay. And do you recall if that methodology changed or stayed the same after this article for the QBIC system?
A. Most likely it evolved.
Q. Do you know for sure?
A. No.
Q. Do you recall any details?
A. No.
Q. All right. Let's talk about shape next.
The similarity between the query image and the database image also used weighted Euclidean

Page 153

distance for shape; correct?
A. Correct.
Q. And do you recall how that worked?
A. There's computing moments. And we used the shape measure, which I mentioned earlier, curvature and turning angle.
(Reporter seeks clarification.)
A. Curvature and turning angle.
Q. And the curvature and turning angle would be compared between the image query and the database image; correct?
A. Correct.
Q. And do you recall if that methodology changed after this article came out in '93?
A. Yes, it did.
Q. How so?
A. Well, in the IEEE computer paper, we described a different way of doing shape measures. And they included the moment calculations.
Q. And you use the moment calculations to determine a distance between the image query and the -- and the database image?
A. Yes.
Q. Sorry?
A. Yeah. Yes.

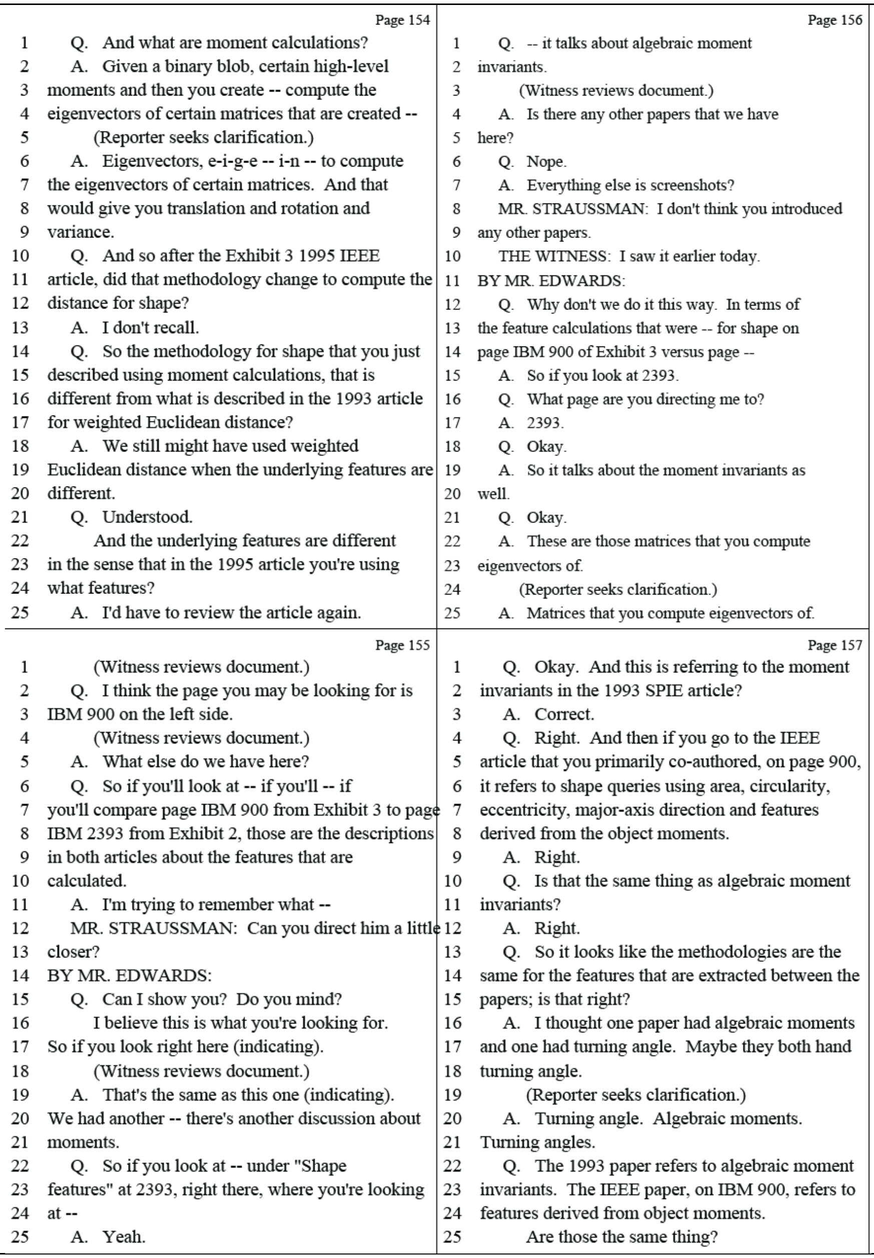

Page 154

Q. And what are moment calculations?
A. Given a binary blob, certain high-level moments and then you create -- compute the eigenvectors of certain matrices that are created --
(Reporter seeks clarification.)
A. Eigenvectors, e-i-g-e -- i-n -- to compute the eigenvectors of certain matrices. And that would give you translation and rotation and variance.
Q. And so after the Exhibit 3 1995 IEEE article, did that methodology change to compute the distance for shape?
A. I don't recall.
Q. So the methodology for shape that you just described using moment calculations, that is different from what is described in the 1993 article for weighted Euclidean distance?
A. We still might have used weighted Euclidean distance when the underlying features are different.
Q. Understood.
And the underlying features are different in the sense that in the 1995 article you're using what features?
A. I'd have to review the article again.

Page 155

(Witness reviews document.)
Q. I think the page you may be looking for is IBM 900 on the left side.
(Witness reviews document.)
A. What else do we have here?
Q. So if you'll look at -- if you'll -- if you'll compare page IBM 900 from Exhibit 3 to page IBM 2393 from Exhibit 2, those are the descriptions in both articles about the features that are calculated.
A. I'm trying to remember what --
MR. STRAUSSMAN: Can you direct him a little closer?
BY MR. EDWARDS:
Q. Can I show you? Do you mind?
I believe this is what you're looking for. So if you look right here (indicating).
(Witness reviews document.)
A. That's the same as this one (indicating). We had another -- there's another discussion about moments.
Q. So if you look at -- under "Shape features" at 2393, right there, where you're looking at --
A. Yeah.

Page 156

Q. -- it talks about algebraic moment invariants.
(Witness reviews document.)
A. Is there any other papers that we have here?
Q. Nope.
A. Everything else is screenshots?
MR. STRAUSSMAN: I don't think you introduced any other papers.
THE WITNESS: I saw it earlier today.
BY MR. EDWARDS:
Q. Why don't we do it this way. In terms of the feature calculations that were -- for shape on page IBM 900 of Exhibit 3 versus page --
A. So if you look at 2393.
Q. What page are you directing me to?
A. 2393.
Q. Okay.
A. So it talks about the moment invariants as well.
Q. Okay.
A. These are those matrices that you compute eigenvectors of.
(Reporter seeks clarification.)
A. Matrices that you compute eigenvectors of.

Page 157

Q. Okay. And this is referring to the moment invariants in the 1993 SPIE article?
A. Correct.
Q. Right. And then if you go to the IEEE article that you primarily co-authored, on page 900, it refers to shape queries using area, circularity, eccentricity, major-axis direction and features derived from the object moments.
A. Right.
Q. Is that the same thing as algebraic moment invariants?
A. Right.
Q. So it looks like the methodologies are the same for the features that are extracted between the papers; is that right?
A. I thought one paper had algebraic moments and one had turning angle. Maybe they both hand turning angle.
(Reporter seeks clarification.)
A. Turning angle. Algebraic moments. Turning angles.
Q. The 1993 paper refers to algebraic moment invariants. The IEEE paper, on IBM 900, refers to features derived from object moments.
Are those the same thing?

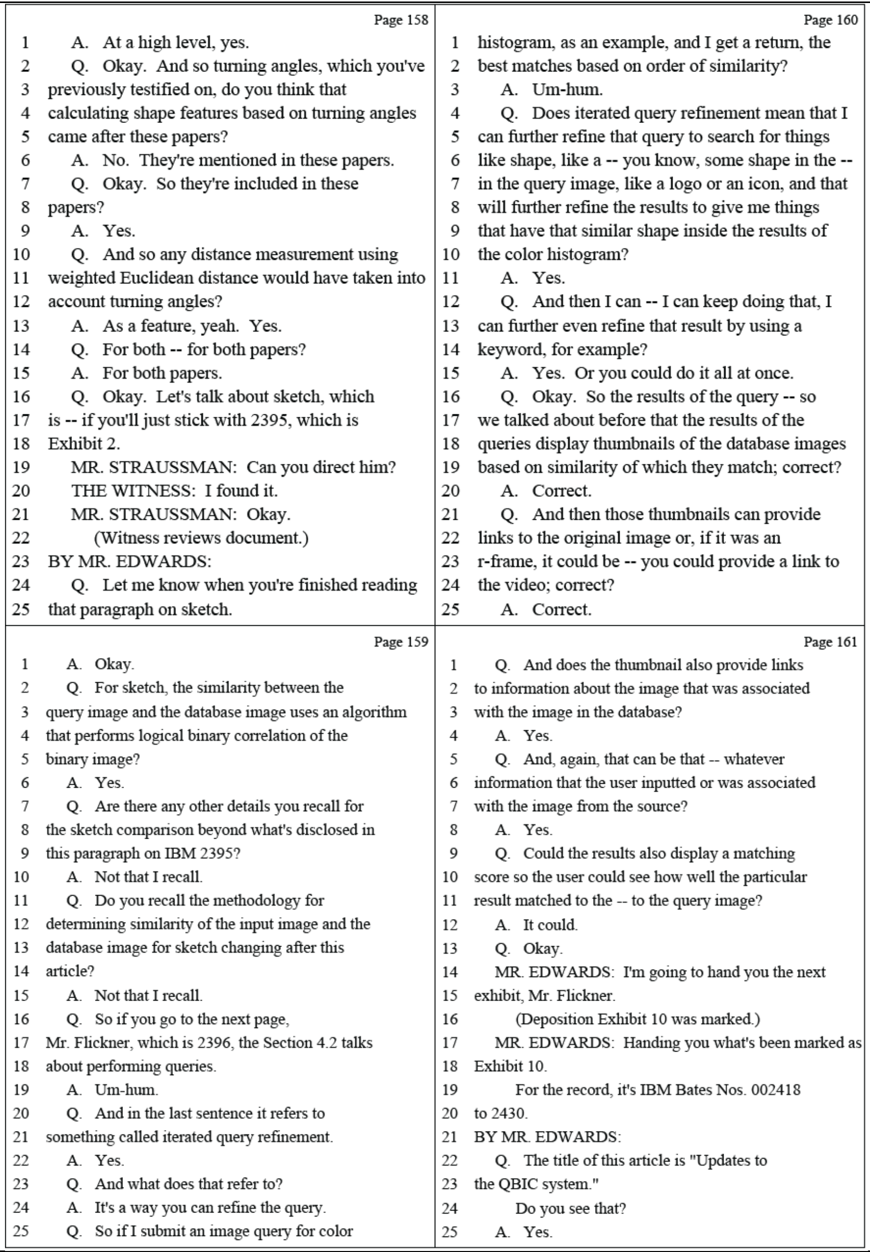

Page 158

A. At a high level, yes.
Q. Okay. And so turning angles, which you've previously testified on, do you think that calculating shape features based on turning angles came after these papers?
A. No. They're mentioned in these papers.
Q. Okay. So they're included in these papers?
A. Yes.
Q. And so any distance measurement using weighted Euclidean distance would have taken into account turning angles?
A. As a feature, yeah. Yes.
Q. For both -- for both papers?
A. For both papers.
Q. Okay. Let's talk about sketch, which is -- if you'll just stick with 2395, which is Exhibit 2.
MR. STRAUSSMAN: Can you direct him?
THE WITNESS: I found it.
MR. STRAUSSMAN: Okay.
(Witness reviews document.)
BY MR. EDWARDS:
Q. Let me know when you're finished reading that paragraph on sketch.

Page 159

A. Okay.
Q. For sketch, the similarity between the query image and the database image uses an algorithm that performs logical binary correlation of the binary image?
A. Yes.
Q. Are there any other details you recall for the sketch comparison beyond what's disclosed in this paragraph on IBM 2395?
A. Not that I recall.
Q. Do you recall the methodology for determining similarity of the input image and the database image for sketch changing after this article?
A. Not that I recall.
Q. So if you go to the next page, Mr. Flickner, which is 2396, the Section 4.2 talks about performing queries.
A. Um-hum.
Q. And in the last sentence it refers to something called iterated query refinement.
A. Yes.
Q. And what does that refer to?
A. It's a way you can refine the query.
Q. So if I submit an image query for color

Page 160

histogram, as an example, and I get a return, the best matches based on order of similarity?
A. Um-hum.
Q. Does iterated query refinement mean that I can further refine that query to search for things like shape, like a -- you know, some shape in the -- in the query image, like a logo or an icon, and that will further refine the results to give me things that have that similar shape inside the results of the color histogram?
A. Yes.
Q. And then I can -- I can keep doing that, I can further even refine that result by using a keyword, for example?
A. Yes. Or you could do it all at once.
Q. Okay. So the results of the query -- so we talked about before that the results of the queries display thumbnails of the database images based on similarity of which they match; correct?
A. Correct.
Q. And then those thumbnails can provide links to the original image or, if it was an r-frame, it could be -- you could provide a link to the video; correct?
A. Correct.

Page 161

Q. And does the thumbnail also provide links to information about the image that was associated with the image in the database?
A. Yes.
Q. And, again, that can be that -- whatever information that the user inputted or was associated with the image from the source?
A. Yes.
Q. Could the results also display a matching score so the user could see how well the particular result matched to the -- to the query image?
A. It could.
Q. Okay.
MR. EDWARDS: I'm going to hand you the next exhibit, Mr. Flickner.
(Deposition Exhibit 10 was marked.)
MR. EDWARDS: Handing you what's been marked as Exhibit 10.
For the record, it's IBM Bates Nos. 002418 to 2430.
BY MR. EDWARDS:
Q. The title of this article is "Updates to the QBIC system."
Do you see that?
A. Yes.

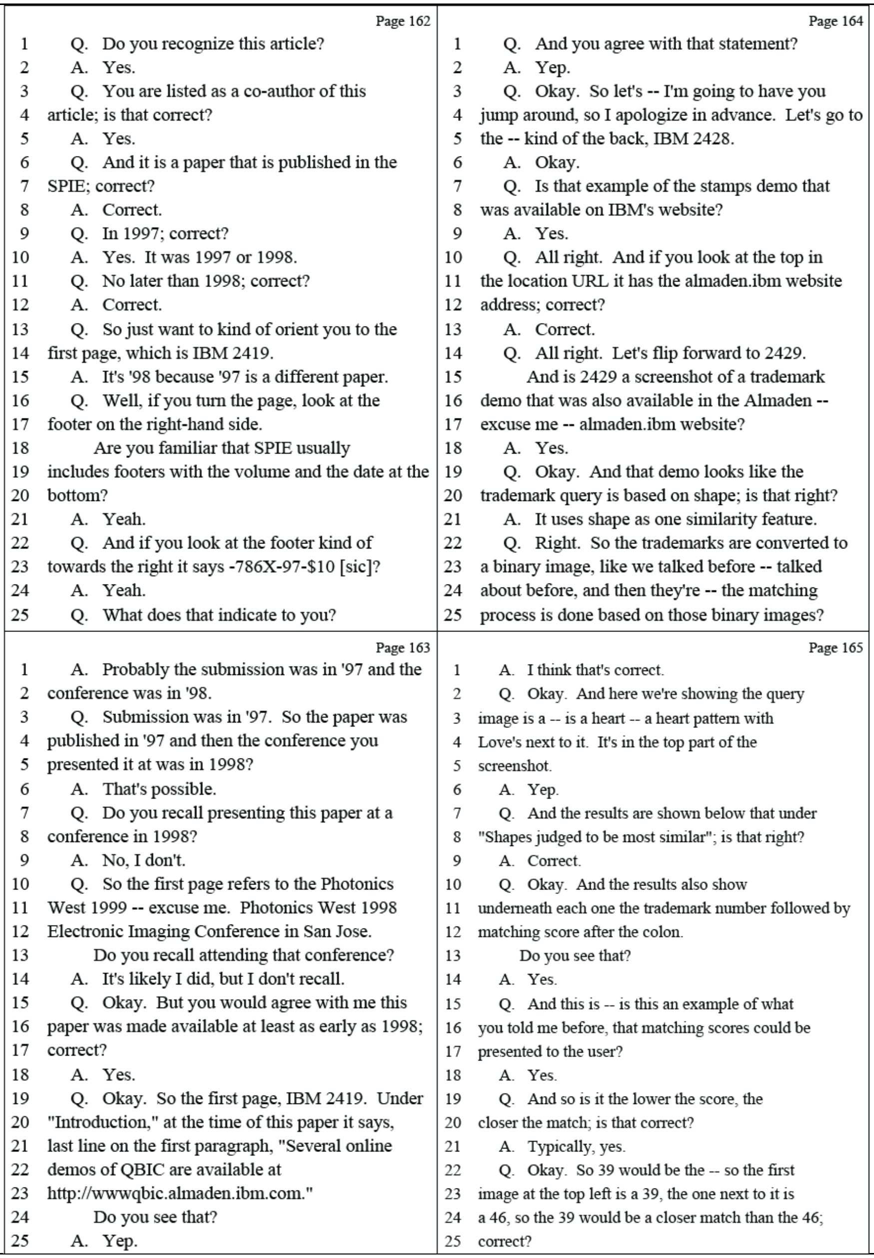

Page 162

Q. Do you recognize this article?
A. Yes.
Q. You are listed as a co-author of this article; is that correct?
A. Yes.
Q. And it is a paper that is published in the SPIE; correct?
A. Correct.
Q. In 1997; correct?
A. Yes. It was 1997 or 1998.
Q. No later than 1998; correct?
A. Correct.
Q. So just want to kind of orient you to the first page, which is IBM 2419.
A. It's '98 because '97 is a different paper.
Q. Well, if you turn the page, look at the footer on the right-hand side.
Are you familiar that SPIE usually includes footers with the volume and the date at the bottom?
A. Yeah.
Q. And if you look at the footer kind of towards the right it says -786X-97-$10 [sic]?
A. Yeah.
Q. What does that indicate to you?

Page 163

A. Probably the submission was in '97 and the conference was in '98.
Q. Submission was in '97. So the paper was published in '97 and then the conference you presented it at was in 1998?
A. That's possible.
Q. Do you recall presenting this paper at a conference in 1998?
A. No, I don't.
Q. So the first page refers to the Photonics West 1999 -- excuse me. Photonics West 1998 Electronic Imaging Conference in San Jose.
Do you recall attending that conference?
A. It's likely I did, but I don't recall.
Q. Okay. But you would agree with me this paper was made available at least as early as 1998; correct?
A. Yes.
Q. Okay. So the first page, IBM 2419. Under "Introduction," at the time of this paper it says, last line on the first paragraph, "Several online demos of QBIC are available at http://wwwqbic.almaden.ibm.com."
Do you see that?
A. Yep.

Page 164

Q. And you agree with that statement?
A. Yep.
Q. Okay. So let's -- I'm going to have you jump around, so I apologize in advance. Let's go to the -- kind of the back, IBM 2428.
A. Okay.
Q. Is that example of the stamps demo that was available on IBM's website?
A. Yes.
Q. All right. And if you look at the top in the location URL it has the almaden.ibm website address; correct?
A. Correct.
Q. All right. Let's flip forward to 2429.
And is 2429 a screenshot of a trademark demo that was also available in the Almaden -- excuse me -- almaden.ibm website?
A. Yes.
Q. Okay. And that demo looks like the trademark query is based on shape; is that right?
A. It uses shape as one similarity feature.
Q. Right. So the trademarks are converted to a binary image, like we talked before -- talked about before, and then they're -- the matching process is done based on those binary images?

Page 165

A. I think that's correct.
Q. Okay. And here we're showing the query image is a -- is a heart -- a heart pattern with Love's next to it. It's in the top part of the screenshot.
A. Yep.
Q. And the results are shown below that under "Shapes judged to be most similar"; is that right?
A. Correct.
Q. Okay. And the results also show underneath each one the trademark number followed by matching score after the colon.
Do you see that?
A. Yes.
Q. And this is -- is this an example of what you told me before, that matching scores could be presented to the user?
A. Yes.
Q. And so is it the lower the score, the closer the match; is that correct?
A. Typically, yes.
Q. Okay. So 39 would be the -- so the first image at the top left is a 39, the one next to it is a 46, so the 39 would be a closer match than the 46; correct?

Element	'278 Claim Recitation	Exemplary Citations in Reference
		Id. at 174–182 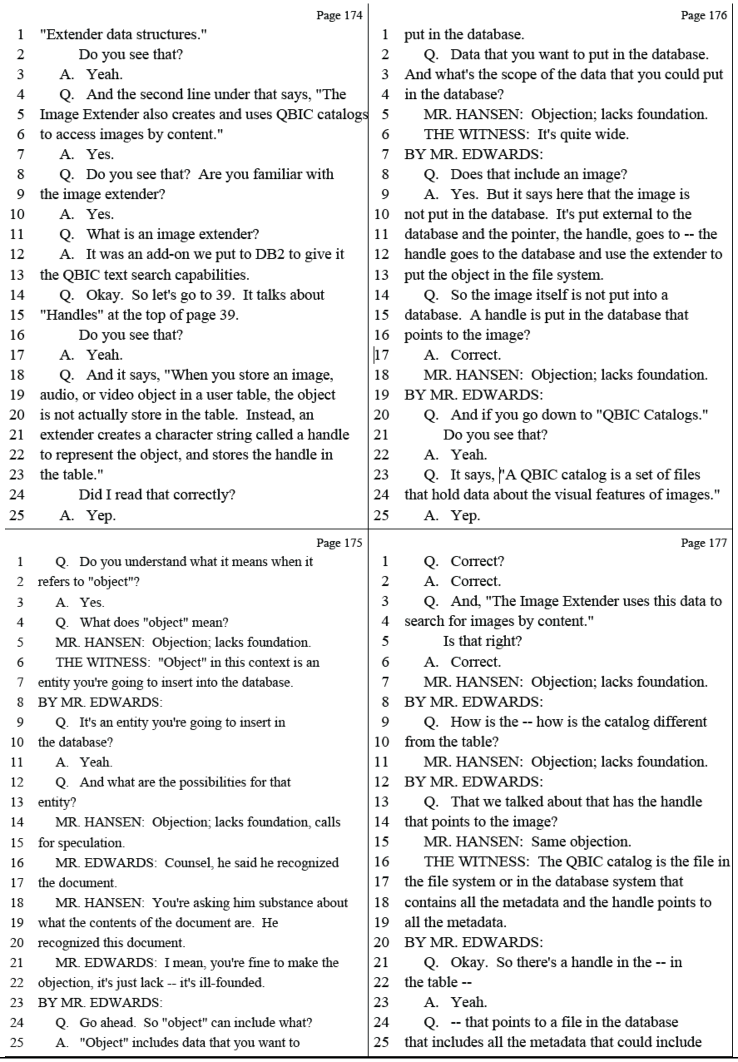

Page 174

"Extender data structures."

Do you see that?

A. Yeah.

Q. And the second line under that says, "The Image Extender also creates and uses QBIC catalogs to access images by content."

A. Yes.

Q. Do you see that? Are you familiar with the image extender?

A. Yes.

Q. What is an image extender?

A. It was an add-on we put to DB2 to give it the QBIC text search capabilities.

Q. Okay. So let's go to 39. It talks about "Handles" at the top of page 39.

Do you see that?

A. Yeah.

Q. And it says, "When you store an image, audio, or video object in a user table, the object is not actually store in the table. Instead, an extender creates a character string called a handle to represent the object, and stores the handle in the table."

Did I read that correctly?

A. Yep.

Page 175

Q. Do you understand what it means when it refers to "object"?

A. Yes.

Q. What does "object" mean?

MR. HANSEN: Objection; lacks foundation.

THE WITNESS: "Object" in this context is an entity you're going to insert into the database.

BY MR. EDWARDS:

Q. It's an entity you're going to insert in the database?

A. Yeah.

Q. And what are the possibilities for that entity?

MR. HANSEN: Objection; lacks foundation, calls for speculation.

MR. EDWARDS: Counsel, he said he recognized the document.

MR. HANSEN: You're asking him substance about what the contents of the document are. He recognized this document.

MR. EDWARDS: I mean, you're fine to make the objection, it's just lack -- it's ill-founded.

BY MR. EDWARDS:

Q. Go ahead. So "object" can include what?

A. "Object" includes data that you want to

Page 176

put in the database.

Q. Data that you want to put in the database. And what's the scope of the data that you could put in the database?

MR. HANSEN: Objection; lacks foundation.

THE WITNESS: It's quite wide.

BY MR. EDWARDS:

Q. Does that include an image?

A. Yes. But it says here that the image is not put in the database. It's put external to the database and the pointer, the handle, goes to -- the handle goes to the database and use the extender to put the object in the file system.

Q. So the image itself is not put into a database. A handle is put in the database that points to the image?

A. Correct.

MR. HANSEN: Objection; lacks foundation.

BY MR. EDWARDS:

Q. And if you go down to "QBIC Catalogs."

Do you see that?

A. Yeah.

Q. It says, "A QBIC catalog is a set of files that hold data about the visual features of images."

A. Yep.

Page 177

Q. Correct?

A. Correct.

Q. And, "The Image Extender uses this data to search for images by content."

Is that right?

A. Correct.

MR. HANSEN: Objection; lacks foundation.

BY MR. EDWARDS:

Q. How is the -- how is the catalog different from the table?

MR. HANSEN: Objection; lacks foundation.

BY MR. EDWARDS:

Q. That we talked about that has the handle that points to the image?

MR. HANSEN: Same objection.

THE WITNESS: The QBIC catalog is the file in the file system or in the database system that contains all the metadata and the handle points to all the metadata.

BY MR. EDWARDS:

Q. Okay. So there's a handle in the -- in the table --

A. Yeah.

Q. -- that points to a file in the database that includes all the metadata that could include

			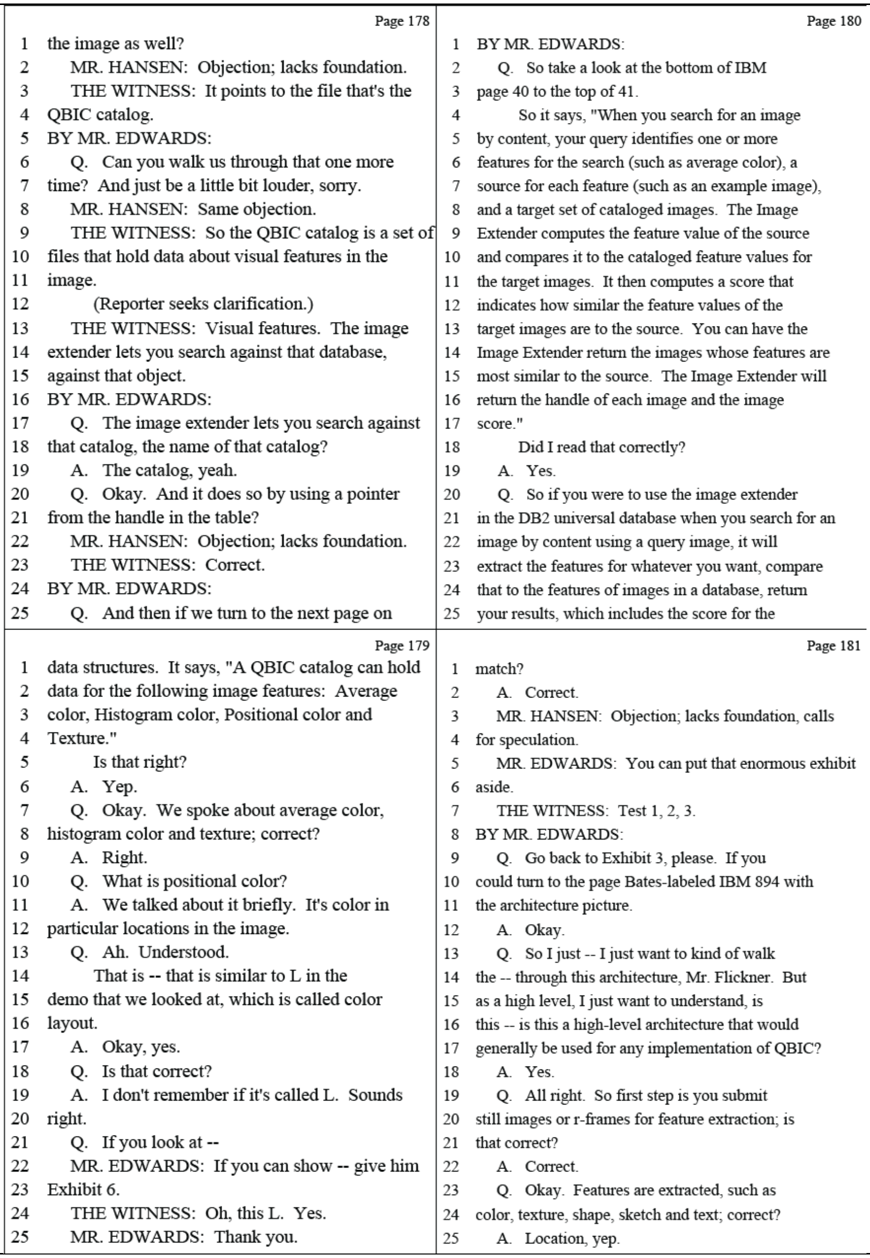 Page 178 the image as well? MR. HANSEN: Objection; lacks foundation. THE WITNESS: It points to the file that's the QBIC catalog. BY MR. EDWARDS: Q. Can you walk us through that one more time? And just be a little bit louder, sorry. MR. HANSEN: Same objection. THE WITNESS: So the QBIC catalog is a set of files that hold data about visual features in the image. (Reporter seeks clarification.) THE WITNESS: Visual features. The image extender lets you search against that database, against that object. BY MR. EDWARDS: Q. The image extender lets you search against that catalog, the name of that catalog? A. The catalog, yeah. Q. Okay. And it does so by using a pointer from the handle in the table? MR. HANSEN: Objection; lacks foundation. THE WITNESS: Correct. BY MR. EDWARDS: Q. And then if we turn to the next page on Page 179 data structures. It says, "A QBIC catalog can hold data for the following image features: Average color, Histogram color, Positional color and Texture." Is that right? A. Yep. Q. Okay. We spoke about average color, histogram color and texture; correct? A. Right. Q. What is positional color? A. We talked about it briefly. It's color in particular locations in the image. Q. Ah. Understood. That is -- that is similar to L in the demo that we looked at, which is called color layout. A. Okay, yes. Q. Is that correct? A. I don't remember if it's called L. Sounds right. Q. If you look at -- MR. EDWARDS: If you can show -- give him Exhibit 6. THE WITNESS: Oh, this L. Yes. MR. EDWARDS: Thank you. Page 180 BY MR. EDWARDS: Q. So take a look at the bottom of IBM page 40 to the top of 41. So it says, "When you search for an image by content, your query identifies one or more features for the search (such as average color), a source for each feature (such as an example image), and a target set of cataloged images. The Image Extender computes the feature value of the source and compares it to the cataloged feature values for the target images. It then computes a score that indicates how similar the feature values of the target images are to the source. You can have the Image Extender return the images whose features are most similar to the source. The Image Extender will return the handle of each image and the image score." Did I read that correctly? A. Yes. Q. So if you were to use the image extender in the DB2 universal database when you search for an image by content using a query image, it will extract the features for whatever you want, compare that to the features of images in a database, return your results, which includes the score for the Page 181 match? A. Correct. MR. HANSEN: Objection; lacks foundation, calls for speculation. MR. EDWARDS: You can put that enormous exhibit aside. THE WITNESS: Test 1, 2, 3. BY MR. EDWARDS: Q. Go back to Exhibit 3, please. If you could turn to the page Bates-labeled IBM 894 with the architecture picture. A. Okay. Q. So I just -- I just want to kind of walk the -- through this architecture, Mr. Flickner. But as a high level, I just want to understand, is this -- is this a high-level architecture that would generally be used for any implementation of QBIC? A. Yes. Q. All right. So first step is you submit still images or r-frames for feature extraction; is that correct? A. Correct. Q. Okay. Features are extracted, such as color, texture, shape, sketch and text; correct? A. Location, yep.

Element	'278 Claim Recitation	Exemplary Citations in Reference
		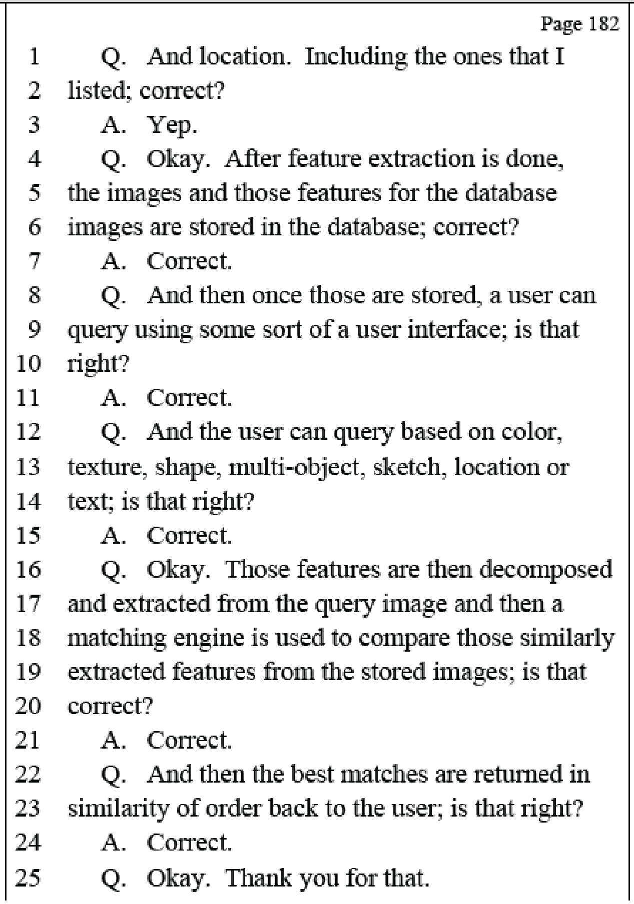 Page 182 Q. And location. Including the ones that I listed; correct? A. Yep. Q. Okay. After feature extraction is done, the images and those features for the database images are stored in the database; correct? A. Correct. Q. And then once those are stored, a user can query using some sort of a user interface; is that right? A. Correct. Q. And the user can query based on color, texture, shape, multi-object, sketch, location or text; is that right? A. Correct. Q. Okay. Those features are then decomposed and extracted from the query image and then a matching engine is used to compare those similarly extracted features from the stored images; is that correct? A. Correct. Q. And then the best matches are returned in similarity of order back to the user; is that right? A. Correct. Q. Okay. Thank you for that. *Id.* at 190–191 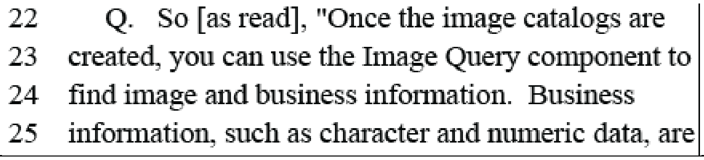 Q. So [as read], "Once the image catalogs are created, you can use the Image Query component to find image and business information. Business information, such as character and numeric data, are

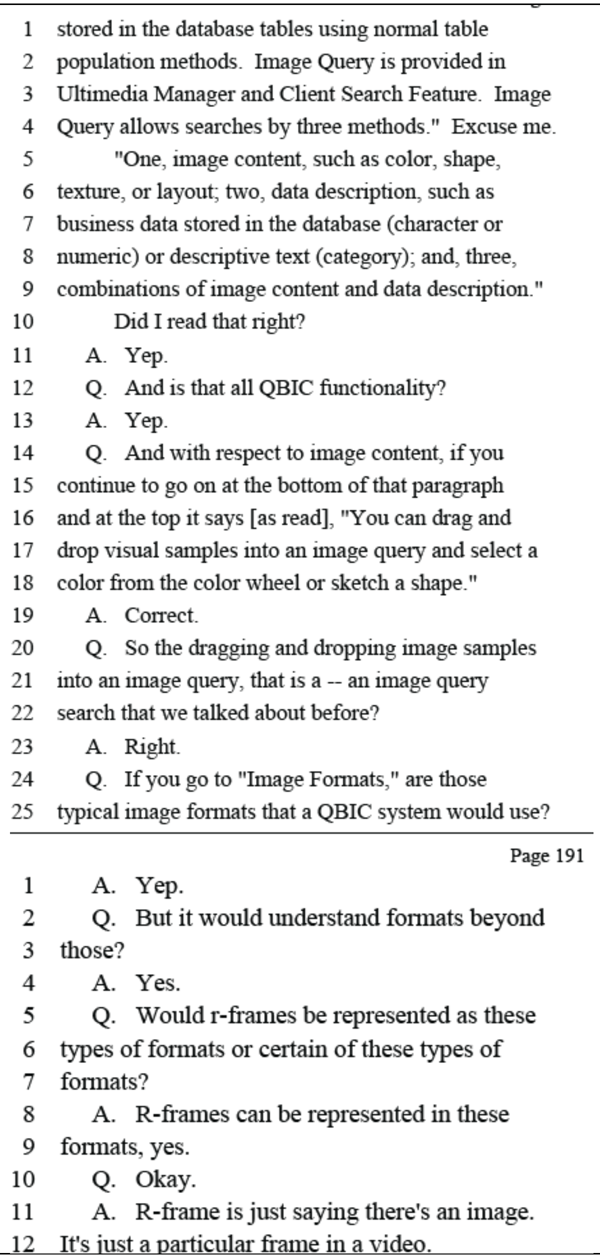

stored in the database tables using normal table population methods. Image Query is provided in Ultimedia Manager and Client Search Feature. Image Query allows searches by three methods." Excuse me.

"One, image content, such as color, shape, texture, or layout; two, data description, such as business data stored in the database (character or numeric) or descriptive text (category); and, three, combinations of image content and data description."

Did I read that right?

A. Yep.

Q. And is that all QBIC functionality?

A. Yep.

Q. And with respect to image content, if you continue to go on at the bottom of that paragraph and at the top it says [as read], "You can drag and drop visual samples into an image query and select a color from the color wheel or sketch a shape."

A. Correct.

Q. So the dragging and dropping image samples into an image query, that is a -- an image query search that we talked about before?

A. Right.

Q. If you go to "Image Formats," are those typical image formats that a QBIC system would use?

Page 191

A. Yep.

Q. But it would understand formats beyond those?

A. Yes.

Q. Would r-frames be represented as these types of formats or certain of these types of formats?

A. R-frames can be represented in these formats, yes.

Q. Okay.

A. R-frame is just saying there's an image. It's just a particular frame in a video.

Element	'278 Claim Recitation	Exemplary Citations in Reference
		Id. at 198 – 199 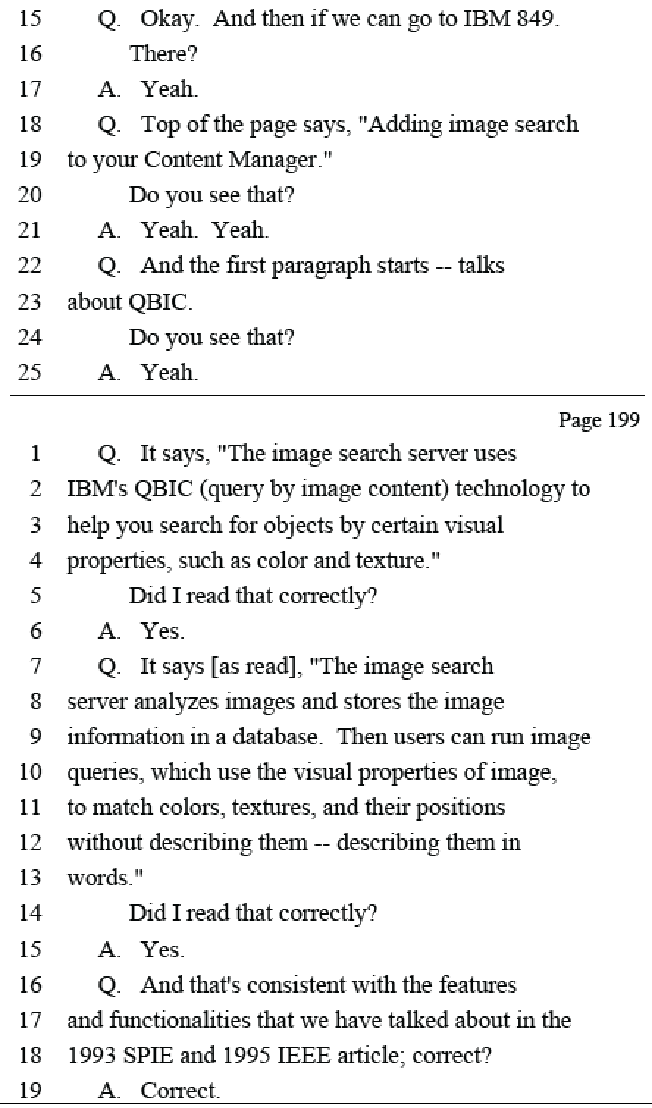 Q. Okay. And then if we can go to IBM 849. There? A. Yeah. Q. Top of the page says, "Adding image search to your Content Manager." Do you see that? A. Yeah. Yeah. Q. And the first paragraph starts -- talks about QBIC. Do you see that? A. Yeah. Page 199 Q. It says, "The image search server uses IBM's QBIC (query by image content) technology to help you search for objects by certain visual properties, such as color and texture." Did I read that correctly? A. Yes. Q. It says [as read], "The image search server analyzes images and stores the image information in a database. Then users can run image queries, which use the visual properties of image, to match colors, textures, and their positions without describing them -- describing them in words." Did I read that correctly? A. Yes. Q. And that's consistent with the features and functionalities that we have talked about in the 1993 SPIE and 1995 IEEE article; correct? A. Correct.

Element	'278 Claim Recitation	Exemplary Citations in Reference
		Id. at 312 – 313 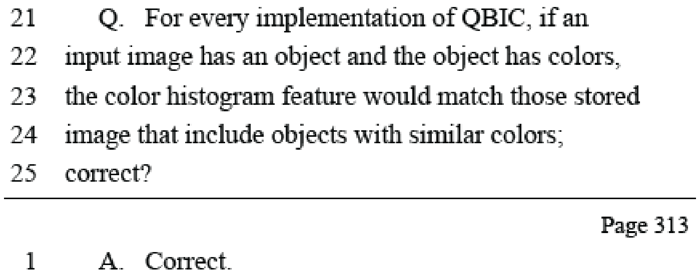 Q. For every implementation of QBIC, if an input image has an object and the object has colors, the color histogram feature would match those stored image that include objects with similar colors; correct? Page 313 A. Correct. *Id.* at 317–322 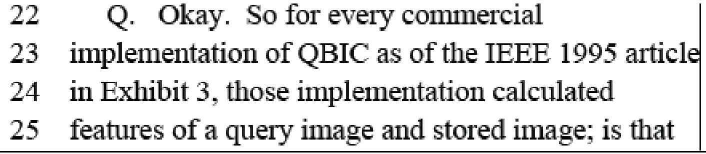 Q. Okay. So for every commercial implementation of QBIC as of the IEEE 1995 article in Exhibit 3, those implementation calculated features of a query image and stored image; is that

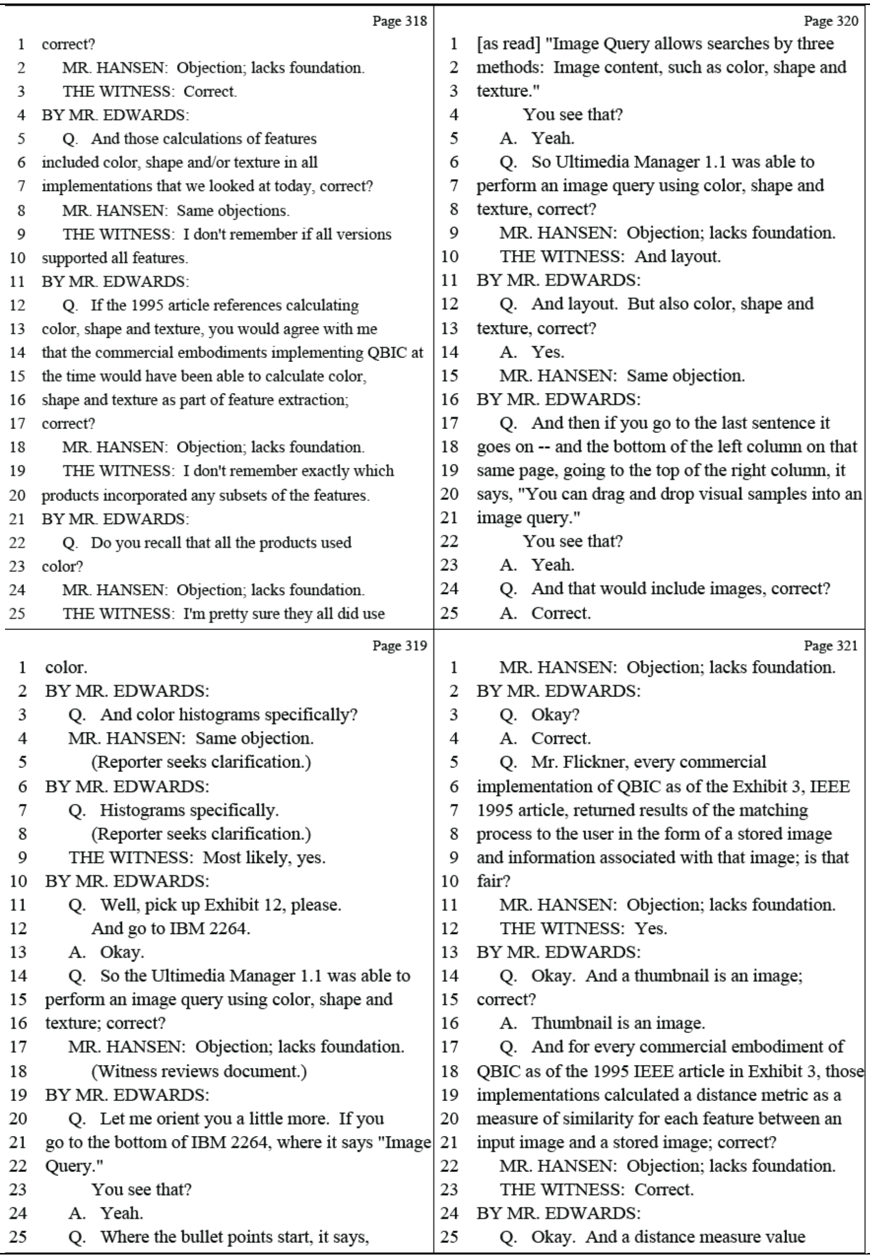

Page 318

correct?
MR. HANSEN: Objection; lacks foundation.
THE WITNESS: Correct.
BY MR. EDWARDS:
Q. And those calculations of features included color, shape and/or texture in all implementations that we looked at today, correct?
MR. HANSEN: Same objections.
THE WITNESS: I don't remember if all versions supported all features.
BY MR. EDWARDS:
Q. If the 1995 article references calculating color, shape and texture, you would agree with me that the commercial embodiments implementing QBIC at the time would have been able to calculate color, shape and texture as part of feature extraction; correct?
MR. HANSEN: Objection; lacks foundation.
THE WITNESS: I don't remember exactly which products incorporated any subsets of the features.
BY MR. EDWARDS:
Q. Do you recall that all the products used color?
MR. HANSEN: Objection; lacks foundation.
THE WITNESS: I'm pretty sure they all did use

Page 319

color.
BY MR. EDWARDS:
Q. And color histograms specifically?
MR. HANSEN: Same objection.
(Reporter seeks clarification.)
BY MR. EDWARDS:
Q. Histograms specifically.
(Reporter seeks clarification.)
THE WITNESS: Most likely, yes.
BY MR. EDWARDS:
Q. Well, pick up Exhibit 12, please.
And go to IBM 2264.
A. Okay.
Q. So the Ultimedia Manager 1.1 was able to perform an image query using color, shape and texture; correct?
MR. HANSEN: Objection; lacks foundation.
(Witness reviews document.)
BY MR. EDWARDS:
Q. Let me orient you a little more. If you go to the bottom of IBM 2264, where it says "Image Query."
You see that?
A. Yeah.
Q. Where the bullet points start, it says,

Page 320

[as read] "Image Query allows searches by three methods: Image content, such as color, shape and texture."
You see that?
A. Yeah.
Q. So Ultimedia Manager 1.1 was able to perform an image query using color, shape and texture, correct?
MR. HANSEN: Objection; lacks foundation.
THE WITNESS: And layout.
BY MR. EDWARDS:
Q. And layout. But also color, shape and texture, correct?
A. Yes.
MR. HANSEN: Same objection.
BY MR. EDWARDS:
Q. And then if you go to the last sentence it goes on -- and the bottom of the left column on that same page, going to the top of the right column, it says, "You can drag and drop visual samples into an image query."
You see that?
A. Yeah.
Q. And that would include images, correct?
A. Correct.

Page 321

MR. HANSEN: Objection; lacks foundation.
BY MR. EDWARDS:
Q. Okay?
A. Correct.
Q. Mr. Flickner, every commercial implementation of QBIC as of the Exhibit 3, IEEE 1995 article, returned results of the matching process to the user in the form of a stored image and information associated with that image; is that fair?
MR. HANSEN: Objection; lacks foundation.
THE WITNESS: Yes.
BY MR. EDWARDS:
Q. Okay. And a thumbnail is an image; correct?
A. Thumbnail is an image.
Q. And for every commercial embodiment of QBIC as of the 1995 IEEE article in Exhibit 3, those implementations calculated a distance metric as a measure of similarity for each feature between an input image and a stored image; correct?
MR. HANSEN: Objection; lacks foundation.
THE WITNESS: Correct.
BY MR. EDWARDS:
Q. Okay. And a distance measure value

Element	'278 Claim Recitation	Exemplary Citations in Reference
		Page 322 indicates how confident the system is that the features of the stored image match the features of the input image; correct? A. Correct. Flickner 8/29/23 Depo at Ex. 6 at 1 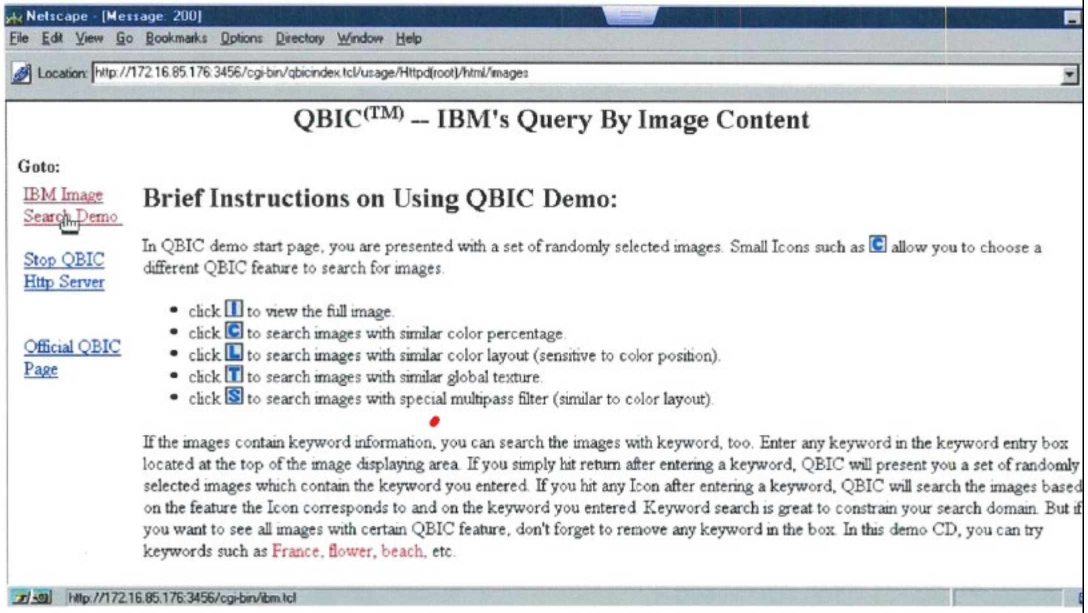 http://172.16.85.176:3456/cgi-bin/qbicindex.tcl/usage/Httpd(root)/html/images QBIC(TM) -- IBM's Query By Image Content **Brief Instructions on Using QBIC Demo:** In QBIC demo start page, you are presented with a set of randomly selected images. Small Icons such as C allow you to choose a different QBIC feature to search for images. - click I to view the full image. - click C to search images with similar color percentage. - click L to search images with similar color layout (sensitive to color position). - click T to search images with similar global texture. - click S to search images with special multipass filter (similar to color layout). If the images contain keyword information, you can search the images with keyword, too. Enter any keyword in the keyword entry box located at the top of the image displaying area. If you simply hit return after entering a keyword, QBIC will present you a set of randomly selected images which contain the keyword you entered. If you hit any Icon after entering a keyword, QBIC will search the images based on the feature the Icon corresponds to and on the keyword you entered. Keyword search is great to constrain your search domain. But if you want to see all images with certain QBIC feature, don't forget to remove any keyword in the box. In this demo CD, you can try keywords such as France, flower, beach, etc. http://172.16.85.176:3456/cgi-bin/ibm.tcl *Id.* at 2.

Element	'278 Claim Recitation	Exemplary Citations in Reference
		 Location: http://172.16.85.176:3456/cgi-bin/ibm.tcl Applet ResultApplet running Flickner 8/29/23 Depo Ex. 7 at 1

Element	'278 Claim Recitation	Exemplary Citations in Reference
		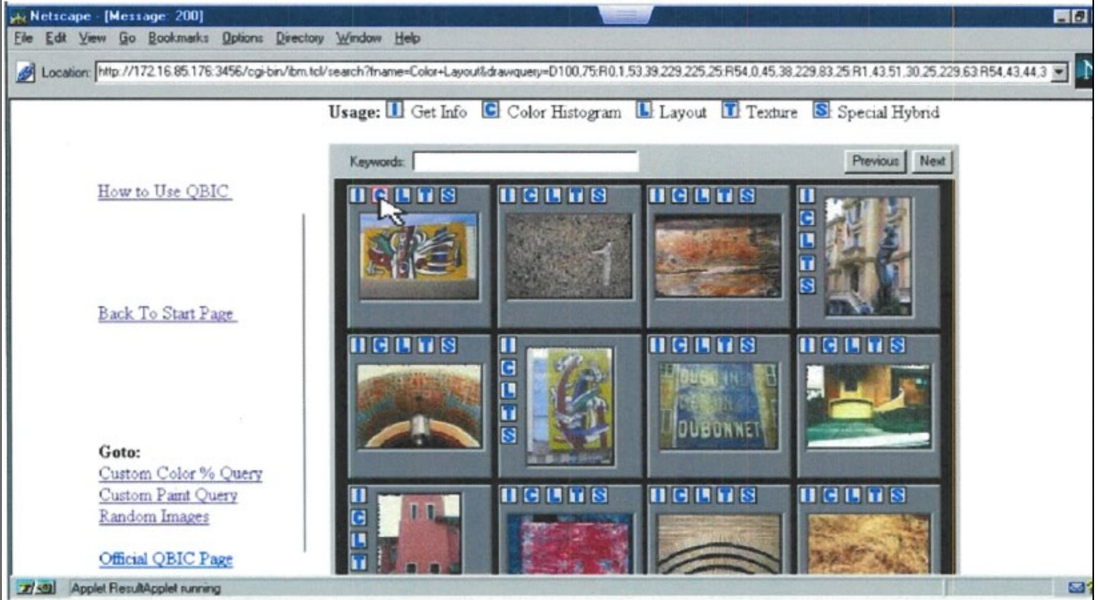 Location: http://172.16.85.176:3456/cgi-bin/ibm.tcl/search?fname=Color+Layout&drawquery=D100,75:R0,1,53,39,229,225,25:R54,0,45,38,229,83,25:R1,43,51,30,25,229,63:R54,43,44,3 *Id.* at 2

Element	'278 Claim Recitation	Exemplary Citations in Reference
		 Flickner 8/29/23 Depo at Ex. 8

Element	’278 Claim Recitation	Exemplary Citations in Reference
		 Location: http://172.16.85.176:3456/cgi-bin/ibm.tcl/search?fname=Color+Histogram&france50%2Ejpg.x=0 *Id.* at 2 IBM000001 at 39 – 41

Element	'278 Claim Recitation	Exemplary Citations in Reference
		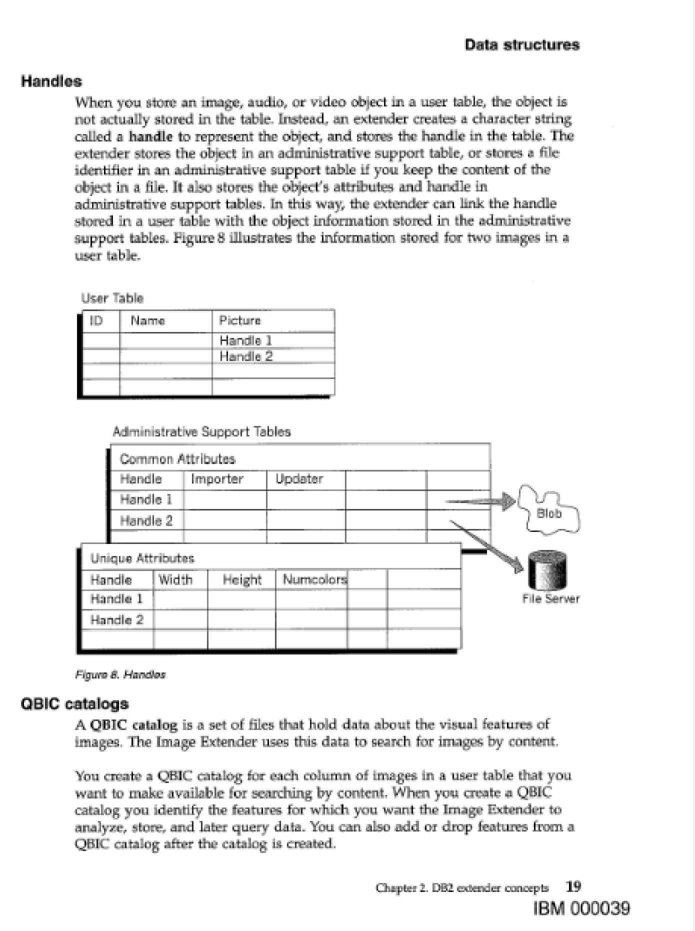 Data structures **Handles** When you store an image, audio, or video object in a user table, the object is not actually stored in the table. Instead, an extender creates a character string called a **handle** to represent the object, and stores the handle in the table. The extender stores the object in an administrative support table, or stores a file identifier in an administrative support table if you keep the content of the object in a file. It also stores the object's attributes and handle in administrative support tables. In this way, the extender can link the handle stored in a user table with the object information stored in the administrative support tables. Figure 8 illustrates the information stored for two images in a user table. *Figure 8. Handles* **QBIC catalogs** A **QBIC catalog** is a set of files that hold data about the visual features of images. The Image Extender uses this data to search for images by content. You create a QBIC catalog for each column of images in a user table that you want to make available for searching by content. When you create a QBIC catalog you identify the features for which you want the Image Extender to analyze, store, and later query data. You can also add or drop features from a QBIC catalog after the catalog is created. Chapter 2. DB2 extender concepts 19 IBM 000039

Element	'278 Claim Recitation	Exemplary Citations in Reference
		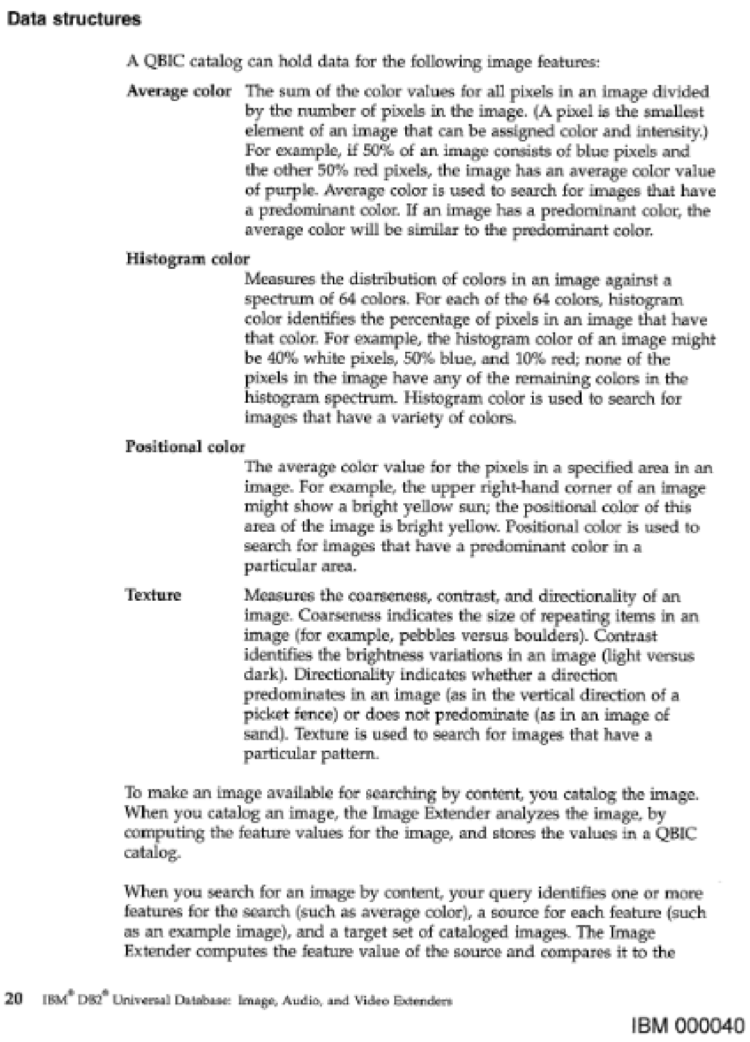 **Data structures** A QBIC catalog can hold data for the following image features: **Average color** The sum of the color values for all pixels in an image divided by the number of pixels in the image. (A pixel is the smallest element of an image that can be assigned color and intensity.) For example, if 50% of an image consists of blue pixels and the other 50% red pixels, the image has an average color value of purple. Average color is used to search for images that have a predominant color. If an image has a predominant color, the average color will be similar to the predominant color. **Histogram color** Measures the distribution of colors in an image against a spectrum of 64 colors. For each of the 64 colors, histogram color identifies the percentage of pixels in an image that have that color. For example, the histogram color of an image might be 40% white pixels, 50% blue, and 10% red; none of the pixels in the image have any of the remaining colors in the histogram spectrum. Histogram color is used to search for images that have a variety of colors. **Positional color** The average color value for the pixels in a specified area in an image. For example, the upper right-hand corner of an image might show a bright yellow sun; the positional color of this area of the image is bright yellow. Positional color is used to search for images that have a predominant color in a particular area. **Texture** Measures the coarseness, contrast, and directionality of an image. Coarseness indicates the size of repeating items in an image (for example, pebbles versus boulders). Contrast identifies the brightness variations in an image (light versus dark). Directionality indicates whether a direction predominates in an image (as in the vertical direction of a picket fence) or does not predominate (as in an image of sand). Texture is used to search for images that have a particular pattern. To make an image available for searching by content, you catalog the image. When you catalog an image, the Image Extender analyzes the image, by computing the feature values for the image, and stores the values in a QBIC catalog. When you search for an image by content, your query identifies one or more features for the search (such as average color), a source for each feature (such as an example image), and a target set of cataloged images. The Image Extender computes the feature value of the source and compares it to the 20 IBM® DB2® Universal Database: Image, Audio, and Video Extenders IBM 000040

Element	'278 Claim Recitation	Exemplary Citations in Reference
		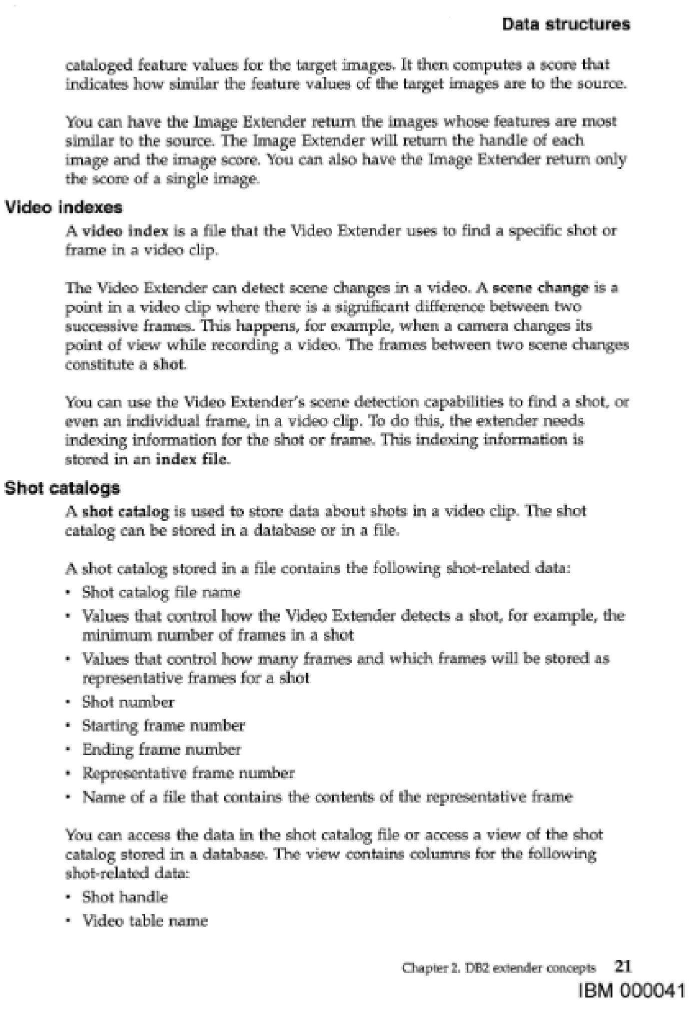 Data structures cataloged feature values for the target images. It then computes a score that indicates how similar the feature values of the target images are to the source. You can have the Image Extender return the images whose features are most similar to the source. The Image Extender will return the handle of each image and the image score. You can also have the Image Extender return only the score of a single image. **Video indexes** A **video index** is a file that the Video Extender uses to find a specific shot or frame in a video clip. The Video Extender can detect scene changes in a video. A **scene change** is a point in a video clip where there is a significant difference between two successive frames. This happens, for example, when a camera changes its point of view while recording a video. The frames between two scene changes constitute a **shot**. You can use the Video Extender's scene detection capabilities to find a shot, or even an individual frame, in a video clip. To do this, the extender needs indexing information for the shot or frame. This indexing information is stored in an **index file**. **Shot catalogs** A **shot catalog** is used to store data about shots in a video clip. The shot catalog can be stored in a database or in a file. A shot catalog stored in a file contains the following shot-related data: • Shot catalog file name • Values that control how the Video Extender detects a shot, for example, the minimum number of frames in a shot • Values that control how many frames and which frames will be stored as representative frames for a shot • Shot number • Starting frame number • Ending frame number • Representative frame number • Name of a file that contains the contents of the representative frame You can access the data in the shot catalog file or access a view of the shot catalog stored in a database. The view contains columns for the following shot-related data: • Shot handle • Video table name Chapter 2. DB2 extender concepts 21 IBM 000041 IBM0002263 at 226

Element	'278 Claim Recitation	Exemplary Citations in Reference
		 **Description** Ultimedia Manager 1.1 opens new dimensions in visual queries, helping you manage your image database more efficiently. With its easy-to-use interface, even entry-level users can perform sophisticated searches. You can search by image content, data description, or a combination of both. Once the images are found, you can view the images in full resolution, use the images in another application, or read the associated textual comments. You can use Ultimedia Manager 1.1 either as a stand-alone product or in a LAN environment. As a stand-alone product, Ultimedia Manager 1.1 employs full image classification and image query capabilities. Within the LAN environment, the Ultimedia Manager image catalogs can be sorted on a DB2 for OS/2 or DB2 for AIX database server. The image catalogs can be accessed by Windows or OS/2 clients using the Client Search Feature or the Visualizer Ultimedia Query. Ultimedia Manager runs as an OS/2 application to a DB2 for OS/2 or DB2 for AIX database server in a LAN, or as a stand-alone application with a local DB2 for OS/2 database. The Client Search Feature runs on DOS or Windows, and accesses files in dBASE flat file format. A database manager is not required to access catalogs in dBASE file format. ***Image Classification*** With the Image Classifier, you can create visual indexes based on color, shape, category, and layout elements, which are areas of the image containing the patterns you want to use for searches. You can store these elements in catalogs in DB2 for OS/2 databases. The catalogs can be used locally or, through the use of utility programs, be transferred to DB2 for OS/2 or AIX servers. In addition, you can convert the catalog data into dBASE flat file format for access by Windows client applications. The Image Classifier and Analyzer allows automatic or interactive classification of images for later image queries, flexibility in storage of catalog information, and access to multiple users. You can organize large image catalogs so that later searches for images can be performed easily. You can classify automatically using a batch program, or interactively by specifying areas of interest in an image and indicating the visual cue (color, shape, or texture) to be saved for later queries. ***Image Query*** Once the image catalogs are created, you can use the Image Query component to find image and business information. Business information, such as character and numeric data, are stored in the database tables using normal table population methods. Image Query is provided in Ultimedia Manager and Client Search Feature. Image Query allows searches by three methods: • Image content, such as color, shape, texture, or layout • Data description, such as business data stored in the database (character or numeric) or descriptive text (category) • Combinations of image content and data description. ***Image Content:*** To search by image content, specify the attributes (color, shape, texture, or layout). You can drag and drop visual samples into an image query, select a color from the color wheel, or sketch a shape. ***Data Description:*** To search by data description, enter character and numeric data values in the database tables using the simple fill-in-the-blanks query panel. You can enter data such as name, number, date, or a range of dates. You can integrate character or numeric data as part of the image query, enabling a more effective search for better results. The simple fill-in-the-blanks method for specifying desired conditions is so easy even non-programmers can use database table information as part of their searches. ***Combinations:*** Image queries using either image content or data description, or combinations of both, allow you to search through thousands of images more quickly and effectively than the traditional text and keyword search methods. ***Image Formats*** Ultimedia Manager 1.1 understands these image formats: • Audio Visual Connection (AVC) images • OS/2 bitmap • Windows 3.x bitmap • PCX • Intel and Motorola TIFF • TARGA TGA • TIFF (FAX) • GIF • Kodak Photo-CD **Product Positioning** The advanced image search technology in Ultimedia Manager, query by image content, was developed in conjunction with the IBM Almaden Research Center. It provides function not available with conventional, text-only image management systems. Ultimedia Manager fulfills image search and retrieval needs of business users and creative professionals through use of online image libraries or collections. It is intended for users who have large image libraries and need fast, flexible access to their image data, users who may also want to add character or numeric data to their image queries, and users who feel limited by text-based systems. Ultimedia Manager 1.1 works with the DB2 Family of products and the Visualizer Ultimedia Query product. Catalogs created by Ultimedia Manager can be stored in DB2 for OS/2 database tables. Image catalogs also can be stored in DB2 for OS/2 or DB2 for AIX database servers in client/server environments. OS/2 and Windows clients in a LAN environment using the Client Search Feature and the Visualizer Ultimedia Query can access catalogs created with Ultimedia Manager 1.1. Visualizer Ultimedia Query provides users with the capability to: • Add multimedia columns (such as image, video, audio, and formatted text) to DB2 for OS/2 database tables • Query and display images • View or listen to video and audio information • Display formatted text Users of Ultimedia Manager 1.1 and Client Search Feature are focused on image management and query applications. Visualizer Ultimedia Query users are 295-042 -2- IBM 0002264

Element	'278 Claim Recitation	Exemplary Citations in Reference
		IBM000777 at 794 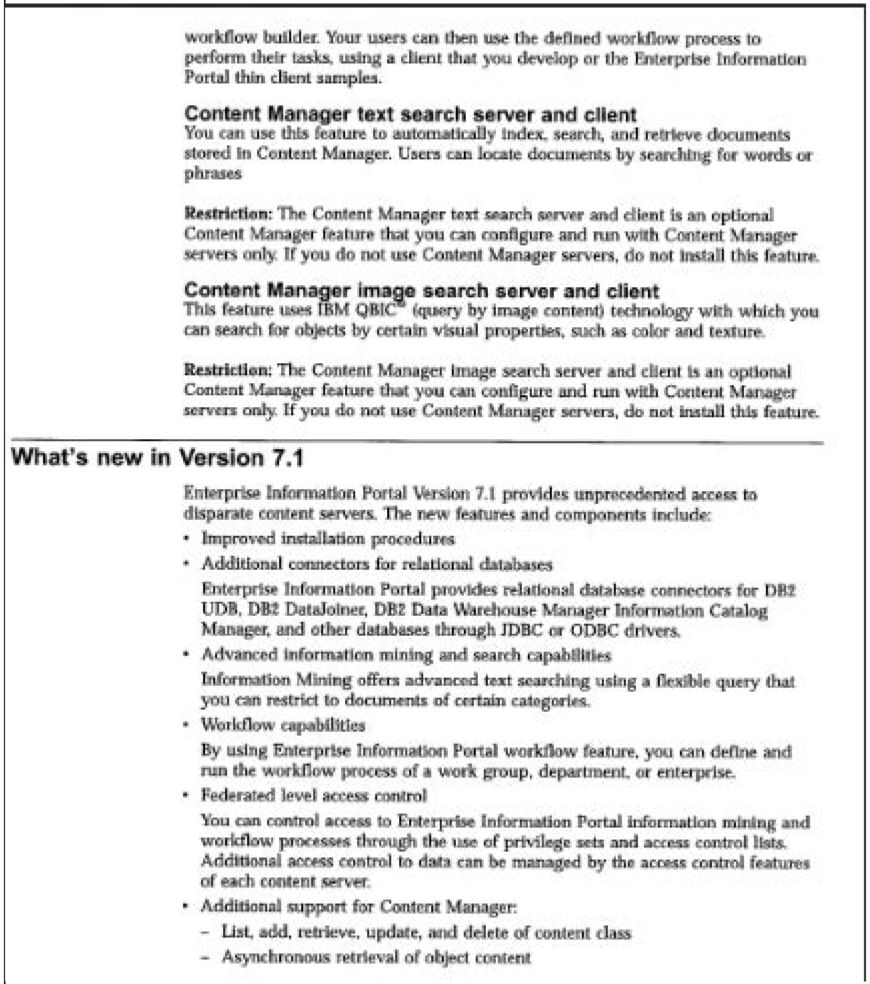 workflow builder. Your users can then use the defined workflow process to perform their tasks, using a client that you develop or the Enterprise Information Portal thin client samples. **Content Manager text search server and client** You can use this feature to automatically index, search, and retrieve documents stored in Content Manager. Users can locate documents by searching for words or phrases **Restriction:** The Content Manager text search server and client is an optional Content Manager feature that you can configure and run with Content Manager servers only. If you do not use Content Manager servers, do not install this feature. **Content Manager image search server and client** This feature uses IBM QBIC® (query by image content) technology with which you can search for objects by certain visual properties, such as color and texture. **Restriction:** The Content Manager image search server and client is an optional Content Manager feature that you can configure and run with Content Manager servers only. If you do not use Content Manager servers, do not install this feature. **What's new in Version 7.1** Enterprise Information Portal Version 7.1 provides unprecedented access to disparate content servers. The new features and components include: • Improved installation procedures • Additional connectors for relational databases Enterprise Information Portal provides relational database connectors for DB2 UDB, DB2 DataJoiner, DB2 Data Warehouse Manager Information Catalog Manager, and other databases through JDBC or ODBC drivers. • Advanced information mining and search capabilities Information Mining offers advanced text searching using a flexible query that you can restrict to documents of certain categories. • Workflow capabilities By using Enterprise Information Portal workflow feature, you can define and run the workflow process of a work group, department, or enterprise. • Federated level access control You can control access to Enterprise Information Portal information mining and workflow processes through the use of privilege sets and access control lists. Additional access control to data can be managed by the access control features of each content server. • Additional support for Content Manager: – List, add, retrieve, update, and delete of content class – Asynchronous retrieval of object content *See also* Flickner Depo at Flickner Depo at 22-25, 25-26, 36, 54-55, 88, 89-90

Element	'278 Claim Recitation	Exemplary Citations in Reference
		To the extent it is found that QBIC does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that QBIC does not anticipate this claim, QBIC renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine QBIC with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1E.i]	the image,	QBIC discloses this limitation. *See* limitation 1E above. To the extent it is found that QBIC does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that QBIC does not anticipate this claim, QBIC renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine QBIC with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

Element	'278 Claim Recitation	Exemplary Citations in Reference
[1E.ii]	the symbol information,	QBIC discloses this limitation. *See* limitations 1D and 1E above. Flickner Depo at 185, 217 To the extent it is found that QBIC does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that QBIC does not anticipate this claim, QBIC renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine QBIC with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1E.iii]	and a query of a database storing target object information associated with a plurality of target objects including the first target object;	QBIC discloses this limitation. *See* limitation 1E above.

Element	'278 Claim Recitation	Exemplary Citations in Reference
		QBIC Demo (IBM000747)  QBIC Demo (IBM000747) 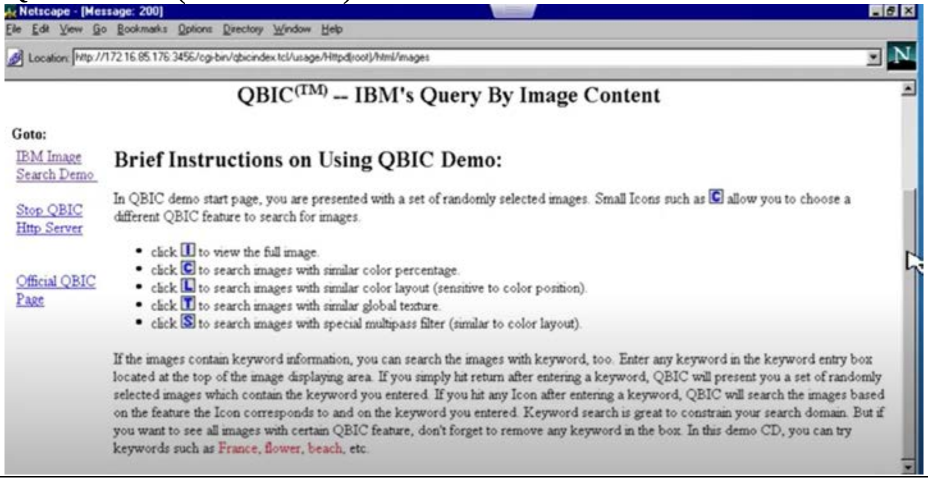 Netscape - [Message: 200] File Edit View Go Bookmarks Options Directory Window Help Location: http://172.16.85.176:3456/cgi-bin/qbicindex.tcl/usage/Httpd(root)/html/images QBIC(TM) -- IBM's Query By Image Content Goto: IBM Image Search Demo Stop QBIC Http Server Official QBIC Page **Brief Instructions on Using QBIC Demo:** In QBIC demo start page, you are presented with a set of randomly selected images. Small Icons such as C allow you to choose a different QBIC feature to search for images. - click I to view the full image. - click C to search images with similar color percentage. - click L to search images with similar color layout (sensitive to color position). - click T to search images with similar global texture. - click S to search images with special multipass filter (similar to color layout). If the images contain keyword information, you can search the images with keyword, too. Enter any keyword in the keyword entry box located at the top of the image displaying area. If you simply hit return after entering a keyword, QBIC will present you a set of randomly selected images which contain the keyword you entered. If you hit any Icon after entering a keyword, QBIC will search the images based on the feature the Icon corresponds to and on the keyword you entered. Keyword search is great to constrain your search domain. But if you want to see all images with certain QBIC feature, don't forget to remove any keyword in the box. In this demo CD, you can try keywords such as France, flower, beach, etc.

Element	'278 Claim Recitation	Exemplary Citations in Reference
		Flickner Depo at 20 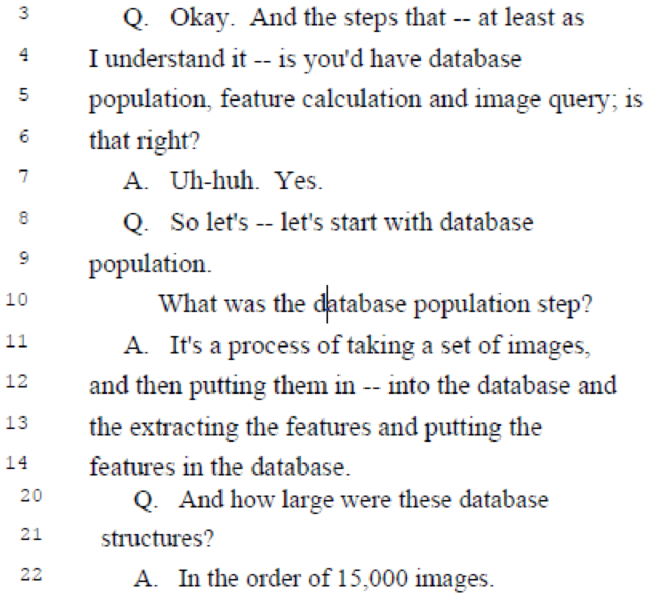 Q. Okay. And the steps that -- at least as I understand it -- is you'd have database population, feature calculation and image query; is that right? A. Uh-huh. Yes. Q. So let's -- let's start with database population. What was the database population step? A. It's a process of taking a set of images, and then putting them in -- into the database and the extracting the features and putting the features in the database. Q. And how large were these database structures? A. In the order of 15,000 images. Flickner Depo at 25-26 Flickner Depo at 18

Element	'278 Claim Recitation	Exemplary Citations in Reference
		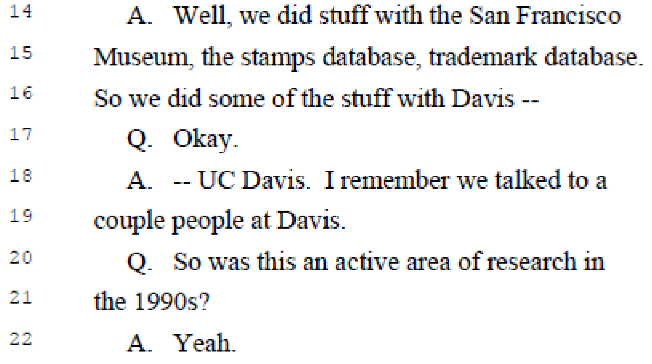 A. Well, we did stuff with the San Francisco Museum, the stamps database, trademark database. So we did some of the stuff with Davis -- Q. Okay. A. -- UC Davis. I remember we talked to a couple people at Davis. Q. So was this an active area of research in the 1990s? A. Yeah. Flickner Depo at 88-89

Element	'278 Claim Recitation	Exemplary Citations in Reference
		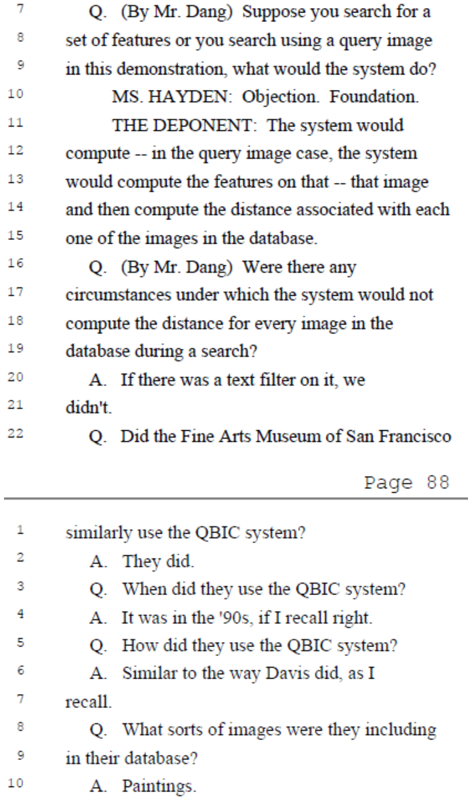 Q. (By Mr. Dang) Suppose you search for a set of features or you search using a query image in this demonstration, what would the system do? MS. HAYDEN: Objection. Foundation. THE DEPONENT: The system would compute -- in the query image case, the system would compute the features on that -- that image and then compute the distance associated with each one of the images in the database. Q. (By Mr. Dang) Were there any circumstances under which the system would not compute the distance for every image in the database during a search? A. If there was a text filter on it, we didn't. Q. Did the Fine Arts Museum of San Francisco Page 88 similarly use the QBIC system? A. They did. Q. When did they use the QBIC system? A. It was in the '90s, if I recall right. Q. How did they use the QBIC system? A. Similar to the way Davis did, as I recall. Q. What sorts of images were they including in their database? A. Paintings. Flickner Depo at 142

Element	'278 Claim Recitation	Exemplary Citations in Reference
		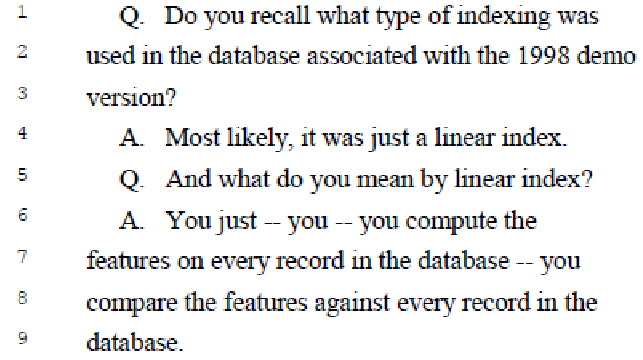 Q. Do you recall what type of indexing was used in the database associated with the 1998 demo version? A. Most likely, it was just a linear index. Q. And what do you mean by linear index? A. You just -- you -- you compute the features on every record in the database -- you compare the features against every record in the database. Flickner Depo at 168-169 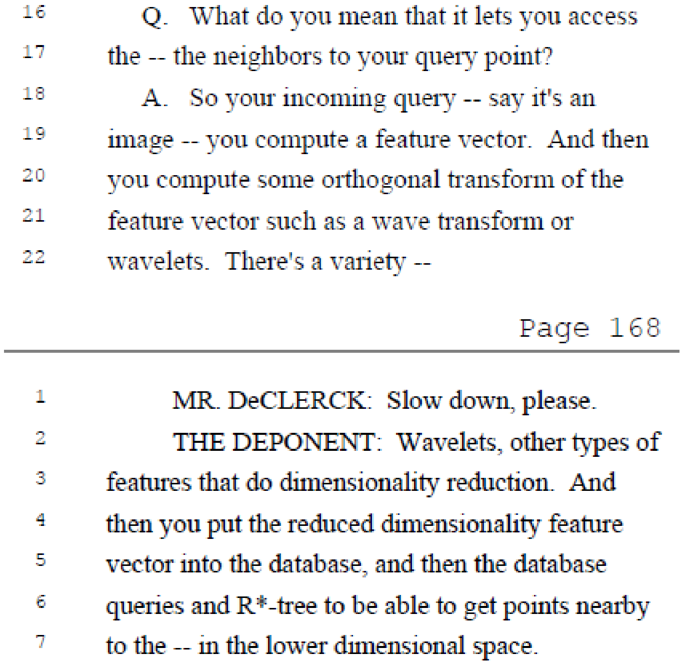 Q. What do you mean that it lets you access the -- the neighbors to your query point? A. So your incoming query -- say it's an image -- you compute a feature vector. And then you compute some orthogonal transform of the feature vector such as a wave transform or wavelets. There's a variety -- Page 168 MR. DeCLERCK: Slow down, please. THE DEPONENT: Wavelets, other types of features that do dimensionality reduction. And then you put the reduced dimensionality feature vector into the database, and then the database queries and R*-tree to be able to get points nearby to the -- in the lower dimensional space.

Element	'278 Claim Recitation	Exemplary Citations in Reference
		Flickner Depo at 181 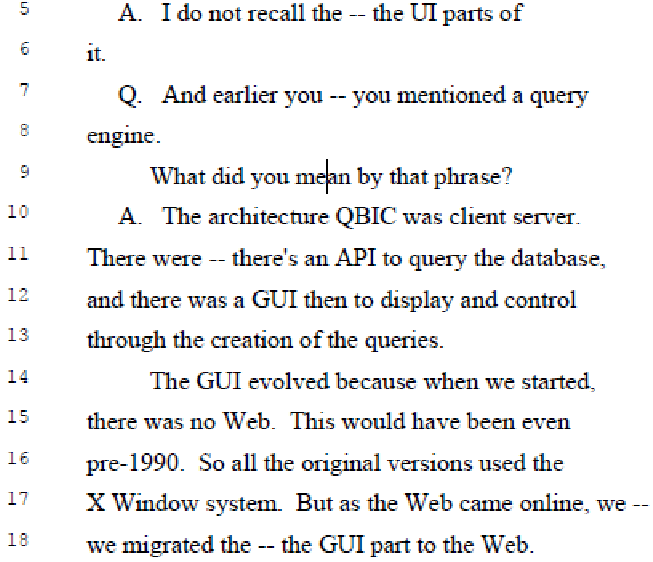 A. I do not recall the -- the UI parts of it. Q. And earlier you -- you mentioned a query engine. What did you mean by that phrase? A. The architecture QBIC was client server. There were -- there's an API to query the database, and there was a GUI then to display and control through the creation of the queries. The GUI evolved because when we started, there was no Web. This would have been even pre-1990. So all the original versions used the X Window system. But as the Web came online, we -- we migrated the -- the GUI part to the Web. Flickner Depo at 91-92 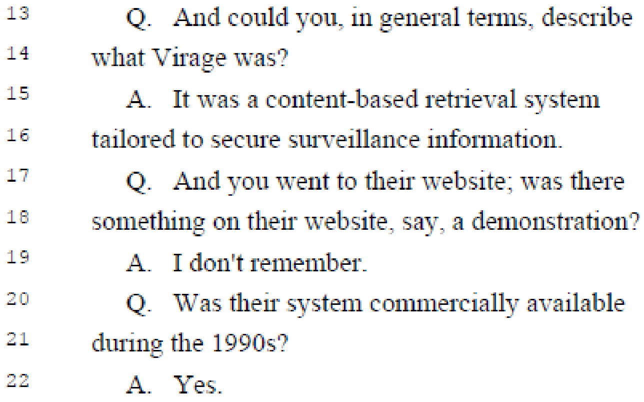 Q. And could you, in general terms, describe what Virage was? A. It was a content-based retrieval system tailored to secure surveillance information. Q. And you went to their website; was there something on their website, say, a demonstration? A. I don't remember. Q. Was their system commercially available during the 1990s? A. Yes.

Element	'278 Claim Recitation	Exemplary Citations in Reference
		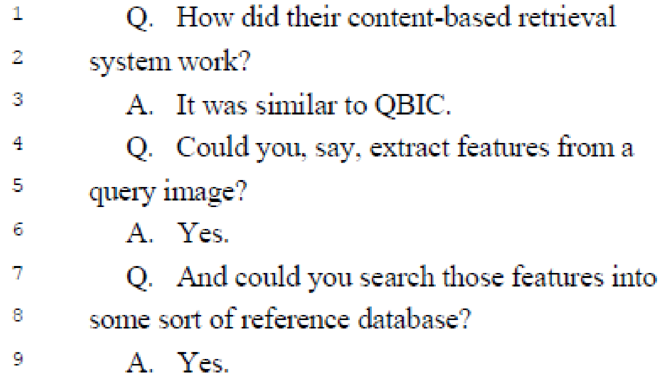 Q. How did their content-based retrieval system work? A. It was similar to QBIC. Q. Could you, say, extract features from a query image? A. Yes. Q. And could you search those features into some sort of reference database? A. Yes. Flickner Depo at 48-49 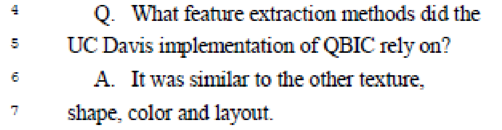 Q. What feature extraction methods did the UC Davis implementation of QBIC rely on? A. It was similar to the other texture, shape, color and layout. Flickner Depo at 209-210 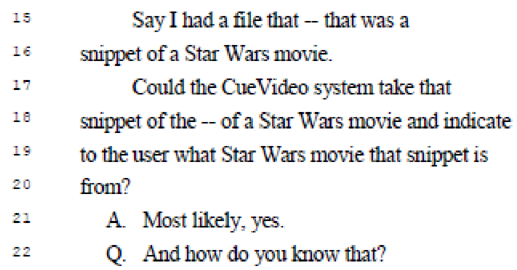 Q. Okay. Let me ask -- let me ask that question using -- using an example. Say I had a file that -- that was a snippet of a Star Wars movie. Could the CueVideo system take that snippet of the -- of a Star Wars movie and indicate to the user what Star Wars movie that snippet is from? A. Most likely, yes. Q. And how do you know that? Flickner 8/29/23 Depo at 53 – 56

Element	'278 Claim Recitation	Exemplary Citations in Reference
		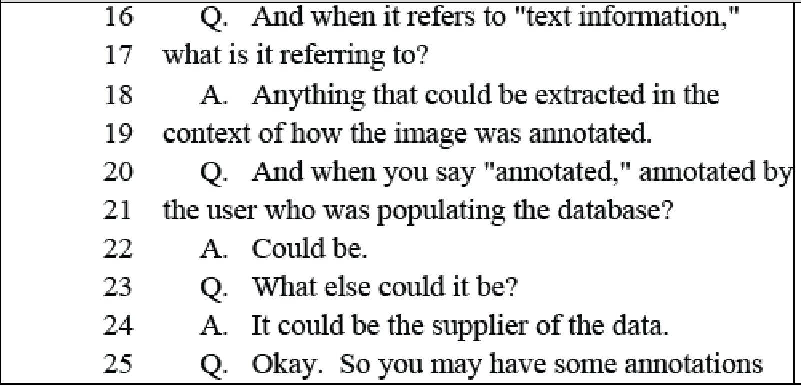 Q. And when it refers to "text information," what is it referring to? A. Anything that could be extracted in the context of how the image was annotated. Q. And when you say "annotated," annotated by the user who was populating the database? A. Could be. Q. What else could it be? A. It could be the supplier of the data. Q. Okay. So you may have some annotations

Element	’278 Claim Recitation	Exemplary Citations in Reference
		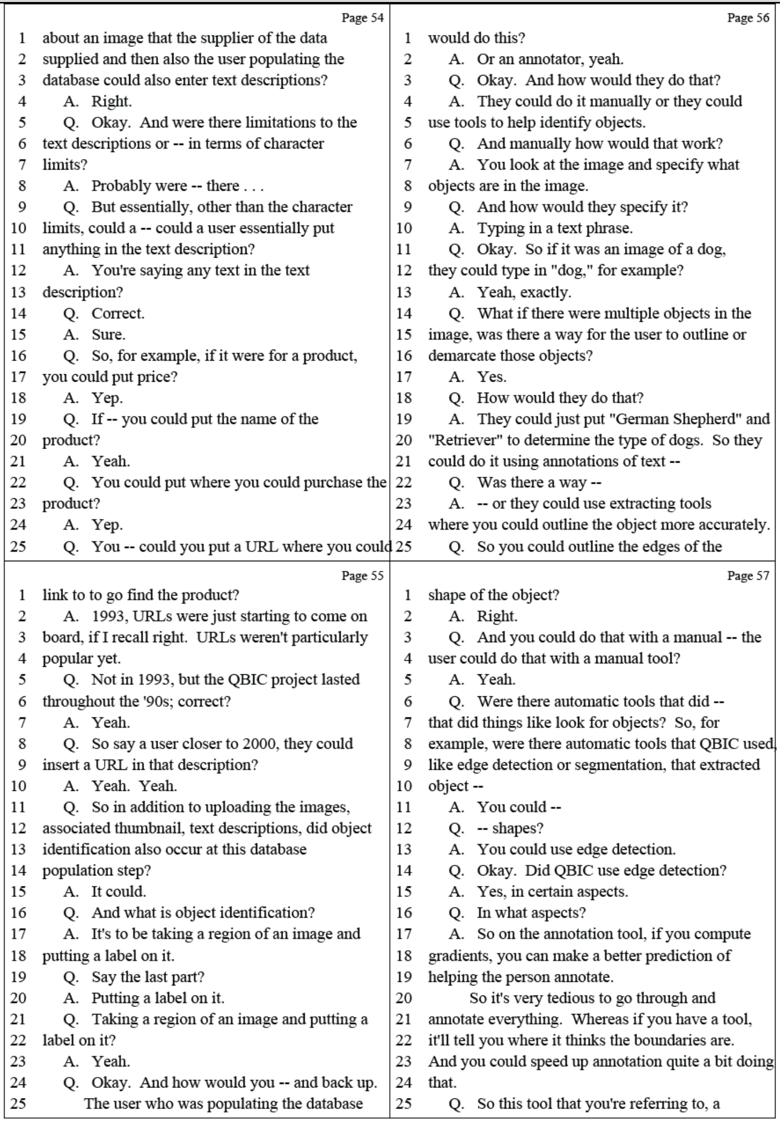 Page 54 about an image that the supplier of the data supplied and then also the user populating the database could also enter text descriptions? A. Right. Q. Okay. And were there limitations to the text descriptions or -- in terms of character limits? A. Probably were -- there . . . Q. But essentially, other than the character limits, could a -- could a user essentially put anything in the text description? A. You're saying any text in the text description? Q. Correct. A. Sure. Q. So, for example, if it were for a product, you could put price? A. Yep. Q. If -- you could put the name of the product? A. Yeah. Q. You could put where you could purchase the product? A. Yep. Q. You -- could you put a URL where you could Page 55 link to to go find the product? A. 1993, URLs were just starting to come on board, if I recall right. URLs weren't particularly popular yet. Q. Not in 1993, but the QBIC project lasted throughout the '90s; correct? A. Yeah. Q. So say a user closer to 2000, they could insert a URL in that description? A. Yeah. Yeah. Q. So in addition to uploading the images, associated thumbnail, text descriptions, did object identification also occur at this database population step? A. It could. Q. And what is object identification? A. It's to be taking a region of an image and putting a label on it. Q. Say the last part? A. Putting a label on it. Q. Taking a region of an image and putting a label on it? A. Yeah. Q. Okay. And how would you -- and back up. The user who was populating the database Page 56 would do this? A. Or an annotator, yeah. Q. Okay. And how would they do that? A. They could do it manually or they could use tools to help identify objects. Q. And manually how would that work? A. You look at the image and specify what objects are in the image. Q. And how would they specify it? A. Typing in a text phrase. Q. Okay. So if it was an image of a dog, they could type in "dog," for example? A. Yeah, exactly. Q. What if there were multiple objects in the image, was there a way for the user to outline or demarcate those objects? A. Yes. Q. How would they do that? A. They could just put "German Shepherd" and "Retriever" to determine the type of dogs. So they could do it using annotations of text -- Q. Was there a way -- A. -- or they could use extracting tools where you could outline the object more accurately. Q. So you could outline the edges of the Page 57 shape of the object? A. Right. Q. And you could do that with a manual -- the user could do that with a manual tool? A. Yeah. Q. Were there automatic tools that did -- that did things like look for objects? So, for example, were there automatic tools that QBIC used, like edge detection or segmentation, that extracted object -- A. You could -- Q. -- shapes? A. You could use edge detection. Q. Okay. Did QBIC use edge detection? A. Yes, in certain aspects. Q. In what aspects? A. So on the annotation tool, if you compute gradients, you can make a better prediction of helping the person annotate. So it's very tedious to go through and annotate everything. Whereas if you have a tool, it'll tell you where it thinks the boundaries are. And you could speed up annotation quite a bit doing that. Q. So this tool that you're referring to, a

Element	'278 Claim Recitation	Exemplary Citations in Reference
		Id. at 60 – 62 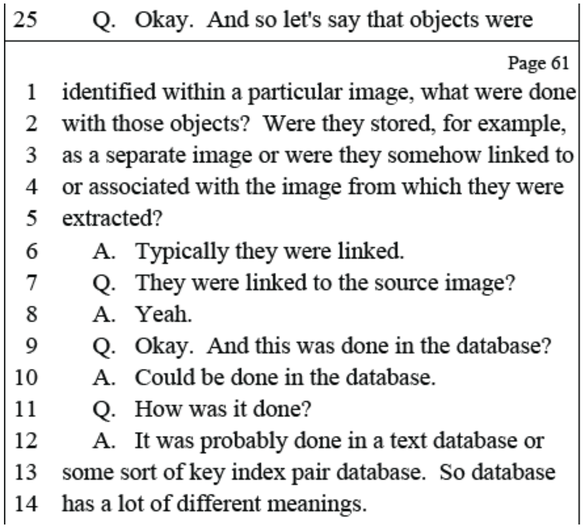 Q. Okay. And so let's say that objects were Page 61 identified within a particular image, what were done with those objects? Were they stored, for example, as a separate image or were they somehow linked to or associated with the image from which they were extracted? A. Typically they were linked. Q. They were linked to the source image? A. Yeah. Q. Okay. And this was done in the database? A. Could be done in the database. Q. How was it done? A. It was probably done in a text database or some sort of key index pair database. So database has a lot of different meanings. 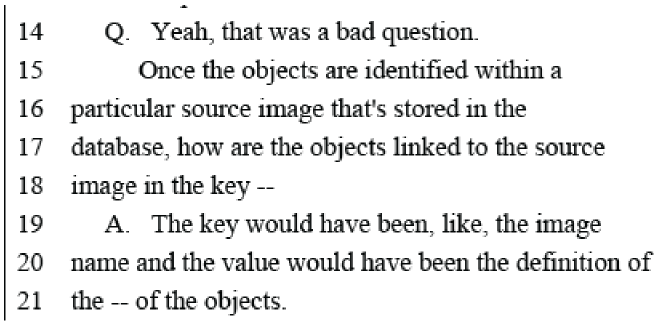 Q. Yeah, that was a bad question. Once the objects are identified within a particular source image that's stored in the database, how are the objects linked to the source image in the key -- A. The key would have been, like, the image name and the value would have been the definition of the -- of the objects. *Id.* at 103

Element	'278 Claim Recitation	Exemplary Citations in Reference
		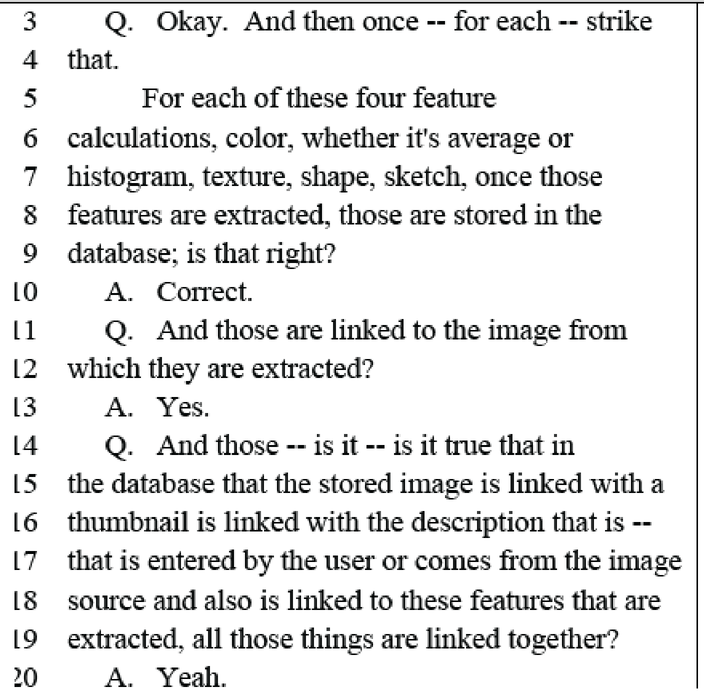 Q. Okay. And then once -- for each -- strike that. For each of these four feature calculations, color, whether it's average or histogram, texture, shape, sketch, once those features are extracted, those are stored in the database; is that right? A. Correct. Q. And those are linked to the image from which they are extracted? A. Yes. Q. And those -- is it -- is it true that in the database that the stored image is linked with a thumbnail is linked with the description that is -- that is entered by the user or comes from the image source and also is linked to these features that are extracted, all those things are linked together? A. Yeah. *Id.* at 129 – 131

Element	'278 Claim Recitation	Exemplary Citations in Reference
		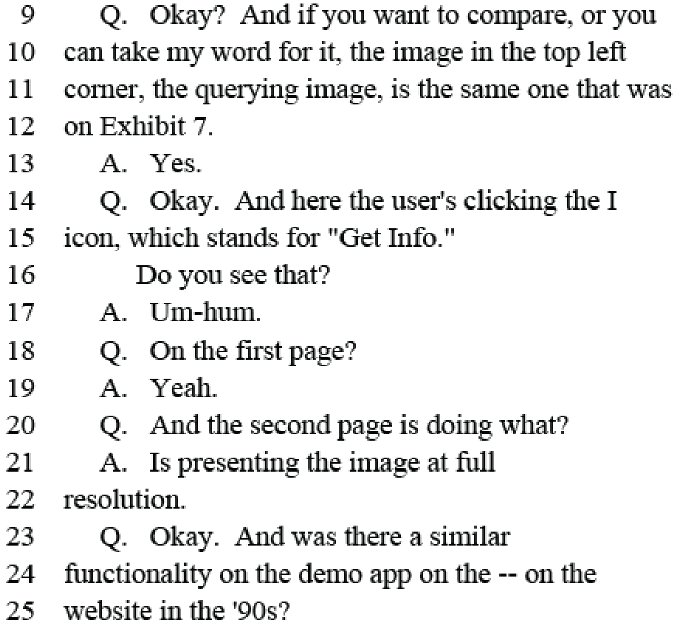 Q. Okay? And if you want to compare, or you can take my word for it, the image in the top left corner, the querying image, is the same one that was on Exhibit 7. A. Yes. Q. Okay. And here the user's clicking the I icon, which stands for "Get Info." Do you see that? A. Um-hum. Q. On the first page? A. Yeah. Q. And the second page is doing what? A. Is presenting the image at full resolution. Q. Okay. And was there a similar functionality on the demo app on the -- on the website in the '90s?

		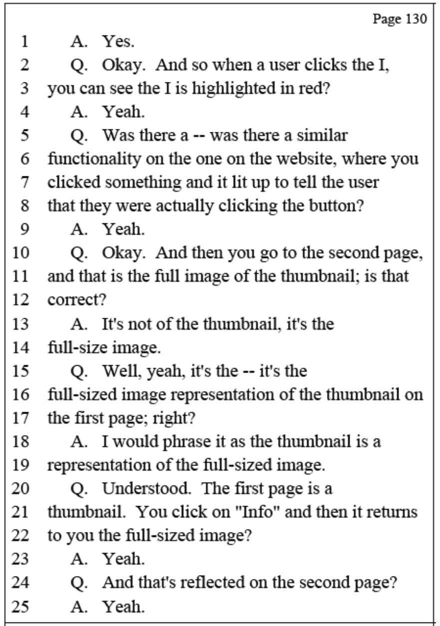 Page 130 A. Yes. Q. Okay. And so when a user clicks the I, you can see the I is highlighted in red? A. Yeah. Q. Was there a -- was there a similar functionality on the one on the website, where you clicked something and it lit up to tell the user that they were actually clicking the button? A. Yeah. Q. Okay. And then you go to the second page, and that is the full image of the thumbnail; is that correct? A. It's not of the thumbnail, it's the full-size image. Q. Well, yeah, it's the -- it's the full-sized image representation of the thumbnail on the first page; right? A. I would phrase it as the thumbnail is a representation of the full-sized image. Q. Understood. The first page is a thumbnail. You click on "Info" and then it returns to you the full-sized image? A. Yeah. Q. And that's reflected on the second page? A. Yeah. 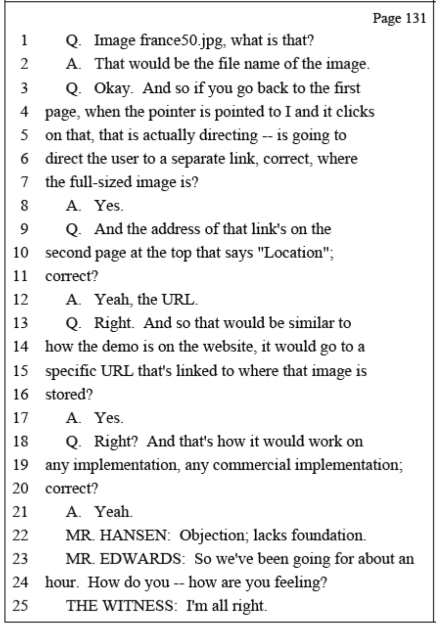 Page 131 Q. Image france50.jpg, what is that? A. That would be the file name of the image. Q. Okay. And so if you go back to the first page, when the pointer is pointed to I and it clicks on that, that is actually directing -- is going to direct the user to a separate link, correct, where the full-sized image is? A. Yes. Q. And the address of that link's on the second page at the top that says "Location"; correct? A. Yeah, the URL. Q. Right. And so that would be similar to how the demo is on the website, it would go to a specific URL that's linked to where that image is stored? A. Yes. Q. Right? And that's how it would work on any implementation, any commercial implementation; correct? A. Yeah. MR. HANSEN: Objection; lacks foundation. MR. EDWARDS: So we've been going for about an hour. How do you -- how are you feeling? THE WITNESS: I'm all right.

Element	'278 Claim Recitation	Exemplary Citations in Reference
		Id. at 161 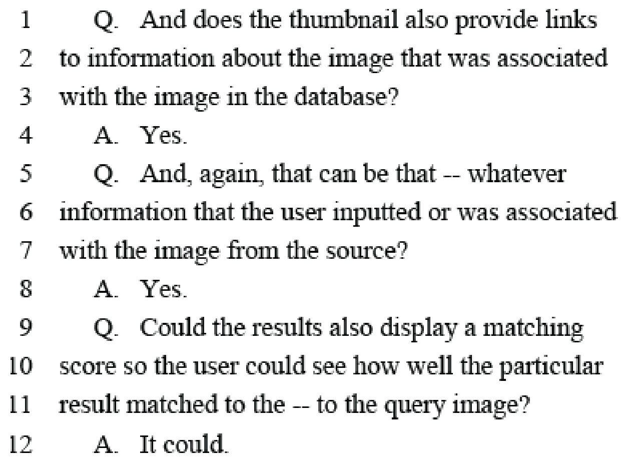 Q. And does the thumbnail also provide links to information about the image that was associated with the image in the database? A. Yes. Q. And, again, that can be that -- whatever information that the user inputted or was associated with the image from the source? A. Yes. Q. Could the results also display a matching score so the user could see how well the particular result matched to the -- to the query image? A. It could. *Id.* at 321 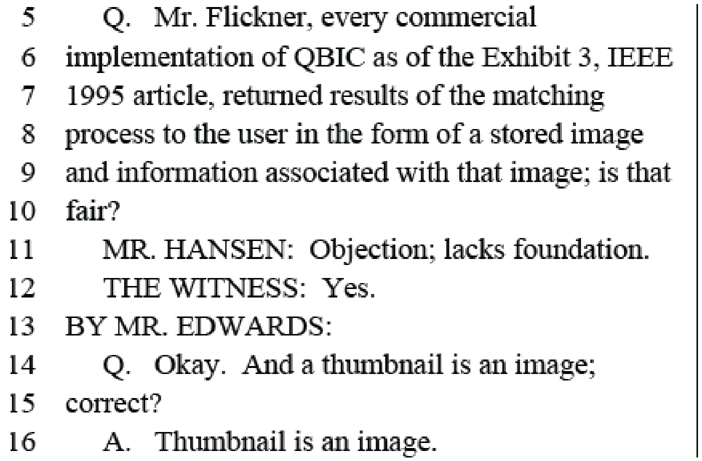 Q. Mr. Flickner, every commercial implementation of QBIC as of the Exhibit 3, IEEE 1995 article, returned results of the matching process to the user in the form of a stored image and information associated with that image; is that fair? MR. HANSEN: Objection; lacks foundation. THE WITNESS: Yes. BY MR. EDWARDS: Q. Okay. And a thumbnail is an image; correct? A. Thumbnail is an image. Flickner 8/29/23 Depo at Ex. 6 at 1

Element	'278 Claim Recitation	Exemplary Citations in Reference
		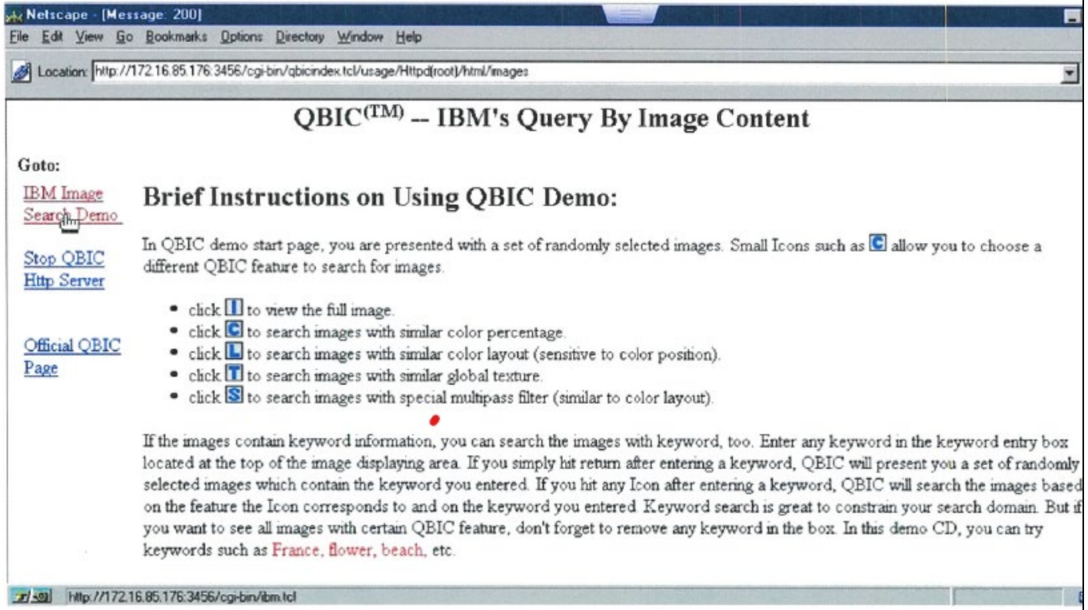 QBIC$^{(TM)}$ -- IBM's Query By Image Content Goto: IBM Image Search Demo Stop QBIC Http Server Official QBIC Page **Brief Instructions on Using QBIC Demo:** In QBIC demo start page, you are presented with a set of randomly selected images. Small Icons such as C allow you to choose a different QBIC feature to search for images. - click I to view the full image. - click C to search images with similar color percentage. - click L to search images with similar color layout (sensitive to color position). - click T to search images with similar global texture. - click S to search images with special multipass filter (similar to color layout). If the images contain keyword information, you can search the images with keyword, too. Enter any keyword in the keyword entry box located at the top of the image displaying area. If you simply hit return after entering a keyword, QBIC will present you a set of randomly selected images which contain the keyword you entered. If you hit any Icon after entering a keyword, QBIC will search the images based on the feature the Icon corresponds to and on the keyword you entered. Keyword search is great to constrain your search domain. But if you want to see all images with certain QBIC feature, don't forget to remove any keyword in the box. In this demo CD, you can try keywords such as France, flower, beach, etc. *Id.* at 2.

Element	'278 Claim Recitation	Exemplary Citations in Reference
		 Flickner 8/29/23 Depo Ex. 7 at 1

Element	**'278 Claim Recitation**	**Exemplary Citations in Reference**
		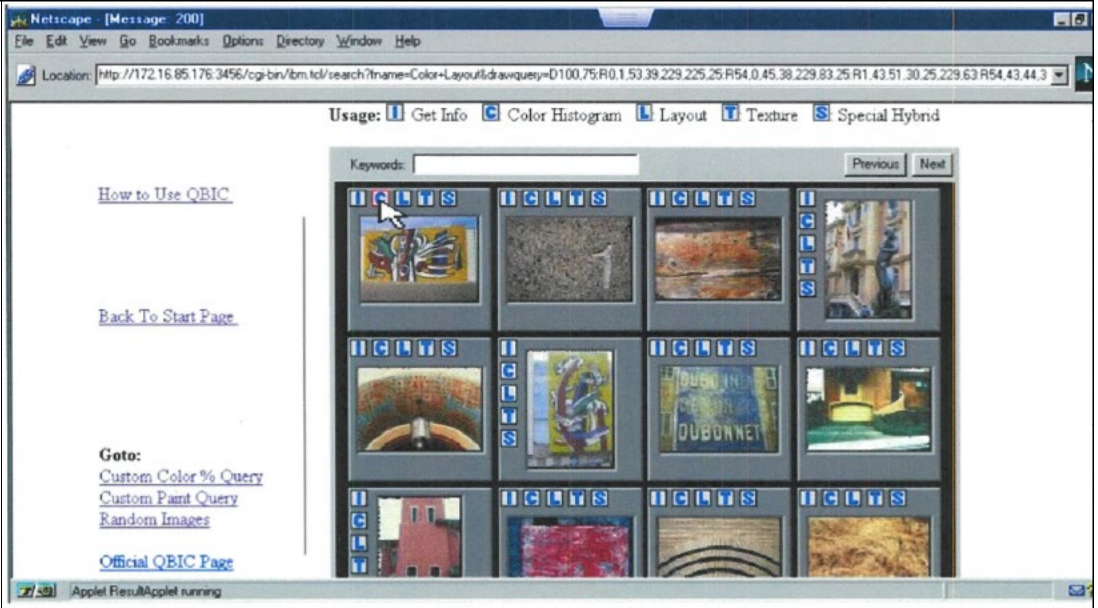 *Id.* at 2

Element	'278 Claim Recitation	Exemplary Citations in Reference
		 Location: http://172.16.85.176:3456/cgi-bin/ibm.tcl/search?fname=Color+Histogram&france50%2Ejpg x=0 Flickner 8/29/23 Depo at Ex. 8

Element	'278 Claim Recitation	Exemplary Citations in Reference
		 Location: http://172.16.85.176:3456/cgi-bin/ibm.tcl/search?fname=Color+Histogram&france50%2Ejpg.x=0 *Id.* at 2 *See also* Flickner Depo at 22-25, 25-26, 36, 39-40, 54-55, 75-76, 78-79, 159, To the extent it is found that QBIC does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that QBIC does not anticipate this claim, QBIC renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine QBIC with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references

Element	'278 Claim Recitation	Exemplary Citations in Reference
		contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1F]	receiving, via an address, first target object information associated with the first target object,	QBIC discloses receiving, via an address, first target object information associated with the first target object. The QBIC Project: Querying Images by Content Using Color, Texture, and Shape (IBM 0002390–404) at 2391 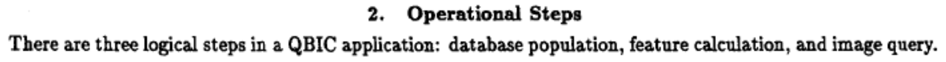 **2. Operational Steps** There are three logical steps in a QBIC application: database population, feature calculation, and image query. IBM0002390 at 2391-92 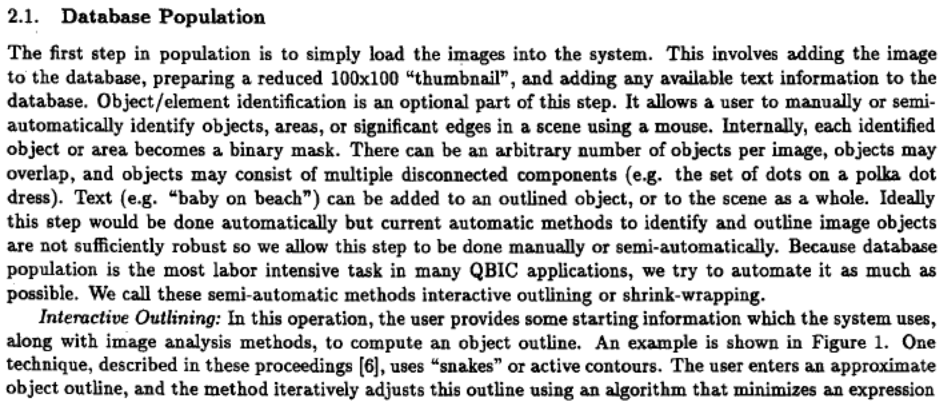 **2.1. Database Population** The first step in population is to simply load the images into the system. This involves adding the image to the database, preparing a reduced 100x100 "thumbnail", and adding any available text information to the database. Object/element identification is an optional part of this step. It allows a user to manually or semi-automatically identify objects, areas, or significant edges in a scene using a mouse. Internally, each identified object or area becomes a binary mask. There can be an arbitrary number of objects per image, objects may overlap, and objects may consist of multiple disconnected components (e.g. the set of dots on a polka dot dress). Text (e.g. "baby on beach") can be added to an outlined object, or to the scene as a whole. Ideally this step would be done automatically but current automatic methods to identify and outline image objects are not sufficiently robust so we allow this step to be done manually or semi-automatically. Because database population is the most labor intensive task in many QBIC applications, we try to automate it as much as possible. We call these semi-automatic methods interactive outlining or shrink-wrapping. *Interactive Outlining:* In this operation, the user provides some starting information which the system uses, along with image analysis methods, to compute an object outline. An example is shown in Figure 1. One technique, described in these proceedings [6], uses "snakes" or active contours. The user enters an approximate object outline, and the method iteratively adjusts this outline using an algorithm that minimizes an expression

Element	'278 Claim Recitation	Exemplary Citations in Reference
		containing two main terms: total image edge strength along the outline, and the outline curvature. Another method, flood fill, can take as starting information a single mouse-clicked image point. Starting from this object point, all adjacent pixels are included in the object whose color is within a specified color distance from the original point. The perimeter of all included points becomes the object boundary. For reasonably uniform objects that are distinct from the background, this operation provides a fast object outliner. In interactive outlining, as in other parts of our system, we have attempted to build the functions in such a way as to carefully distribute the tasks between the human and machine. For example, the tasks of giving the starting outline for an object or of attaching semantic meaning to objects is left to the user since they are difficult for machines. On the other hand, computation of quantitative image/object features is ideally suited for machines and so is done automatically.  Figure 1: Original image, approximate user outline, and automatically refined outline of sun. IBM0002390 at 2392-93:

Element	'278 Claim Recitation	Exemplary Citations in Reference
		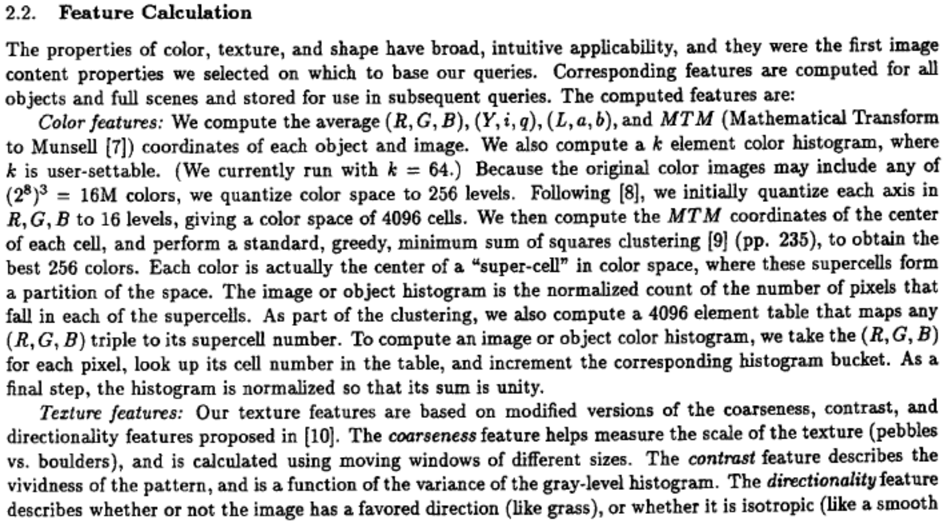 2.2. **Feature Calculation** The properties of color, texture, and shape have broad, intuitive applicability, and they were the first image content properties we selected on which to base our queries. Corresponding features are computed for all objects and full scenes and stored for use in subsequent queries. The computed features are: *Color features:* We compute the average (R,G,B), (Y,i,q), (L,a,b), and MTM (Mathematical Transform to Munsell [7]) coordinates of each object and image. We also compute a k element color histogram, where k is user-settable. (We currently run with $k = 64$.) Because the original color images may include any of $(2^8)^3 = 16$M colors, we quantize color space to 256 levels. Following [8], we initially quantize each axis in R, G, B to 16 levels, giving a color space of 4096 cells. We then compute the MTM coordinates of the center of each cell, and perform a standard, greedy, minimum sum of squares clustering [9] (pp. 235), to obtain the best 256 colors. Each color is actually the center of a "super-cell" in color space, where these supercells form a partition of the space. The image or object histogram is the normalized count of the number of pixels that fall in each of the supercells. As part of the clustering, we also compute a 4096 element table that maps any (R,G,B) triple to its supercell number. To compute an image or object color histogram, we take the (R,G,B) for each pixel, look up its cell number in the table, and increment the corresponding histogram bucket. As a final step, the histogram is normalized so that its sum is unity. *Texture features:* Our texture features are based on modified versions of the coarseness, contrast, and directionality features proposed in [10]. The ***coarseness*** feature helps measure the scale of the texture (pebbles vs. boulders), and is calculated using moving windows of different sizes. The ***contrast*** feature describes the vividness of the pattern, and is a function of the variance of the gray-level histogram. The ***directionality*** feature describes whether or not the image has a favored direction (like grass), or whether it is isotropic (like a smooth

Element	'278 Claim Recitation	Exemplary Citations in Reference
		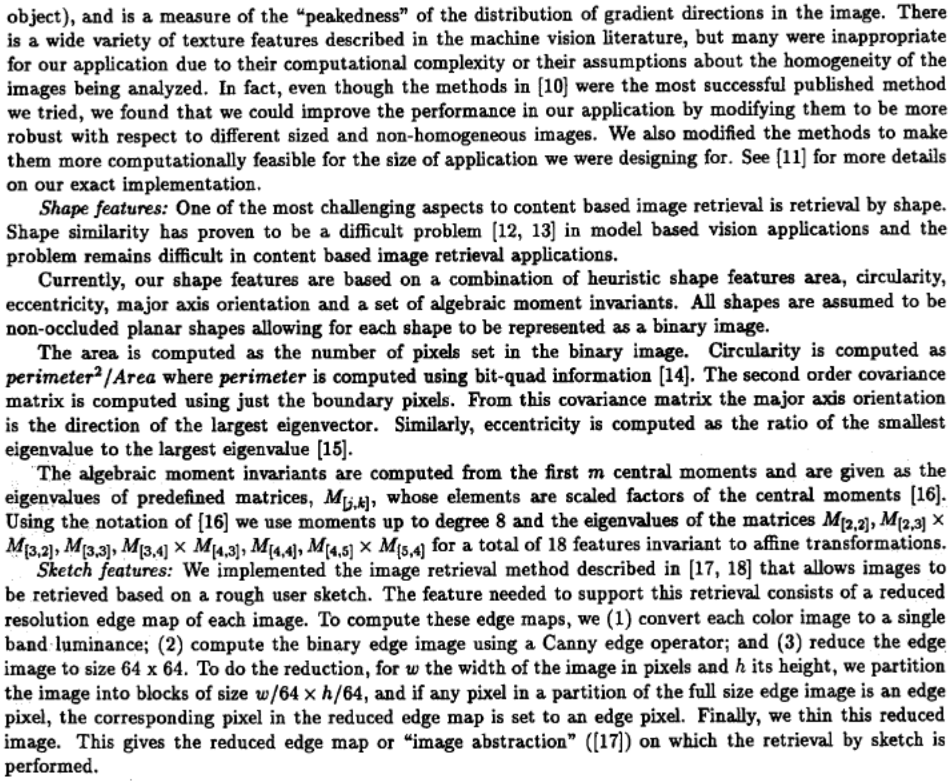 object), and is a measure of the "peakedness" of the distribution of gradient directions in the image. There is a wide variety of texture features described in the machine vision literature, but many were inappropriate for our application due to their computational complexity or their assumptions about the homogeneity of the images being analyzed. In fact, even though the methods in [10] were the most successful published method we tried, we found that we could improve the performance in our application by modifying them to be more robust with respect to different sized and non-homogeneous images. We also modified the methods to make them more computationally feasible for the size of application we were designing for. See [11] for more details on our exact implementation. *Shape features:* One of the most challenging aspects to content based image retrieval is retrieval by shape. Shape similarity has proven to be a difficult problem [12, 13] in model based vision applications and the problem remains difficult in content based image retrieval applications. Currently, our shape features are based on a combination of heuristic shape features area, circularity, eccentricity, major axis orientation and a set of algebraic moment invariants. All shapes are assumed to be non-occluded planar shapes allowing for each shape to be represented as a binary image. The area is computed as the number of pixels set in the binary image. Circularity is computed as $perimeter^2/Area$ where $perimeter$ is computed using bit-quad information [14]. The second order covariance matrix is computed using just the boundary pixels. From this covariance matrix the major axis orientation is the direction of the largest eigenvector. Similarly, eccentricity is computed as the ratio of the smallest eigenvalue to the largest eigenvalue [15]. The algebraic moment invariants are computed from the first m central moments and are given as the eigenvalues of predefined matrices, $M_{[j,k]}$, whose elements are scaled factors of the central moments [16]. Using the notation of [16] we use moments up to degree 8 and the eigenvalues of the matrices $M_{[2,2]}, M_{[2,3]} \times M_{[3,2]}, M_{[3,3]}, M_{[3,4]} \times M_{[4,3]}, M_{[4,4]}, M_{[4,5]} \times M_{[5,4]}$ for a total of 18 features invariant to affine transformations. *Sketch features:* We implemented the image retrieval method described in [17, 18] that allows images to be retrieved based on a rough user sketch. The feature needed to support this retrieval consists of a reduced resolution edge map of each image. To compute these edge maps, we (1) convert each color image to a single band luminance; (2) compute the binary edge image using a Canny edge operator; and (3) reduce the edge image to size 64 x 64. To do the reduction, for w the width of the image in pixels and h its height, we partition the image into blocks of size $w/64 \times h/64$, and if any pixel in a partition of the full size edge image is an edge pixel, the corresponding pixel in the reduced edge map is set to an edge pixel. Finally, we thin this reduced image. This gives the reduced edge map or "image abstraction" ([17]) on which the retrieval by sketch is performed. IBM0002390 at 2393-95:

Element	'278 Claim Recitation	Exemplary Citations in Reference
		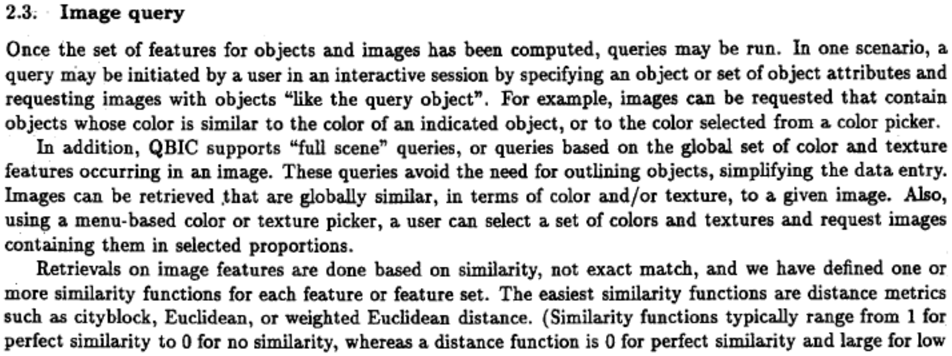 **2.3. Image query** Once the set of features for objects and images has been computed, queries may be run. In one scenario, a query may be initiated by a user in an interactive session by specifying an object or set of object attributes and requesting images with objects "like the query object". For example, images can be requested that contain objects whose color is similar to the color of an indicated object, or to the color selected from a color picker. In addition, QBIC supports "full scene" queries, or queries based on the global set of color and texture features occurring in an image. These queries avoid the need for outlining objects, simplifying the data entry. Images can be retrieved that are globally similar, in terms of color and/or texture, to a given image. Also, using a menu-based color or texture picker, a user can select a set of colors and textures and request images containing them in selected proportions. Retrievals on image features are done based on similarity, not exact match, and we have defined one or more similarity functions for each feature or feature set. The easiest similarity functions are distance metrics such as cityblock, Euclidean, or weighted Euclidean distance. (Similarity functions typically range from 1 for perfect similarity to 0 for no similarity, whereas a distance function is 0 for perfect similarity and large for low

Element	'278 Claim Recitation	Exemplary Citations in Reference
		similarity, so we must convert as necessary to obtain consistent measures.) Different similarity measures may be used, and they may vary from application to application. The matching methods and similarity functions for the different feature types are described below: *Color:* For average color, the distance between a query object and database object is weighted Euclidean distance. The weights are user adjustable so that we may, for example, downweight the L component of an (L, a, b) query. Best results, as judged by our subjective evaluation in on-going experiments, are obtained by representing the average color in the MTM color space, and by inversely weighting each component by its standard deviation over the samples in the database. Improved results are obtained when color matching is done not on average color, but on the distribution of colors that occur in an object or image. In this way, a red and blue object matches other red and blue objects better than to a purple object. To match color histograms, we use a method similar to the one described in [8]. Let X be the query histogram and Y the histogram of an item in the database, both normalized. We compute the element by element difference histogram Z. Then the similarity between X and Y is given by $\Vert Z \Vert = Z^T A Z$, where A is a symmetric color similarity matrix with $$a(i,j) \equiv 1 - d(c_i, c_j)/d_{max}$$ where c_i and c_j are the ith and jth colors in the color histograms, and $d(c_i, c_j)$ is the MTM color distance, and d_{max} is the maximum distance between any two colors ([8] used Luv distance.) This metric gives the (weighted) length of the difference vector between X and Y, weighted by A, which accounts for the perceptual distance between different pairs of colors. The result is a measure that accounts for both the difference between the amounts in any given color (e.g. a particular shade of red), as well as differences between similar colors (e.g. red and orange). Based on our experience, we also found it useful for the user to be able to request images with $x\%$ of color 1, $y\%$ of color 2, etc. where both the percentages and the colors are explicitly specified by the user. This allows searches such as for a beach scene (25% white (for the sand) and 50% blue (for the sky and water)), and for many scenes for which the user has seen before, and for which he/she can remember approximate color percentages. The color picker that allows this type of query is shown in Figure 2. 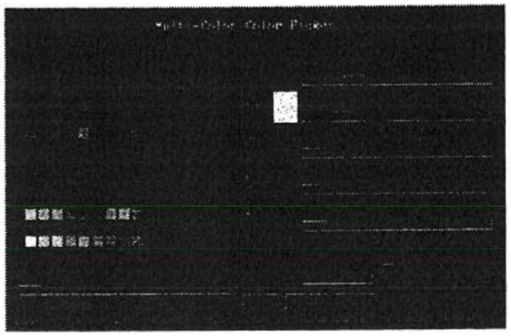 Figure 2: Multi-Color Color Picker with two selected colors. The palette on the left displays the set of selectable colors. Once a color is selected, it appears in a box on the right, and the amount of the color can be specified by the sliders. In this case, 25% light blue and 15% white have been selected. *Texture:* Texture distance is computed as weighted Euclidean distance in the three dimensional texture space. The most common weighting factors are the inverse variances for each component, computed over the

Element	'278 Claim Recitation	Exemplary Citations in Reference
		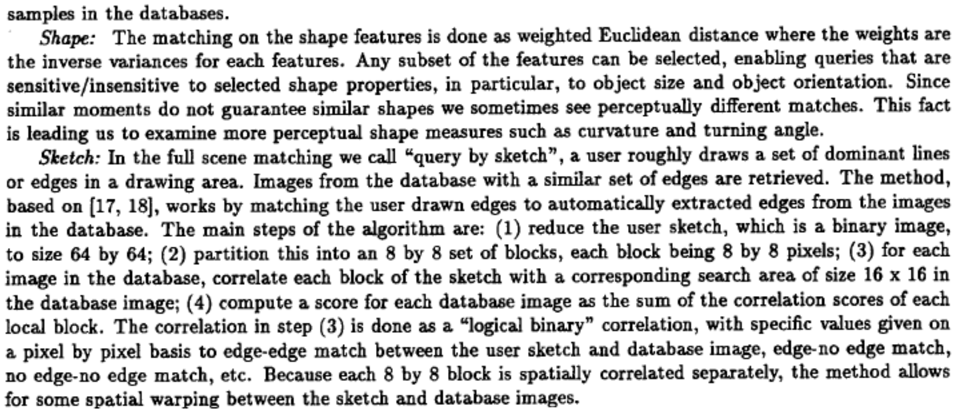 samples in the databases. *Shape:* The matching on the shape features is done as weighted Euclidean distance where the weights are the inverse variances for each features. Any subset of the features can be selected, enabling queries that are sensitive/insensitive to selected shape properties, in particular, to object size and object orientation. Since similar moments do not guarantee similar shapes we sometimes see perceptually different matches. This fact is leading us to examine more perceptual shape measures such as curvature and turning angle. *Sketch:* In the full scene matching we call "query by sketch", a user roughly draws a set of dominant lines or edges in a drawing area. Images from the database with a similar set of edges are retrieved. The method, based on [17, 18], works by matching the user drawn edges to automatically extracted edges from the images in the database. The main steps of the algorithm are: (1) reduce the user sketch, which is a binary image, to size 64 by 64; (2) partition this into an 8 by 8 set of blocks, each block being 8 by 8 pixels; (3) for each image in the database, correlate each block of the sketch with a corresponding search area of size 16 x 16 in the database image; (4) compute a score for each database image as the sum of the correlation scores of each local block. The correlation in step (3) is done as a "logical binary" correlation, with specific values given on a pixel by pixel basis to edge-edge match between the user sketch and database image, edge-no edge match, no edge-no edge match, etc. Because each 8 by 8 block is spatially correlated separately, the method allows for some spatial warping between the sketch and database images. IBM0002390 at 2395: 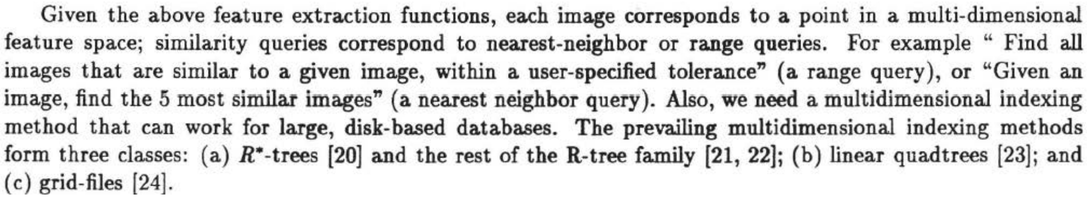 Given the above feature extraction functions, each image corresponds to a point in a multi-dimensional feature space; similarity queries correspond to nearest-neighbor or range queries. For example " Find all images that are similar to a given image, within a user-specified tolerance" (a range query), or "Given an image, find the 5 most similar images" (a nearest neighbor query). Also, we need a multidimensional indexing method that can work for large, disk-based databases. The prevailing multidimensional indexing methods form three classes: (a) R^*-trees [20] and the rest of the R-tree family [21, 22]; (b) linear quadtrees [23]; and (c) grid-files [24]. IBM0002390 at 2396-98

Element	'278 Claim Recitation	Exemplary Citations in Reference
		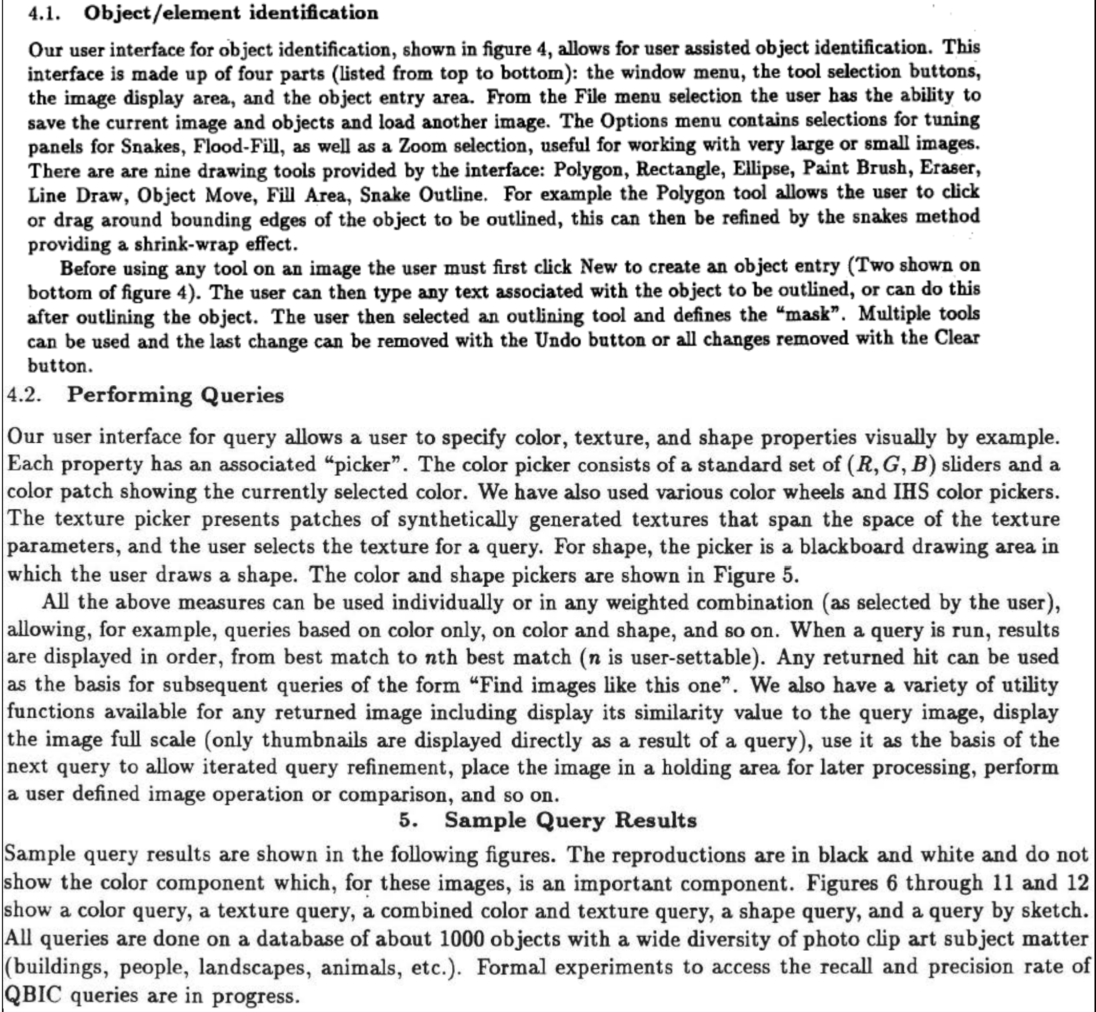 4.1. **Object/element identification** Our user interface for object identification, shown in figure 4, allows for user assisted object identification. This interface is made up of four parts (listed from top to bottom): the window menu, the tool selection buttons, the image display area, and the object entry area. From the File menu selection the user has the ability to save the current image and objects and load another image. The Options menu contains selections for tuning panels for Snakes, Flood-Fill, as well as a Zoom selection, useful for working with very large or small images. There are are nine drawing tools provided by the interface: Polygon, Rectangle, Ellipse, Paint Brush, Eraser, Line Draw, Object Move, Fill Area, Snake Outline. For example the Polygon tool allows the user to click or drag around bounding edges of the object to be outlined, this can then be refined by the snakes method providing a shrink-wrap effect. Before using any tool on an image the user must first click New to create an object entry (Two shown on bottom of figure 4). The user can then type any text associated with the object to be outlined, or can do this after outlining the object. The user then selected an outlining tool and defines the "mask". Multiple tools can be used and the last change can be removed with the Undo button or all changes removed with the Clear button. 4.2. **Performing Queries** Our user interface for query allows a user to specify color, texture, and shape properties visually by example. Each property has an associated "picker". The color picker consists of a standard set of (R, G, B) sliders and a color patch showing the currently selected color. We have also used various color wheels and IHS color pickers. The texture picker presents patches of synthetically generated textures that span the space of the texture parameters, and the user selects the texture for a query. For shape, the picker is a blackboard drawing area in which the user draws a shape. The color and shape pickers are shown in Figure 5. All the above measures can be used individually or in any weighted combination (as selected by the user), allowing, for example, queries based on color only, on color and shape, and so on. When a query is run, results are displayed in order, from best match to nth best match (n is user-settable). Any returned hit can be used as the basis for subsequent queries of the form "Find images like this one". We also have a variety of utility functions available for any returned image including display its similarity value to the query image, display the image full scale (only thumbnails are displayed directly as a result of a query), use it as the basis of the next query to allow iterated query refinement, place the image in a holding area for later processing, perform a user defined image operation or comparison, and so on. 5. **Sample Query Results** Sample query results are shown in the following figures. The reproductions are in black and white and do not show the color component which, for these images, is an important component. Figures 6 through 11 and 12 show a color query, a texture query, a combined color and texture query, a shape query, and a query by sketch. All queries are done on a database of about 1000 objects with a wide diversity of photo clip art subject matter (buildings, people, landscapes, animals, etc.). Formal experiments to access the recall and precision rate of QBIC queries are in progress. The QBIC Project: Querying Images by Content Using Color, Texture, and Shape (IBM 0002390–404) at Figs. 7–9

Element	'278 Claim Recitation	Exemplary Citations in Reference
		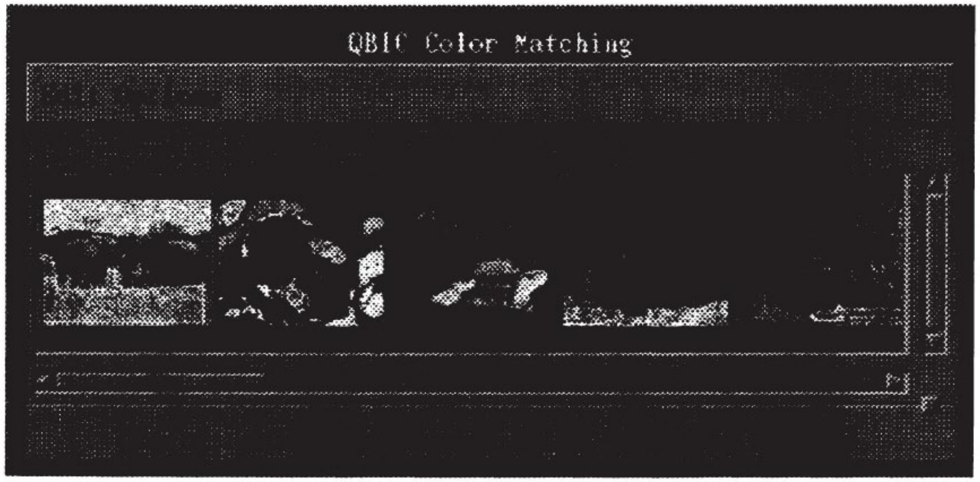 Figure 7: Result of query by example using color only.

Element	'278 Claim Recitation	Exemplary Citations in Reference
		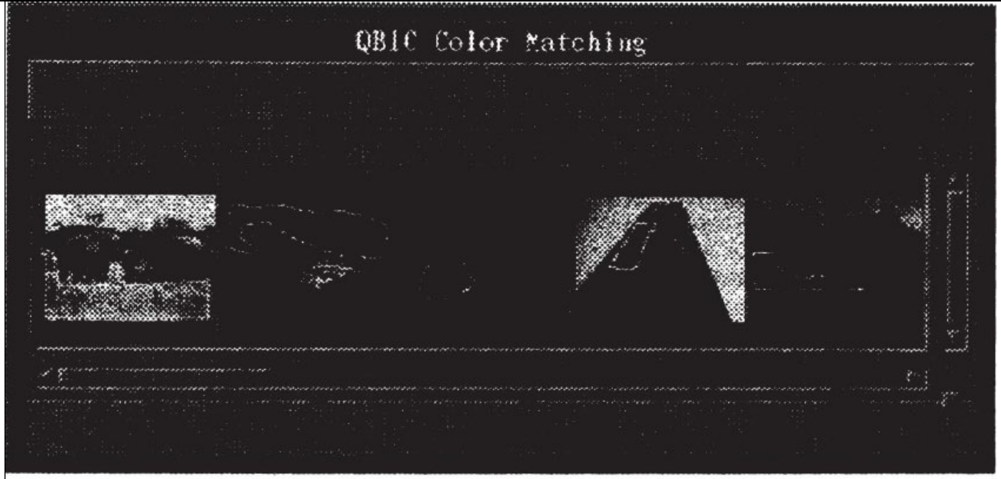 Figure 8: Result of query by example using texture only.

Element	'278 Claim Recitation	Exemplary Citations in Reference
		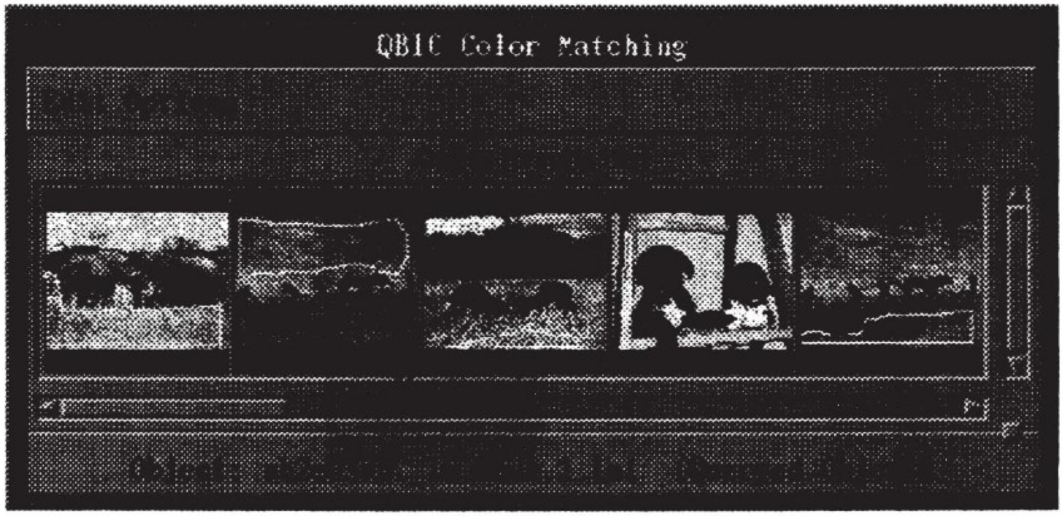 Figure 9: Result of query by example using color and texture. QBIC (Query by Image Content) at 1 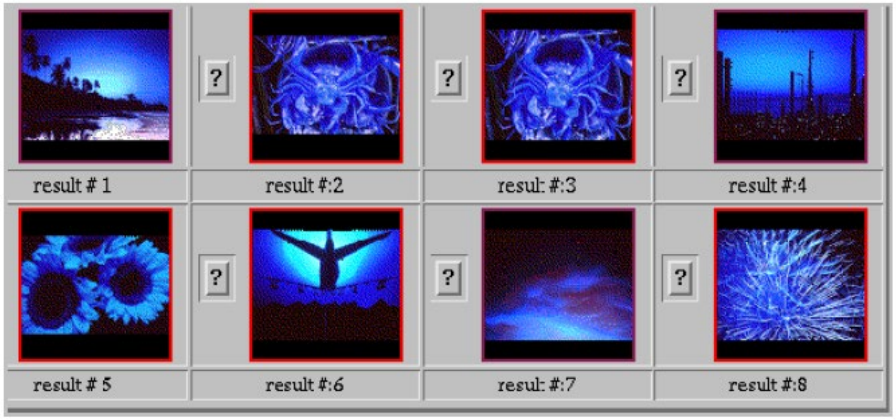 **Result of QBIC Search**

Element	'278 Claim Recitation	Exemplary Citations in Reference
		Query by Image and Video Content: The QBIC System (IBM 0000893–902): "To achieve this functionality, QBIC has two main components: database population (the process of creating an image database) and database query. During the population, images and videos are processed to extract features describing their content-colors, textures, shapes, and camera and object motion- and the features are stored in a database. During the query, the user composes a query graphically. Features are generated from the graphical query and then input to a matching engine that finds images or videos from the database with similar features. Figure 2 shows the system architecture." *Id*. at p. 3.

Element	'278 Claim Recitation	Exemplary Citations in Reference
		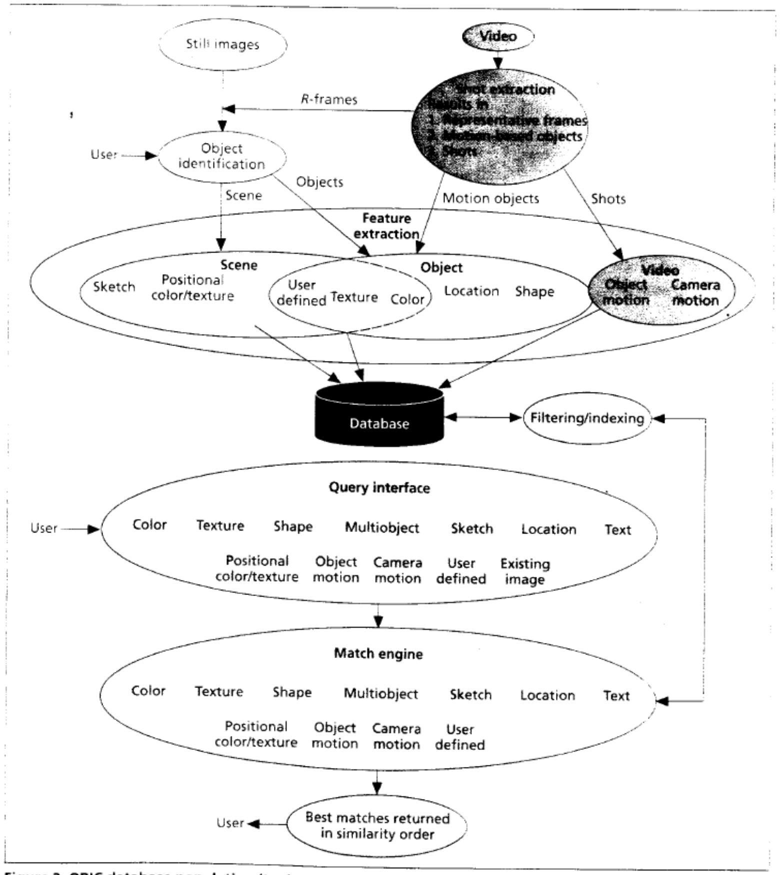 **Figure 2. QBIC database population (top) and query (bottom) architecture.** *Id*. at Fig. 2. "INTEGRATION WITH TEXT AND PARAMETRIC ANNOTATION. Query by visual content complements and extends existing query methods. Systems must be able to integrate queries combining date, subject matter, price, and availability with content properties such as color, texture, and shape." *Id*. at p. 8.

Element	'278 Claim Recitation	Exemplary Citations in Reference
		Image, Audio, and Video Extenders Administration and Programming (IBM 000001–617): "To make an image available for searching by content, you catalog the image. When you catalog an image, the Image Extender analyzes the image, by computing the feature values for the image, and stores the values in a QBIC catalog. When you search for an image by content, your query identifies one or more features for the search (such as average color), a source for each feature (such as an example image), and a target set of cataloged images. The Image Extender computes the feature value of the source and compares it to the cataloged feature values for the target images. It then computes a score that indicates how similar the feature values of the target images are to the source. You can have the Image Extender return the images whose features are most similar to the source. The Image Extender will return the handle of each image and the image score. You can also have the Image Extender return only the score of a single image." *Id.* at pp. 20–21. QBIC Demo (IBM000747) 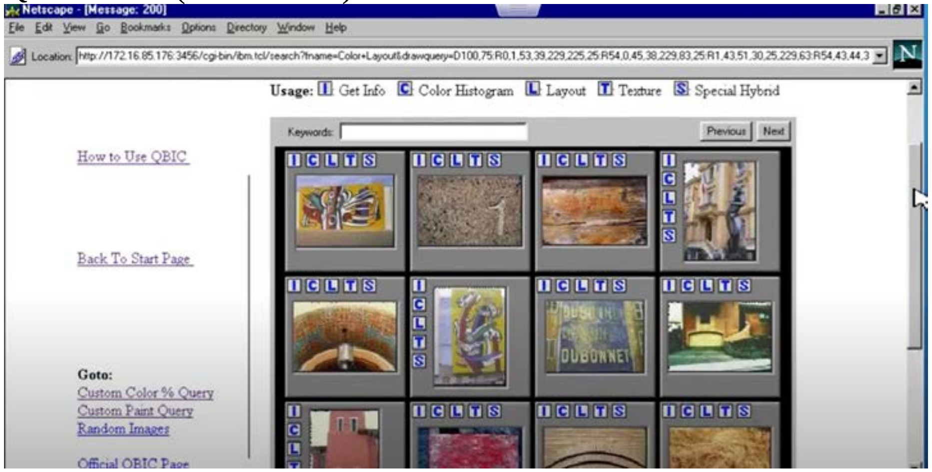 Location: http://172.16.85.176:3456/cgi-bin/ibm.tcl/search?fname=Color+Layout&drawquery=D100,75:R0,1,53,39,229,225,25:R54,0,45,38,229,83,25:R1,43,51,30,25,229,63:R54,43,44,3

Element	'278 Claim Recitation	Exemplary Citations in Reference
		QBIC Demo (IBM000747)  Flickner Depo at 201-202

Element	'278 Claim Recitation	Exemplary Citations in Reference
		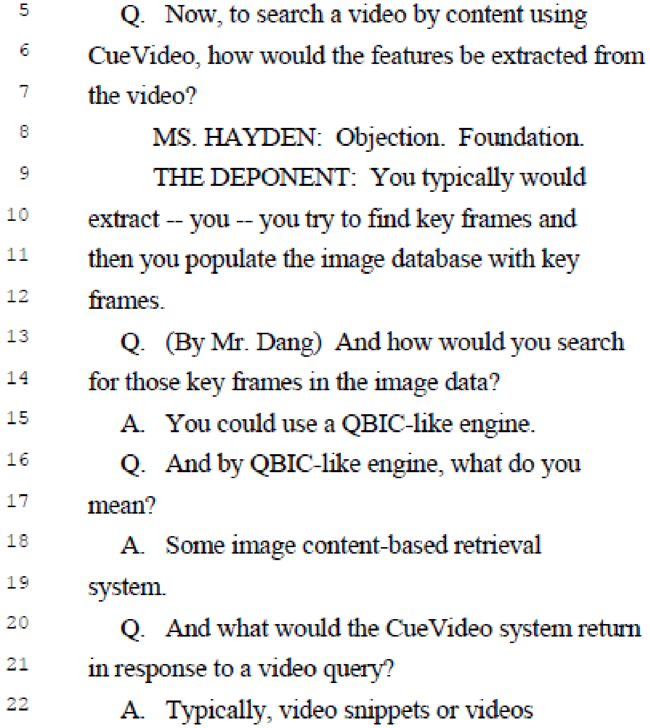 Q. Now, to search a video by content using CueVideo, how would the features be extracted from the video? MS. HAYDEN: Objection. Foundation. THE DEPONENT: You typically would extract -- you -- you try to find key frames and then you populate the image database with key frames. Q. (By Mr. Dang) And how would you search for those key frames in the image data? A. You could use a QBIC-like engine. Q. And by QBIC-like engine, what do you mean? A. Some image content-based retrieval system. Q. And what would the CueVideo system return in response to a video query? A. Typically, video snippets or videos

Element	'278 Claim Recitation	Exemplary Citations in Reference
		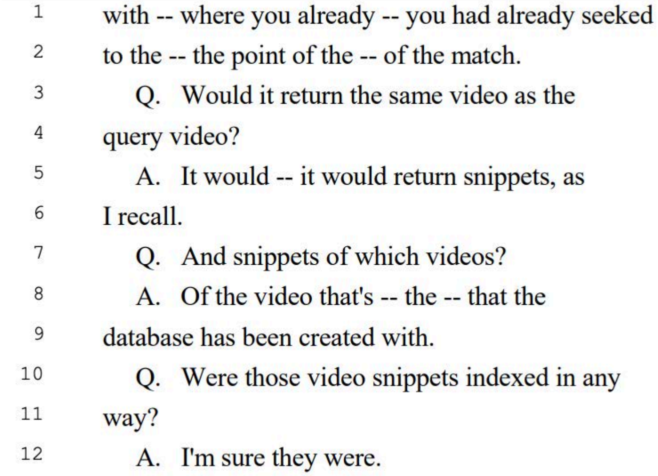with -- where you already -- you had already seeked to the -- the point of the -- of the match. Q. Would it return the same video as the query video? A. It would -- it would return snippets, as I recall. Q. And snippets of which videos? A. Of the video that's -- the -- that the database has been created with. Q. Were those video snippets indexed in any way? A. I'm sure they were. Flickner Depo at 131-132 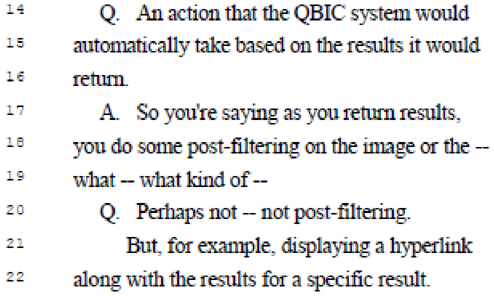Q. An action that the QBIC system would automatically take based on the results it would return. A. So you're saying as you return results, you do some post-filtering on the image or the -- what -- what kind of -- Q. Perhaps not -- not post-filtering. But, for example, displaying a hyperlink along with the results for a specific result. A. Yeah, you could do that. Flickner 8/29/23 Depo at 53–56

Element	'278 Claim Recitation	Exemplary Citations in Reference
		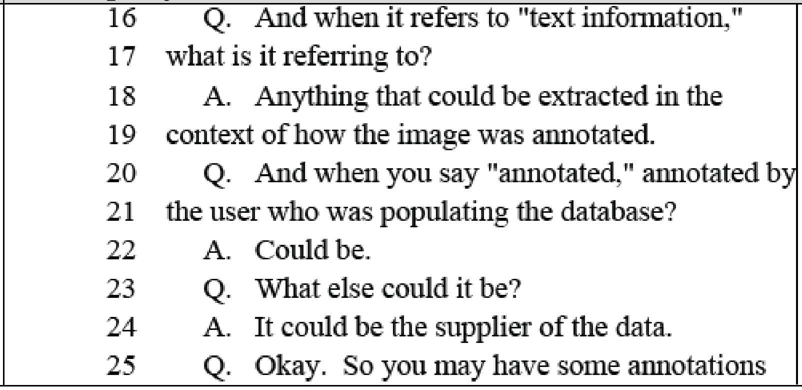 Q. And when it refers to "text information," what is it referring to? A. Anything that could be extracted in the context of how the image was annotated. Q. And when you say "annotated," annotated by the user who was populating the database? A. Could be. Q. What else could it be? A. It could be the supplier of the data. Q. Okay. So you may have some annotations

Element	'278 Claim Recitation	Exemplary Citations in Reference
		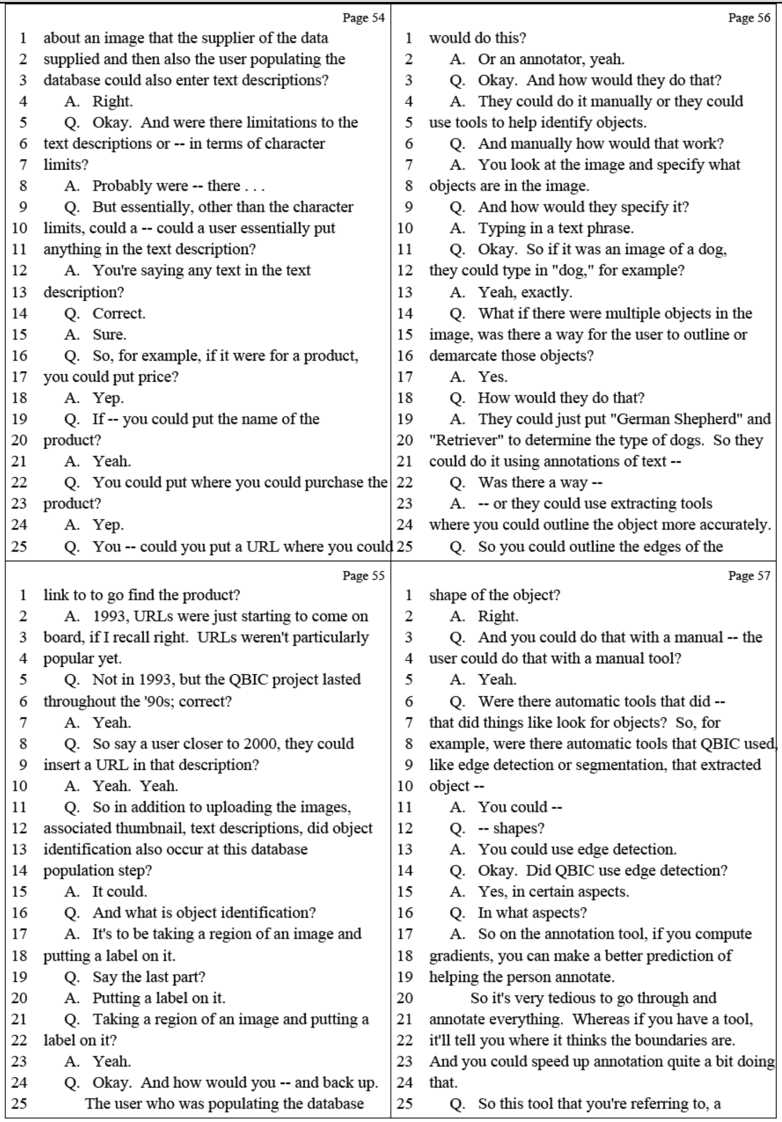

Page 54

about an image that the supplier of the data supplied and then also the user populating the database could also enter text descriptions?
A. Right.
Q. Okay. And were there limitations to the text descriptions or -- in terms of character limits?
A. Probably were -- there . . .
Q. But essentially, other than the character limits, could a -- could a user essentially put anything in the text description?
A. You're saying any text in the text description?
Q. Correct.
A. Sure.
Q. So, for example, if it were for a product, you could put price?
A. Yep.
Q. If -- you could put the name of the product?
A. Yeah.
Q. You could put where you could purchase the product?
A. Yep.
Q. You -- could you put a URL where you could

Page 55

link to to go find the product?
A. 1993, URLs were just starting to come on board, if I recall right. URLs weren't particularly popular yet.
Q. Not in 1993, but the QBIC project lasted throughout the '90s; correct?
A. Yeah.
Q. So say a user closer to 2000, they could insert a URL in that description?
A. Yeah. Yeah.
Q. So in addition to uploading the images, associated thumbnail, text descriptions, did object identification also occur at this database population step?
A. It could.
Q. And what is object identification?
A. It's to be taking a region of an image and putting a label on it.
Q. Say the last part?
A. Putting a label on it.
Q. Taking a region of an image and putting a label on it?
A. Yeah.
Q. Okay. And how would you -- and back up.
The user who was populating the database

Page 56

would do this?
A. Or an annotator, yeah.
Q. Okay. And how would they do that?
A. They could do it manually or they could use tools to help identify objects.
Q. And manually how would that work?
A. You look at the image and specify what objects are in the image.
Q. And how would they specify it?
A. Typing in a text phrase.
Q. Okay. So if it was an image of a dog, they could type in "dog," for example?
A. Yeah, exactly.
Q. What if there were multiple objects in the image, was there a way for the user to outline or demarcate those objects?
A. Yes.
Q. How would they do that?
A. They could just put "German Shepherd" and "Retriever" to determine the type of dogs. So they could do it using annotations of text --
Q. Was there a way --
A. -- or they could use extracting tools where you could outline the object more accurately.
Q. So you could outline the edges of the

Page 57

shape of the object?
A. Right.
Q. And you could do that with a manual -- the user could do that with a manual tool?
A. Yeah.
Q. Were there automatic tools that did -- that did things like look for objects? So, for example, were there automatic tools that QBIC used, like edge detection or segmentation, that extracted object --
A. You could --
Q. -- shapes?
A. You could use edge detection.
Q. Okay. Did QBIC use edge detection?
A. Yes, in certain aspects.
Q. In what aspects?
A. So on the annotation tool, if you compute gradients, you can make a better prediction of helping the person annotate.
So it's very tedious to go through and annotate everything. Whereas if you have a tool, it'll tell you where it thinks the boundaries are. And you could speed up annotation quite a bit doing that.
Q. So this tool that you're referring to, a

Element	’278 Claim Recitation	Exemplary Citations in Reference
		Id. at 68–72 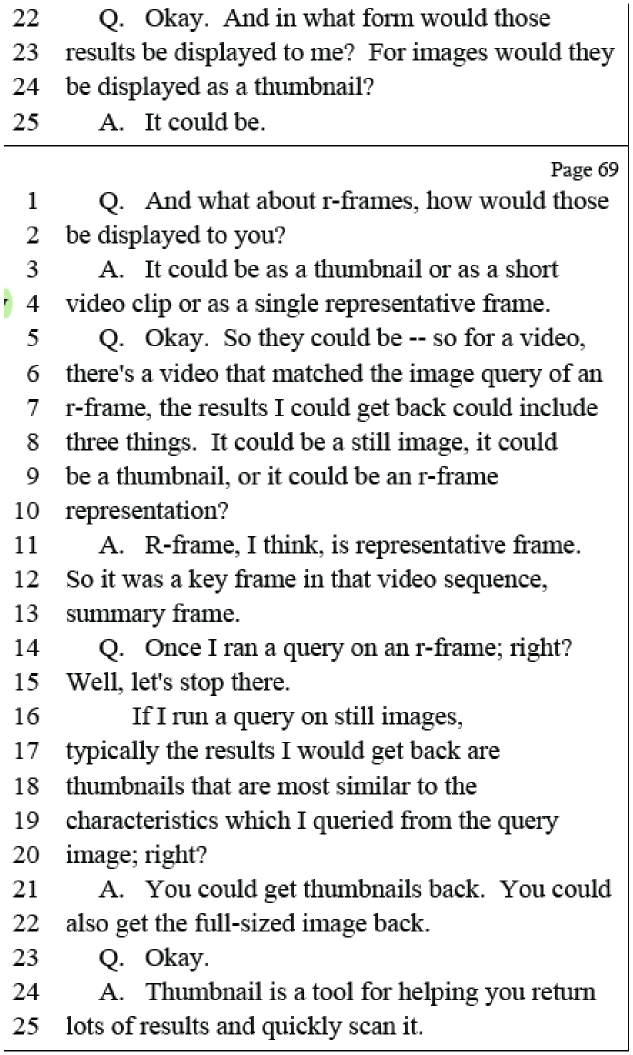 Q. Okay. And in what form would those results be displayed to me? For images would they be displayed as a thumbnail? A. It could be. Page 69 Q. And what about r-frames, how would those be displayed to you? A. It could be as a thumbnail or as a short video clip or as a single representative frame. Q. Okay. So they could be -- so for a video, there's a video that matched the image query of an r-frame, the results I could get back could include three things. It could be a still image, it could be a thumbnail, or it could be an r-frame representation? A. R-frame, I think, is representative frame. So it was a key frame in that video sequence, summary frame. Q. Once I ran a query on an r-frame; right? Well, let's stop there. If I run a query on still images, typically the results I would get back are thumbnails that are most similar to the characteristics which I queried from the query image; right? A. You could get thumbnails back. You could also get the full-sized image back. Q. Okay. A. Thumbnail is a tool for helping you return lots of results and quickly scan it.

Element	'278 Claim Recitation	Exemplary Citations in Reference
		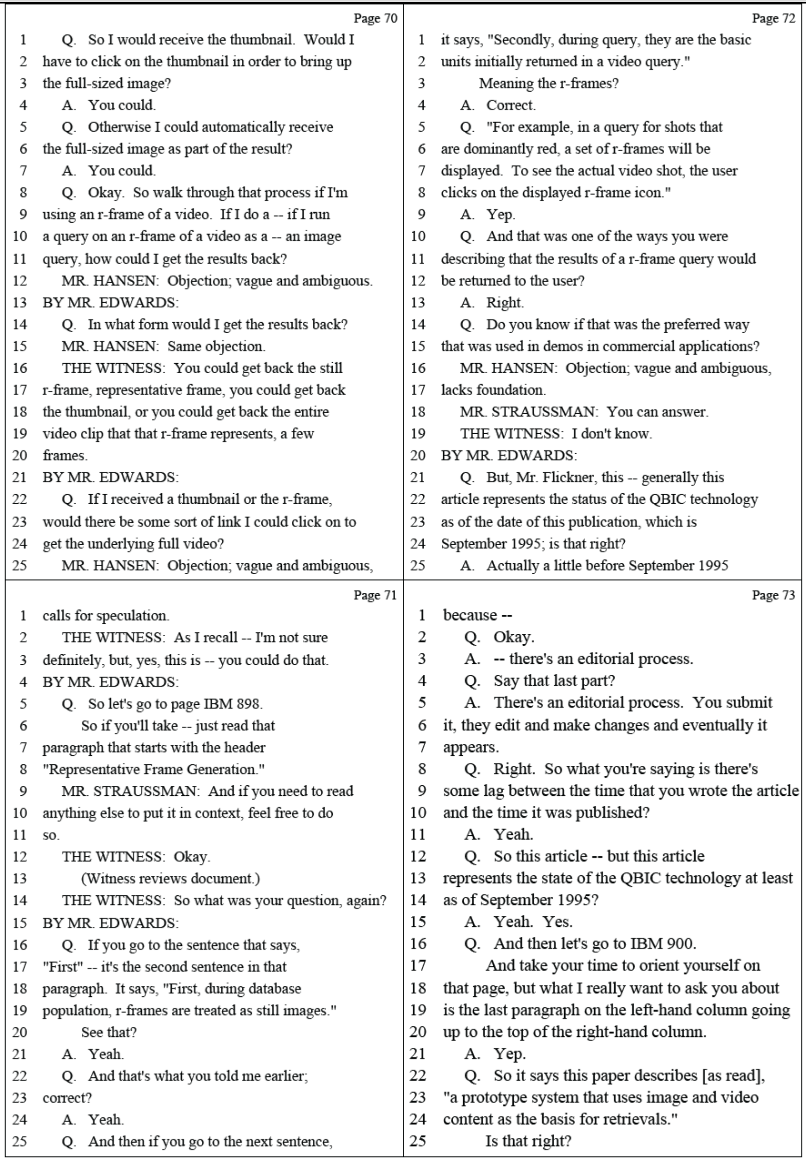 Page 70 Q. So I would receive the thumbnail. Would I have to click on the thumbnail in order to bring up the full-sized image? A. You could. Q. Otherwise I could automatically receive the full-sized image as part of the result? A. You could. Q. Okay. So walk through that process if I'm using an r-frame of a video. If I do a -- if I run a query on an r-frame of a video as a -- an image query, how could I get the results back? MR. HANSEN: Objection; vague and ambiguous. BY MR. EDWARDS: Q. In what form would I get the results back? MR. HANSEN: Same objection. THE WITNESS: You could get back the still r-frame, representative frame, you could get back the thumbnail, or you could get back the entire video clip that that r-frame represents, a few frames. BY MR. EDWARDS: Q. If I received a thumbnail or the r-frame, would there be some sort of link I could click on to get the underlying full video? MR. HANSEN: Objection; vague and ambiguous, Page 71 calls for speculation. THE WITNESS: As I recall -- I'm not sure definitely, but, yes, this is -- you could do that. BY MR. EDWARDS: Q. So let's go to page IBM 898. So if you'll take -- just read that paragraph that starts with the header "Representative Frame Generation." MR. STRAUSSMAN: And if you need to read anything else to put it in context, feel free to do so. THE WITNESS: Okay. (Witness reviews document.) THE WITNESS: So what was your question, again? BY MR. EDWARDS: Q. If you go to the sentence that says, "First" -- it's the second sentence in that paragraph. It says, "First, during database population, r-frames are treated as still images." See that? A. Yeah. Q. And that's what you told me earlier; correct? A. Yeah. Q. And then if you go to the next sentence, Page 72 it says, "Secondly, during query, they are the basic units initially returned in a video query." Meaning the r-frames? A. Correct. Q. "For example, in a query for shots that are dominantly red, a set of r-frames will be displayed. To see the actual video shot, the user clicks on the displayed r-frame icon." A. Yep. Q. And that was one of the ways you were describing that the results of a r-frame query would be returned to the user? A. Right. Q. Do you know if that was the preferred way that was used in demos in commercial applications? MR. HANSEN: Objection; vague and ambiguous, lacks foundation. MR. STRAUSSMAN: You can answer. THE WITNESS: I don't know. BY MR. EDWARDS: Q. But, Mr. Flickner, this -- generally this article represents the status of the QBIC technology as of the date of this publication, which is September 1995; is that right? A. Actually a little before September 1995 Page 73 because -- Q. Okay. A. -- there's an editorial process. Q. Say that last part? A. There's an editorial process. You submit it, they edit and make changes and eventually it appears. Q. Right. So what you're saying is there's some lag between the time that you wrote the article and the time it was published? A. Yeah. Q. So this article -- but this article represents the state of the QBIC technology at least as of September 1995? A. Yeah. Yes. Q. And then let's go to IBM 900. And take your time to orient yourself on that page, but what I really want to ask you about is the last paragraph on the left-hand column going up to the top of the right-hand column. A. Yep. Q. So it says this paper describes [as read], "a prototype system that uses image and video content as the basis for retrievals." Is that right?

Element	'278 Claim Recitation	Exemplary Citations in Reference
		Id. at 118–119 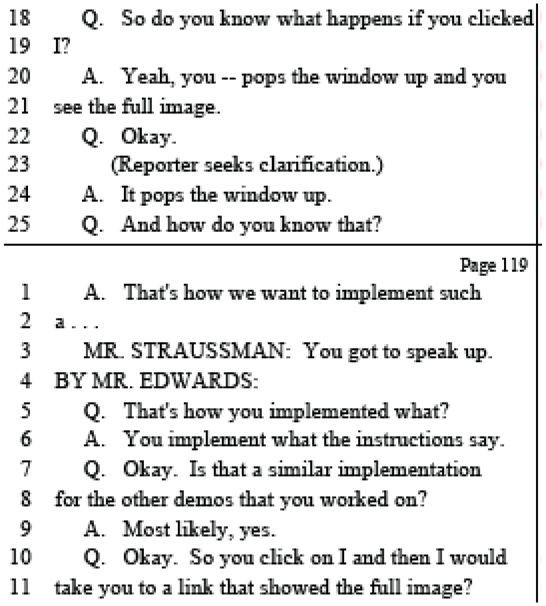 Q. So do you know what happens if you clicked I? A. Yeah, you -- pops the window up and you see the full image. Q. Okay. (Reporter seeks clarification.) A. It pops the window up. Q. And how do you know that? Page 119 A. That's how we want to implement such a . . . MR. STRAUSSMAN: You got to speak up. BY MR. EDWARDS: Q. That's how you implemented what? A. You implement what the instructions say. Q. Okay. Is that a similar implementation for the other demos that you worked on? A. Most likely, yes. Q. Okay. So you click on I and then I would take you to a link that showed the full image? *Id.* at 129 – 131

Element	'278 Claim Recitation	Exemplary Citations in Reference
		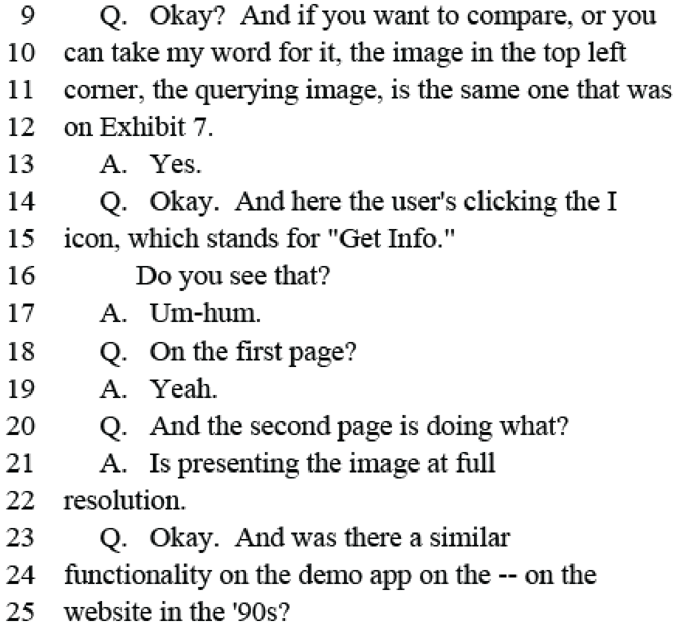 Q. Okay? And if you want to compare, or you can take my word for it, the image in the top left corner, the querying image, is the same one that was on Exhibit 7. A. Yes. Q. Okay. And here the user's clicking the I icon, which stands for "Get Info." Do you see that? A. Um-hum. Q. On the first page? A. Yeah. Q. And the second page is doing what? A. Is presenting the image at full resolution. Q. Okay. And was there a similar functionality on the demo app on the -- on the website in the '90s?

		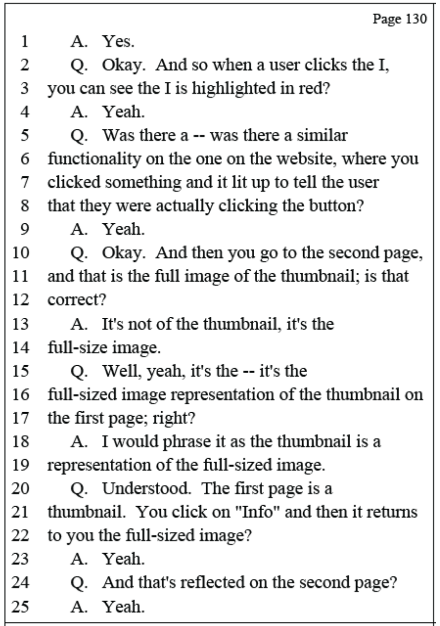 Page 130 A. Yes. Q. Okay. And so when a user clicks the I, you can see the I is highlighted in red? A. Yeah. Q. Was there a -- was there a similar functionality on the one on the website, where you clicked something and it lit up to tell the user that they were actually clicking the button? A. Yeah. Q. Okay. And then you go to the second page, and that is the full image of the thumbnail; is that correct? A. It's not of the thumbnail, it's the full-size image. Q. Well, yeah, it's the -- it's the full-sized image representation of the thumbnail on the first page; right? A. I would phrase it as the thumbnail is a representation of the full-sized image. Q. Understood. The first page is a thumbnail. You click on "Info" and then it returns to you the full-sized image? A. Yeah. Q. And that's reflected on the second page? A. Yeah. 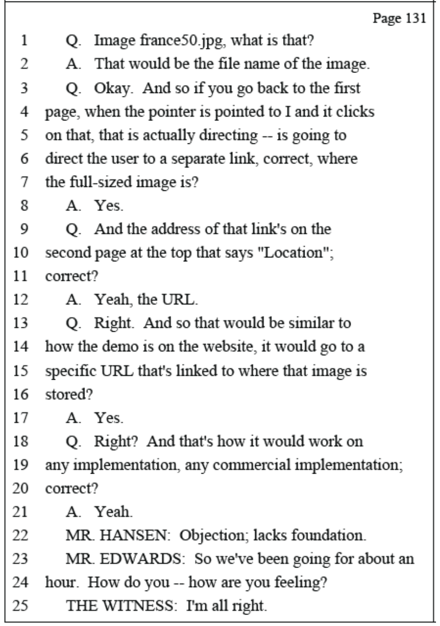 Page 131 Q. Image france50.jpg, what is that? A. That would be the file name of the image. Q. Okay. And so if you go back to the first page, when the pointer is pointed to I and it clicks on that, that is actually directing -- is going to direct the user to a separate link, correct, where the full-sized image is? A. Yes. Q. And the address of that link's on the second page at the top that says "Location"; correct? A. Yeah, the URL. Q. Right. And so that would be similar to how the demo is on the website, it would go to a specific URL that's linked to where that image is stored? A. Yes. Q. Right? And that's how it would work on any implementation, any commercial implementation; correct? A. Yeah. MR. HANSEN: Objection; lacks foundation. MR. EDWARDS: So we've been going for about an hour. How do you -- how are you feeling? THE WITNESS: I'm all right.

Element	'278 Claim Recitation	Exemplary Citations in Reference
		Id. at 160–161 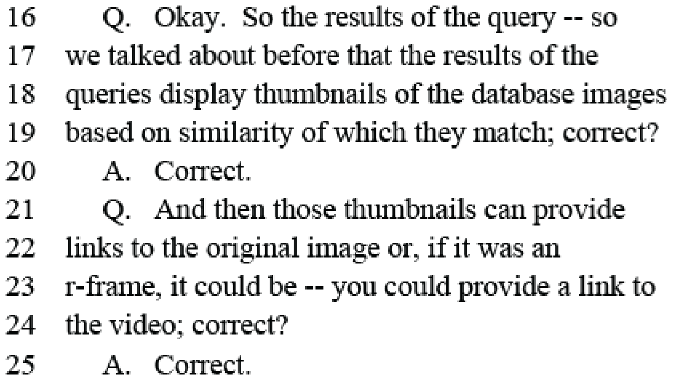 Q. Okay. So the results of the query -- so we talked about before that the results of the queries display thumbnails of the database images based on similarity of which they match; correct? A. Correct. Q. And then those thumbnails can provide links to the original image or, if it was an r-frame, it could be -- you could provide a link to the video; correct? A. Correct. 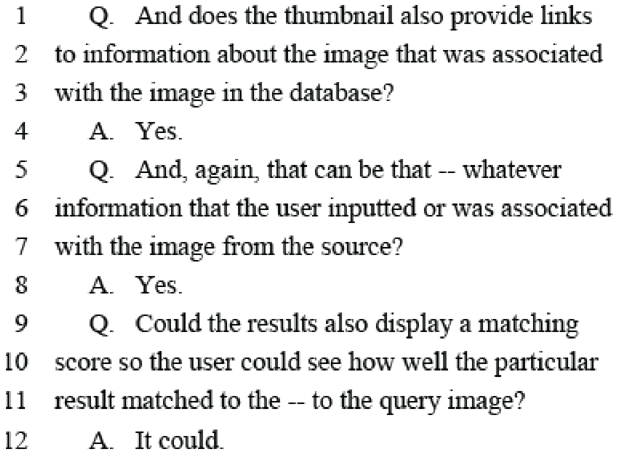 Q. And does the thumbnail also provide links to information about the image that was associated with the image in the database? A. Yes. Q. And, again, that can be that -- whatever information that the user inputted or was associated with the image from the source? A. Yes. Q. Could the results also display a matching score so the user could see how well the particular result matched to the -- to the query image? A. It could. *Id.* at 321

Element	'278 Claim Recitation	Exemplary Citations in Reference
		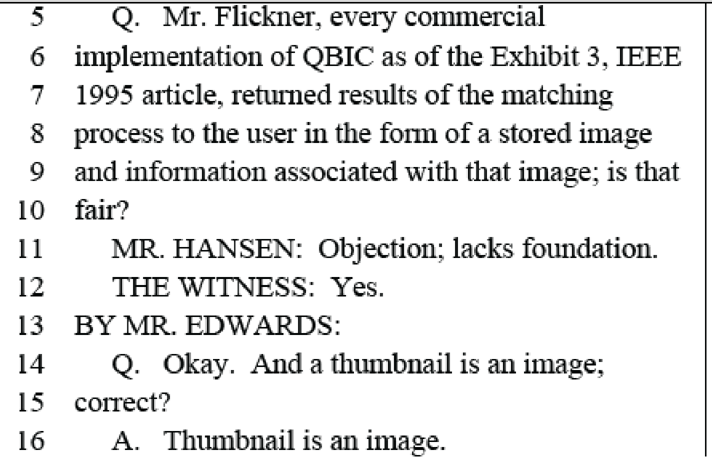 Q. Mr. Flickner, every commercial implementation of QBIC as of the Exhibit 3, IEEE 1995 article, returned results of the matching process to the user in the form of a stored image and information associated with that image; is that fair? MR. HANSEN: Objection; lacks foundation. THE WITNESS: Yes. BY MR. EDWARDS: Q. Okay. And a thumbnail is an image; correct? A. Thumbnail is an image. *See also* Flickner Depo at 54-55, 75-76, 78-79, 88, 89-90, 159, 209-210 Flickner 8/29/23 Depo at Ex. 6 at 1

Element	'278 Claim Recitation	Exemplary Citations in Reference
		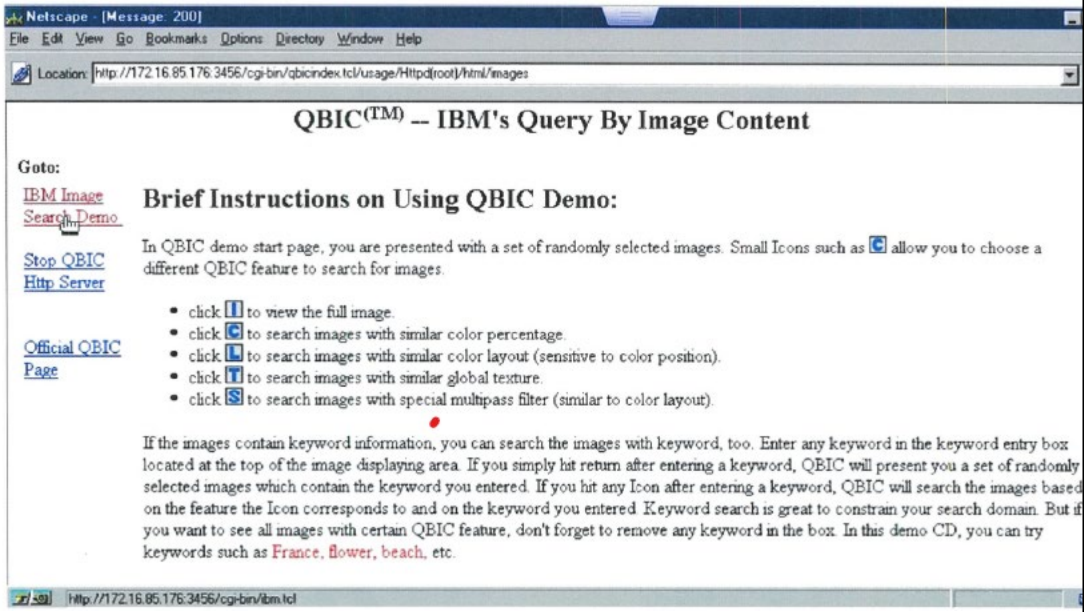 **QBIC$^{(TM)}$ -- IBM's Query By Image Content** **Brief Instructions on Using QBIC Demo:** In QBIC demo start page, you are presented with a set of randomly selected images. Small Icons such as C allow you to choose a different QBIC feature to search for images. • click I to view the full image. • click C to search images with similar color percentage. • click L to search images with similar color layout (sensitive to color position). • click T to search images with similar global texture. • click S to search images with special multipass filter (similar to color layout). If the images contain keyword information, you can search the images with keyword, too. Enter any keyword in the keyword entry box located at the top of the image displaying area. If you simply hit return after entering a keyword, QBIC will present you a set of randomly selected images which contain the keyword you entered. If you hit any Icon after entering a keyword, QBIC will search the images based on the feature the Icon corresponds to and on the keyword you entered. Keyword search is great to constrain your search domain. But if you want to see all images with certain QBIC feature, don't forget to remove any keyword in the box. In this demo CD, you can try keywords such as France, flower, beach, etc. *Id.* at 2.

Element	'278 Claim Recitation	Exemplary Citations in Reference
		 Location: http://172.16.85.176:3456/cgi-bin/ibm.tcl Flickner 8/29/23 Depo Ex. 7 at 1

Element	'278 Claim Recitation	Exemplary Citations in Reference
		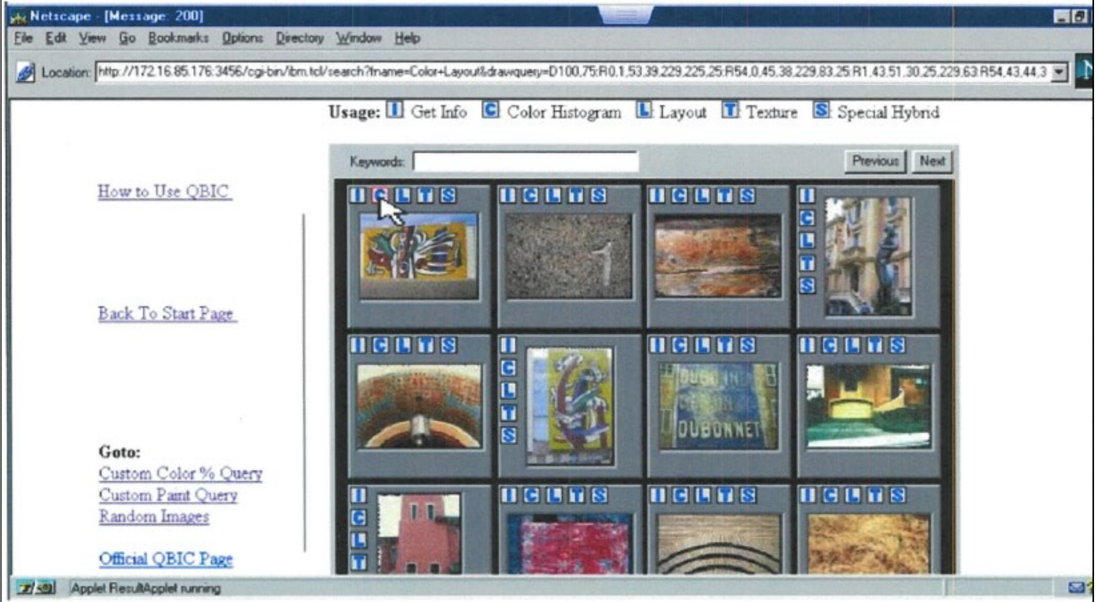 Netscape - [Message: 200] Location: http://172.16.85.176:3456/cgi-bin/ibm.tcl/search?fname=Color+Layout&drawquery=D100,75:R0,1,53,39,229,225,25:R54,0,45,38,229,83,25:R1,43,51,30,25,229,63:R54,43,44,3 **Usage:** I Get Info C Color Histogram L Layout T Texture S Special Hybrid *Id.* at 2

Element	'278 Claim Recitation	Exemplary Citations in Reference
		 Location: http://172.16.85.176:3456/cgi-bin/ibm.tcl/search?fname=Color+Histogram&france50%2Ejpg.x=0 Flickner 8/29/23 Depo at Ex. 8

Element	'278 Claim Recitation	Exemplary Citations in Reference
		 *Id.* at 2

Element	'278 Claim Recitation	Exemplary Citations in Reference
		 http://172.16.85.176:3456/cgi-bin/ibm.tcl/info?france50%2Ejpg Image france50.jpg *Id.* at 3

Element	'278 Claim Recitation	Exemplary Citations in Reference
		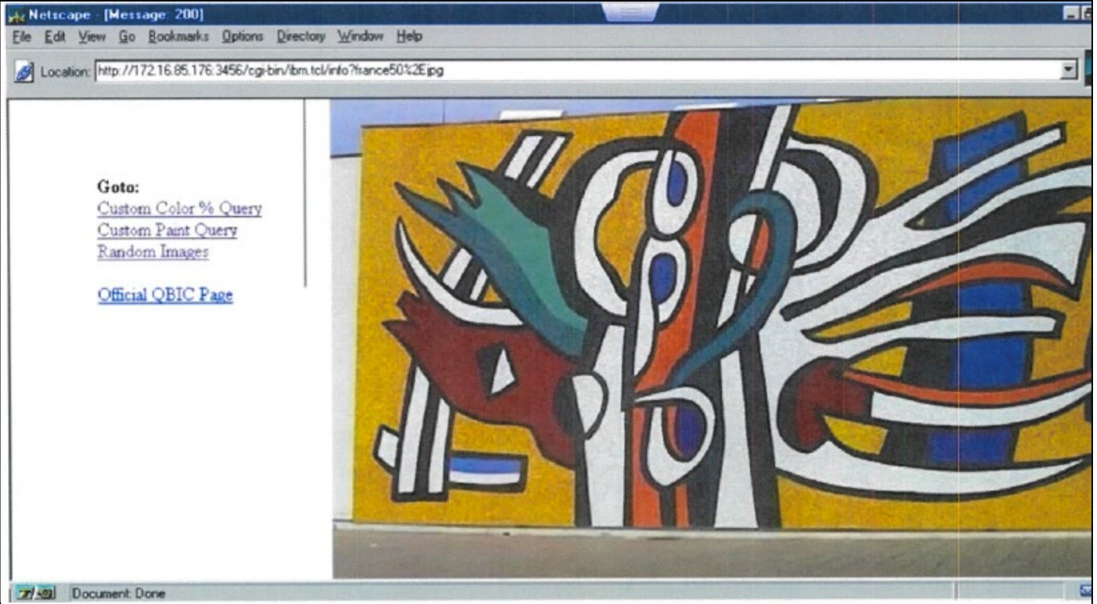 Flickner 8/29/23 Depo at Ex. 9 To the extent it is found that QBIC does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that QBIC does not anticipate this claim, QBIC renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine QBIC with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar

Element	'278 Claim Recitation	Exemplary Citations in Reference
		known problems, according to known methods which would yield predictable results.
[1F.i]	wherein the first target object information comprises a response regarding the validity of the document.	QBIC discloses that the first target object information comprises a response regarding the validity of the document. Image, Audio, and Video Extenders Administration and Programming (IBM 000001–617): "DB2 (DB2) Universal Database (UDB) is a powerful, object-relational database manager. It stores and protects traditional numeric and character data, as well as large, complex objects (LOBs). DB2 extenders help you exploit DB2 's object-relational features. The extenders define distinct data types and special functions for image, audio, video, and text objects. … Because of that, the extenders give your applications a single point of access to any or all of these types of data, along with traditional numeric and character data." *Id*. at p. 3. "A trigger defines a set of actions that are activated by a change to a table. Triggers can be used to perform actions such as validating input data, automatically generating a value for a newly inserted row, reading from other tables for cross-referencing purposes, or writing to other tables for auditing purposes. Triggers are often used for integrity checking or to enforce business rules." *Id*. at 16–17. "When you search for an image by content, your query identifies one or more features for the search (such as average color), a source for each feature (such as an example image), and a target set of cataloged images. The Image Extender computes the feature value of the source and compares it to the cataloged feature values for the target images. It then computes a score that indicates how similar the feature values of the target images are to the source. You can have the Image Extender return the images whose features are most similar to the source. The Image Extender will return the handle of each image

Element	'278 Claim Recitation	Exemplary Citations in Reference
		and the image score. You can also have the Image Extender return only the score of a single image." *Id.* at pp. 20–21. Managing Enterprise Information Portal (IBM 000777–880): "Many paper-intensive enterprises, such as insurance companies and financial institutions, administer large volumes of business-related content. The need for an enterprise solution for managing and accessing business information spans many industries. A content server is a software system that stores multimedia, business forms, documents, and related data, along with metadata that allows employees to process and work with the content." *Id.* at p. 3. "The image search server uses IBM's QBIC (query by image content) technology to help you search for objects by certain visual properties, such as color and texture. The image search server analyzes images and stores the image information in a database. Then users can run image queries, which use the visual properties of images, to match colors, textures, and their positions without describing them in words. You can combine content-based queries with text and keyword searches for more powerful retrieval of image and multimedia data." *Id.* at p. 63. Flickner 8/29/23 Depo at 187–188

Element	'278 Claim Recitation	Exemplary Citations in Reference
		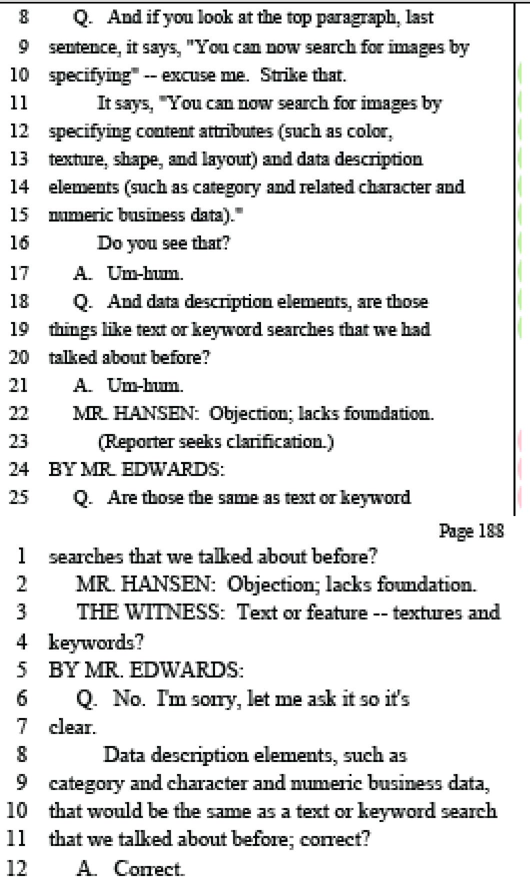 Q. And if you look at the top paragraph, last sentence, it says, "You can now search for images by specifying" -- excuse me. Strike that. It says, "You can now search for images by specifying content attributes (such as color, texture, shape, and layout) and data description elements (such as category and related character and numeric business data)." Do you see that? A. Um-hum. Q. And data description elements, are those things like text or keyword searches that we had talked about before? A. Um-hum. MR. HANSEN: Objection; lacks foundation. (Reporter seeks clarification.) BY MR. EDWARDS: Q. Are those the same as text or keyword Page 188 searches that we talked about before? MR. HANSEN: Objection; lacks foundation. THE WITNESS: Text or feature -- textures and keywords? BY MR. EDWARDS: Q. No. I'm sorry, let me ask it so it's clear. Data description elements, such as category and character and numeric business data, that would be the same as a text or keyword search that we talked about before; correct? A. Correct. *Id.* at 190–191 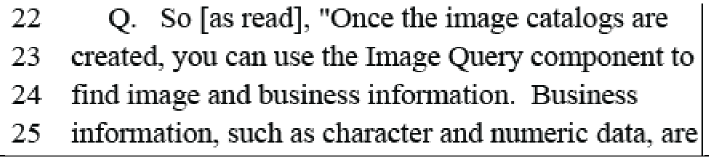 Q. So [as read], "Once the image catalogs are created, you can use the Image Query component to find image and business information. Business information, such as character and numeric data, are

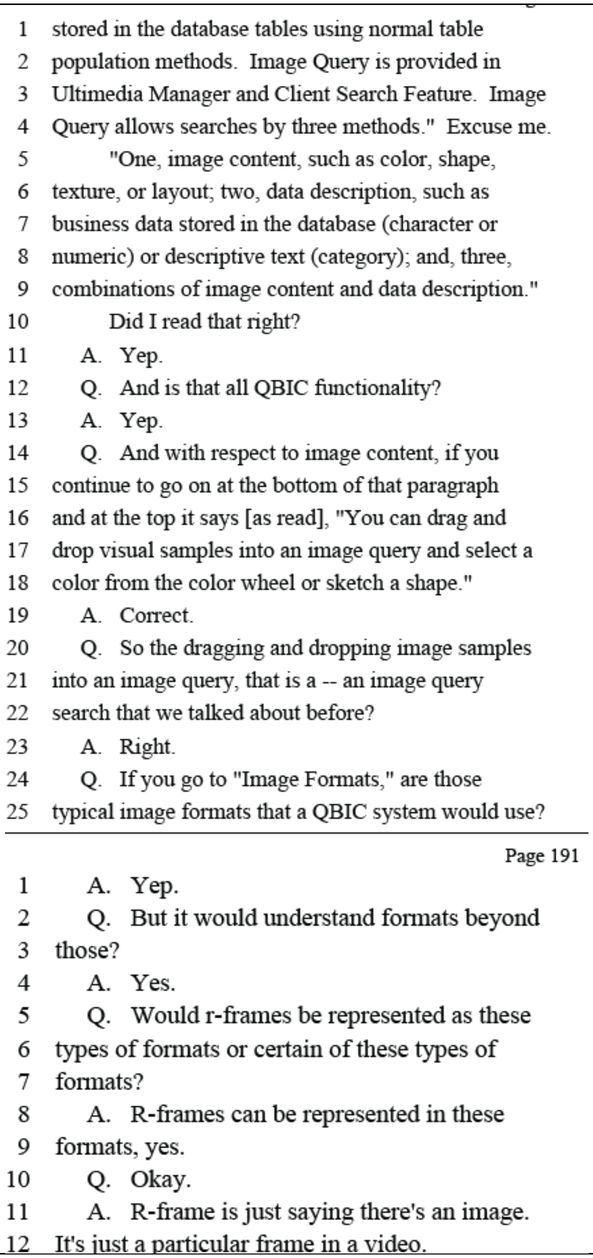

stored in the database tables using normal table population methods. Image Query is provided in Ultimedia Manager and Client Search Feature. Image Query allows searches by three methods." Excuse me.

"One, image content, such as color, shape, texture, or layout; two, data description, such as business data stored in the database (character or numeric) or descriptive text (category); and, three, combinations of image content and data description."

Did I read that right?

A. Yep.

Q. And is that all QBIC functionality?

A. Yep.

Q. And with respect to image content, if you continue to go on at the bottom of that paragraph and at the top it says [as read], "You can drag and drop visual samples into an image query and select a color from the color wheel or sketch a shape."

A. Correct.

Q. So the dragging and dropping image samples into an image query, that is a -- an image query search that we talked about before?

A. Right.

Q. If you go to "Image Formats," are those typical image formats that a QBIC system would use?

Page 191

A. Yep.

Q. But it would understand formats beyond those?

A. Yes.

Q. Would r-frames be represented as these types of formats or certain of these types of formats?

A. R-frames can be represented in these formats, yes.

Q. Okay.

A. R-frame is just saying there's an image. It's just a particular frame in a video.

Element	'278 Claim Recitation	Exemplary Citations in Reference
		Id. at 197–198 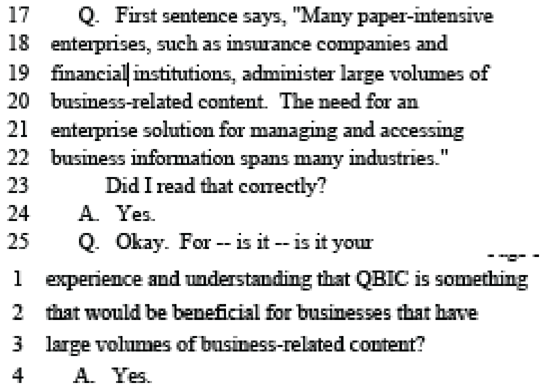 Q. First sentence says, "Many paper-intensive enterprises, such as insurance companies and financial institutions, administer large volumes of business-related content. The need for an enterprise solution for managing and accessing business information spans many industries." Did I read that correctly? A. Yes. Q. Okay. For -- is it -- is it your experience and understanding that QBIC is something that would be beneficial for businesses that have large volumes of business-related content? A. Yes. *Id.* at 307 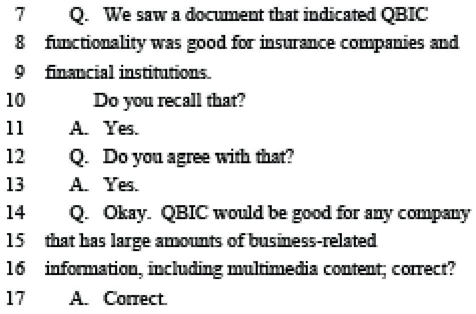 Q. We saw a document that indicated QBIC functionality was good for insurance companies and financial institutions. Do you recall that? A. Yes. Q. Do you agree with that? A. Yes. Q. Okay. QBIC would be good for any company that has large amounts of business-related information, including multimedia content; correct? A. Correct. *See also* Flickner Depo at 22-23, 54-55, 89-90, 143. To the extent it is found that QBIC does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that QBIC does not anticipate this claim, QBIC renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other

Element	'278 Claim Recitation	Exemplary Citations in Reference
		prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine QBIC with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

2. Claim 3 (Non-Asserted)

Claim	'278 Claim Recitation	Exemplary Citations in Reference
3	The method of claim 1, wherein the symbolic content comprises a unique code.	QBIC discloses the method of claim 1, wherein the symbolic content comprises a unique code. Query by Image and Video Content: The QBIC System (IBM 0000893–902): "We have developed the QBIC (Query by Image Content) system to explore content-based retrieval methods. QBIC allows queries on large image and video databases based on • example images, • user-constructed sketches and drawings, • selected color and texture patterns, • camera and object motion, and • other graphical information." *Id.* at pp. 1–3. "INTEGRATION WITH TEXT AND PARAMETRIC ANNOTATION. Query by visual content complements and extends existing query methods. Systems must be able to integrate queries combining date, subject matter, price, and availability with content properties such as color, texture, and shape." *Id.* at p. 8. Image, Audio, and Video Extenders Administration and Programming (IBM 000001–617): "DB2 (DB2) Universal Database (UDB) is a powerful, object-relational database manager. It stores and protects traditional numeric and character data, as well as large, complex objects (LOBs). DB2 extenders help you exploit DB2 's object-relational features. The extenders define distinct data types and special functions for image, audio, video, and text objects. … Because of that, the extenders give your applications a single point of access to any or all of these types of data, along with traditional numeric and character data." *Id.* at p. 3. "A trigger defines a set of actions that are activated by a change to a table. Triggers can be used to perform actions such as validating input data, automatically generating a value for a newly inserted row, reading from other

		tables for cross-referencing purposes, or writing to other tables for auditing purposes. Triggers are often used for integrity checking or to enforce business rules." *Id.* at 16–17. Managing Enterprise Information Portal (IBM 000777–880): "Many paper-intensive enterprises, such as insurance companies and financial institutions, administer large volumes of business-related content. The need for an enterprise solution for managing and accessing business information spans many industries. A content server is a software system that stores multimedia, business forms, documents, and related data, along with metadata that allows employees to process and work with the content." *Id.* at p. 3. "The image search server uses IBM's QBIC (query by image content) technology to help you search for objects by certain visual properties, such as color and texture. The image search server analyzes images and stores the image information in a database. Then users can run image queries, which use the visual properties of images, to match colors, textures, and their positions without describing them in words. You can combine content-based queries with text and keyword searches for more powerful retrieval of image and multimedia data." *Id.* at p. 63. Flickner Depo at 185, 217. Flickner 8/29/31 Depo at 44 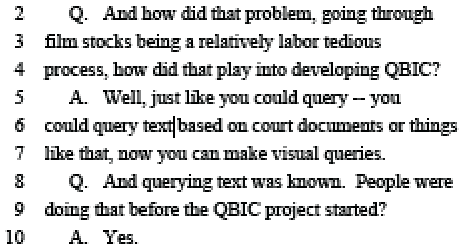 Q. And how did that problem, going through film stocks being a relatively labor tedious process, how did that play into developing QBIC? A. Well, just like you could query -- you could query text based on court documents or things like that, now you can make visual queries. Q. And querying text was known. People were doing that before the QBIC project started? A. Yes.

		Id. at 45–47 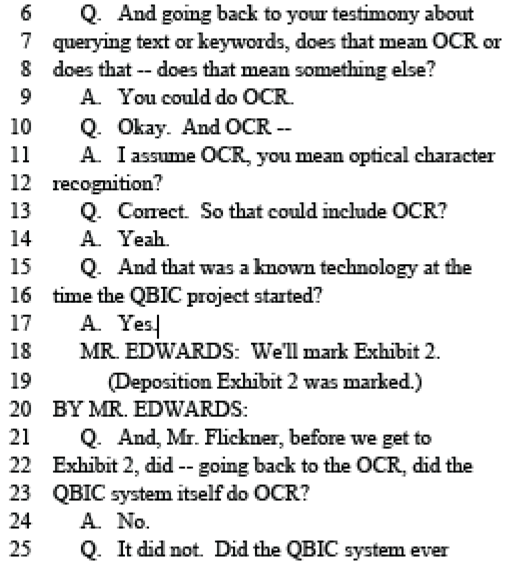 Q. And going back to your testimony about querying text or keywords, does that mean OCR or does that -- does that mean something else? A. You could do OCR. Q. Okay. And OCR -- A. I assume OCR, you mean optical character recognition? Q. Correct. So that could include OCR? A. Yeah. Q. And that was a known technology at the time the QBIC project started? A. Yes. MR. EDWARDS: We'll mark Exhibit 2. (Deposition Exhibit 2 was marked.) BY MR. EDWARDS: Q. And, Mr. Flickner, before we get to Exhibit 2, did -- going back to the OCR, did the QBIC system itself do OCR? A. No. Q. It did not. Did the QBIC system ever

		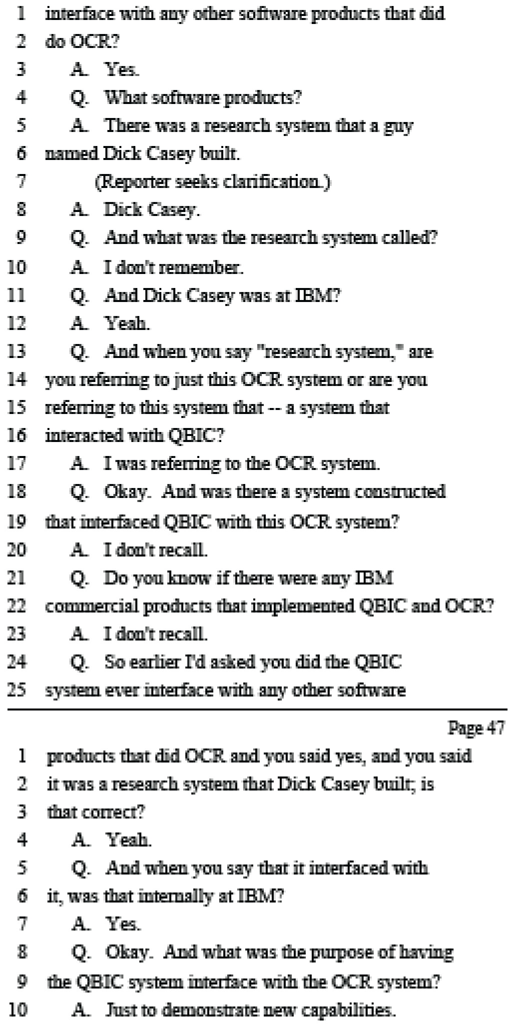 interface with any other software products that did do OCR? A. Yes. Q. What software products? A. There was a research system that a guy named Dick Casey built. (Reporter seeks clarification.) A. Dick Casey. Q. And what was the research system called? A. I don't remember. Q. And Dick Casey was at IBM? A. Yeah. Q. And when you say "research system," are you referring to just this OCR system or are you referring to this system that -- a system that interacted with QBIC? A. I was referring to the OCR system. Q. Okay. And was there a system constructed that interfaced QBIC with this OCR system? A. I don't recall. Q. Do you know if there were any IBM commercial products that implemented QBIC and OCR? A. I don't recall. Q. So earlier I'd asked you did the QBIC system ever interface with any other software Page 47 products that did OCR and you said yes, and you said it was a research system that Dick Casey built; is that correct? A. Yeah. Q. And when you say that it interfaced with it, was that internally at IBM? A. Yes. Q. Okay. And what was the purpose of having the QBIC system interface with the OCR system? A. Just to demonstrate new capabilities. *Id.* at 76

		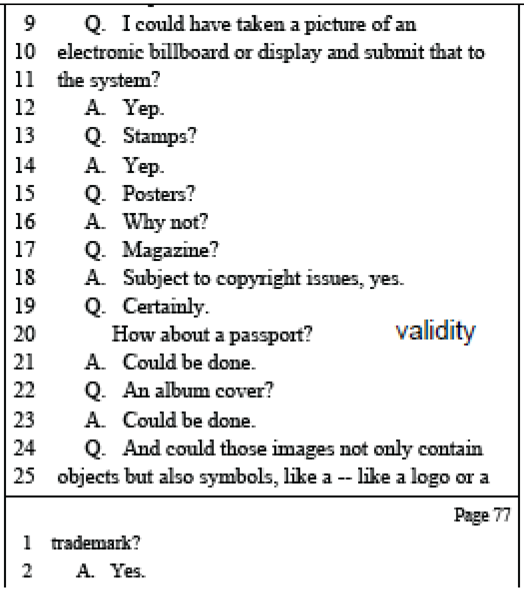 Q. I could have taken a picture of an electronic billboard or display and submit that to the system? A. Yep. Q. Stamps? A. Yep. Q. Posters? A. Why not? Q. Magazine? A. Subject to copyright issues, yes. Q. Certainly. How about a passport? validity A. Could be done. Q. An album cover? A. Could be done. Q. And could those images not only contain objects but also symbols, like a -- like a logo or a Page 77 trademark? A. Yes. *Id.* at 187–188 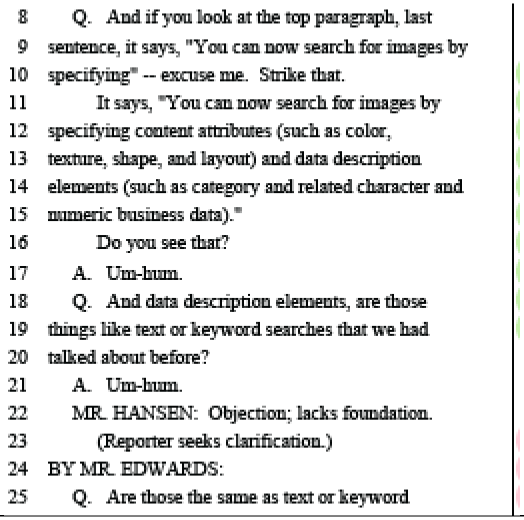 Q. And if you look at the top paragraph, last sentence, it says, "You can now search for images by specifying" -- excuse me. Strike that. It says, "You can now search for images by specifying content attributes (such as color, texture, shape, and layout) and data description elements (such as category and related character and numeric business data)." Do you see that? A. Um-hum. Q. And data description elements, are those things like text or keyword searches that we had talked about before? A. Um-hum. MR. HANSEN: Objection; lacks foundation. (Reporter seeks clarification.) BY MR. EDWARDS: Q. Are those the same as text or keyword	

		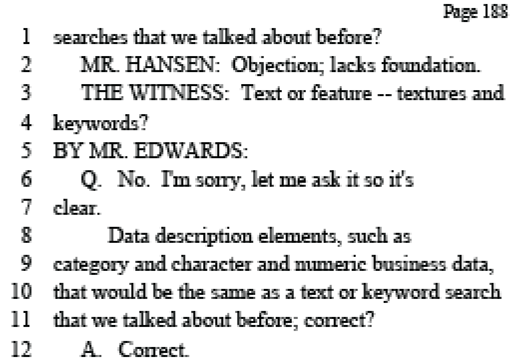 Page 188 searches that we talked about before? MR. HANSEN: Objection; lacks foundation. THE WITNESS: Text or feature -- textures and keywords? BY MR. EDWARDS: Q. No. I'm sorry, let me ask it so it's clear. Data description elements, such as category and character and numeric business data, that would be the same as a text or keyword search that we talked about before; correct? A. Correct. *Id.* at 190–191 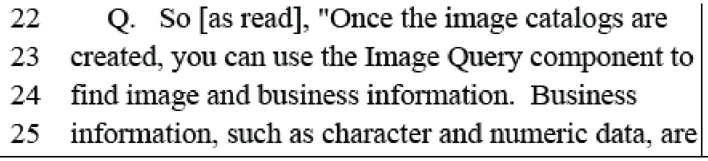 Q. So [as read], "Once the image catalogs are created, you can use the Image Query component to find image and business information. Business information, such as character and numeric data, are

		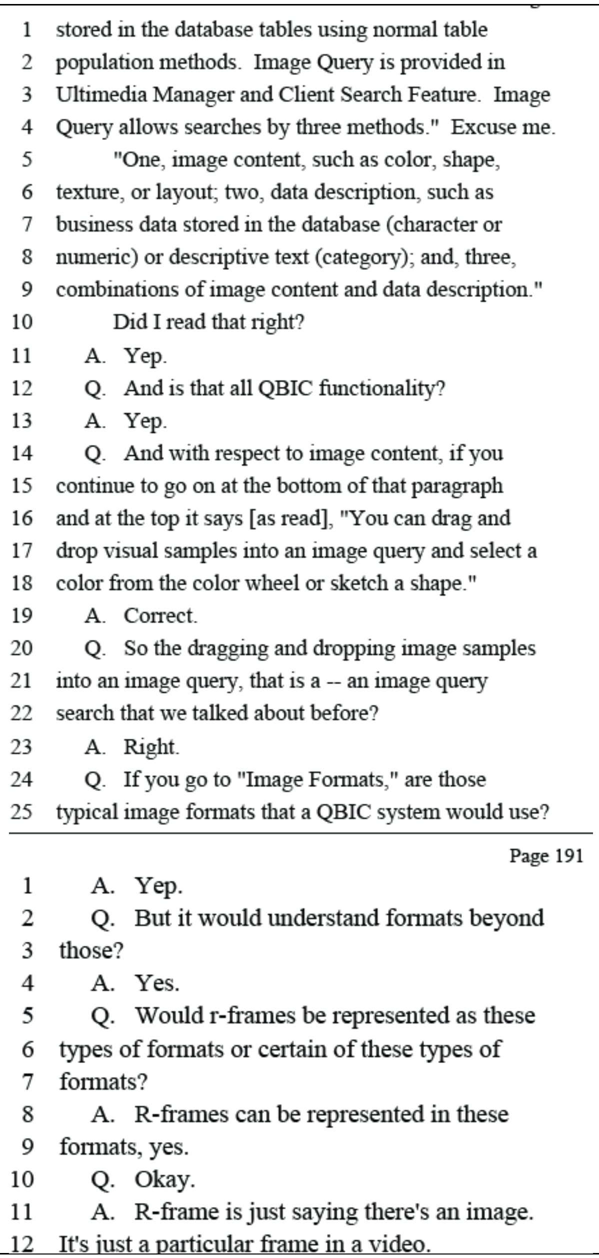 stored in the database tables using normal table population methods. Image Query is provided in Ultimedia Manager and Client Search Feature. Image Query allows searches by three methods." Excuse me. "One, image content, such as color, shape, texture, or layout; two, data description, such as business data stored in the database (character or numeric) or descriptive text (category); and, three, combinations of image content and data description." Did I read that right? A. Yep. Q. And is that all QBIC functionality? A. Yep. Q. And with respect to image content, if you continue to go on at the bottom of that paragraph and at the top it says [as read], "You can drag and drop visual samples into an image query and select a color from the color wheel or sketch a shape." A. Correct. Q. So the dragging and dropping image samples into an image query, that is a -- an image query search that we talked about before? A. Right. Q. If you go to "Image Formats," are those typical image formats that a QBIC system would use? Page 191 A. Yep. Q. But it would understand formats beyond those? A. Yes. Q. Would r-frames be represented as these types of formats or certain of these types of formats? A. R-frames can be represented in these formats, yes. Q. Okay. A. R-frame is just saying there's an image. It's just a particular frame in a video.

Id. at 197–198

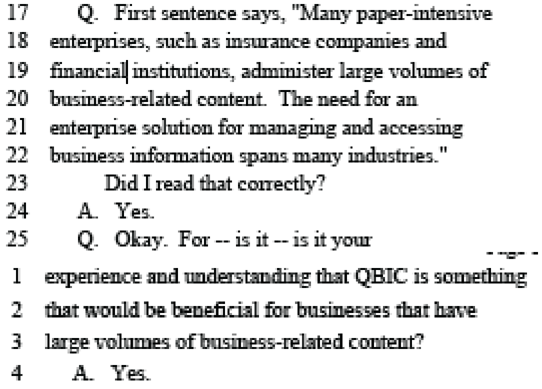

Q. First sentence says, "Many paper-intensive enterprises, such as insurance companies and financial institutions, administer large volumes of business-related content. The need for an enterprise solution for managing and accessing business information spans many industries."
Did I read that correctly?
A. Yes.
Q. Okay. For -- is it -- is it your experience and understanding that QBIC is something that would be beneficial for businesses that have large volumes of business-related content?
A. Yes.

Id. at 307

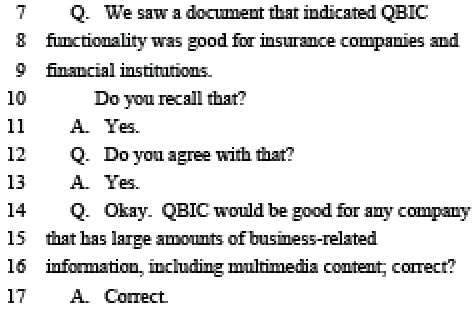

Q. We saw a document that indicated QBIC functionality was good for insurance companies and financial institutions.
Do you recall that?
A. Yes.
Q. Do you agree with that?
A. Yes.
Q. Okay. QBIC would be good for any company that has large amounts of business-related information, including multimedia content; correct?
A. Correct.

To the extent it is found that QBIC does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that QBIC does not anticipate this claim, QBIC renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine QBIC with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each

		of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

3. Claim 4

Claim	’278 Claim Recitation	Exemplary Citations in Reference
4	The method of claim 3, wherein the image includes a representation of the code.	QBIC discloses the method of claim 3, wherein the image includes a representation of the code. IBM0002390 at 2391-92 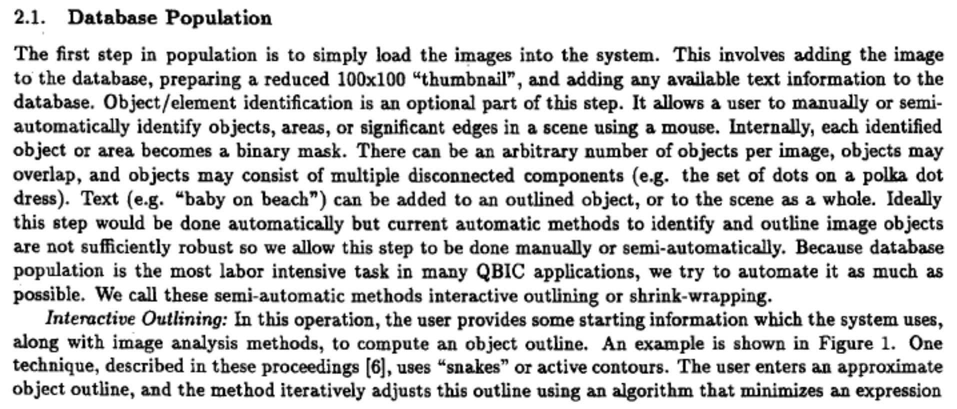 **2.1. Database Population** The first step in population is to simply load the images into the system. This involves adding the image to the database, preparing a reduced 100x100 “thumbnail”, and adding any available text information to the database. Object/element identification is an optional part of this step. It allows a user to manually or semi-automatically identify objects, areas, or significant edges in a scene using a mouse. Internally, each identified object or area becomes a binary mask. There can be an arbitrary number of objects per image, objects may overlap, and objects may consist of multiple disconnected components (e.g. the set of dots on a polka dot dress). Text (e.g. “baby on beach”) can be added to an outlined object, or to the scene as a whole. Ideally this step would be done automatically but current automatic methods to identify and outline image objects are not sufficiently robust so we allow this step to be done manually or semi-automatically. Because database population is the most labor intensive task in many QBIC applications, we try to automate it as much as possible. We call these semi-automatic methods interactive outlining or shrink-wrapping. *Interactive Outlining:* In this operation, the user provides some starting information which the system uses, along with image analysis methods, to compute an object outline. An example is shown in Figure 1. One technique, described in these proceedings [6], uses “snakes” or active contours. The user enters an approximate object outline, and the method iteratively adjusts this outline using an algorithm that minimizes an expression

containing two main terms: total image edge strength along the outline, and the outline curvature. Another method, flood fill, can take as starting information a single mouse-clicked image point. Starting from this object point, all adjacent pixels are included in the object whose color is within a specified color distance from the original point. The perimeter of all included points becomes the object boundary. For reasonably uniform objects that are distinct from the background, this operation provides a fast object outliner.

In interactive outlining, as in other parts of our system, we have attempted to build the functions in such a way as to carefully distribute the tasks between the human and machine. For example, the tasks of giving the starting outline for an object or of attaching semantic meaning to objects is left to the user since they are difficult for machines. On the other hand, computation of quantitative image/object features is ideally suited for machines and so is done automatically.

Figure 1: Original image, approximate user outline, and automatically refined outline of sun.

Query by Image and Video Content: The QBIC System (IBM 0000893–902):

> "INTEGRATION WITH TEXT AND PARAMETRIC ANNOTATION. Query by visual content complements and extends existing query methods. Systems must be able to integrate queries combining date, subject matter, price, and availability with content properties such as color, texture, and shape." *Id.* at p. 8.

Image, Audio, and Video Extenders Administration and Programming (IBM 000001–617):

> "DB2 (DB2) Universal Database (UDB) is a powerful, object-relational database manager. It stores and protects traditional numeric and character data, as well as large, complex objects (LOBs). DB2 extenders help you exploit DB2 's object-relational features. The extenders define distinct data

		types and special functions for image, audio, video, and text objects. … Because of that, the extenders give your applications a single point of access to any or all of these types of data, along with traditional numeric and character data." *Id*. at p. 3. "A trigger defines a set of actions that are activated by a change to a table. Triggers can be used to perform actions such as validating input data, automatically generating a value for a newly inserted row, reading from other tables for cross-referencing purposes, or writing to other tables for auditing purposes. Triggers are often used for integrity checking or to enforce business rules." *Id*. at 16–17. Managing Enterprise Information Portal (IBM 000777–880): "Many paper-intensive enterprises, such as insurance companies and financial institutions, administer large volumes of business-related content. The need for an enterprise solution for managing and accessing business information spans many industries. A content server is a software system that stores multimedia, business forms, documents, and related data, along with metadata that allows employees to process and work with the content." *Id*. at p. 3. "The image search server uses IBM's QBIC (query by image content) technology to help you search for objects by certain visual properties, such as color and texture. The image search server analyzes images and stores the image information in a database. Then users can run image queries, which use the visual properties of images, to match colors, textures, and their positions without describing them in words. You can combine content-based queries with text and keyword searches for more powerful retrieval of image and multimedia data." *Id*. at p. 63. Flickner Depo at 185, 217.

		See claim 3 above. Flickner 8/29/23 Depo at 53 – 56 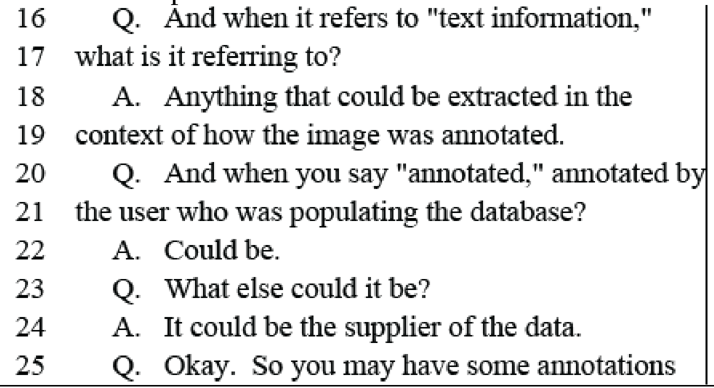 Q. And when it refers to "text information," what is it referring to? A. Anything that could be extracted in the context of how the image was annotated. Q. And when you say "annotated," annotated by the user who was populating the database? A. Could be. Q. What else could it be? A. It could be the supplier of the data. Q. Okay. So you may have some annotations

		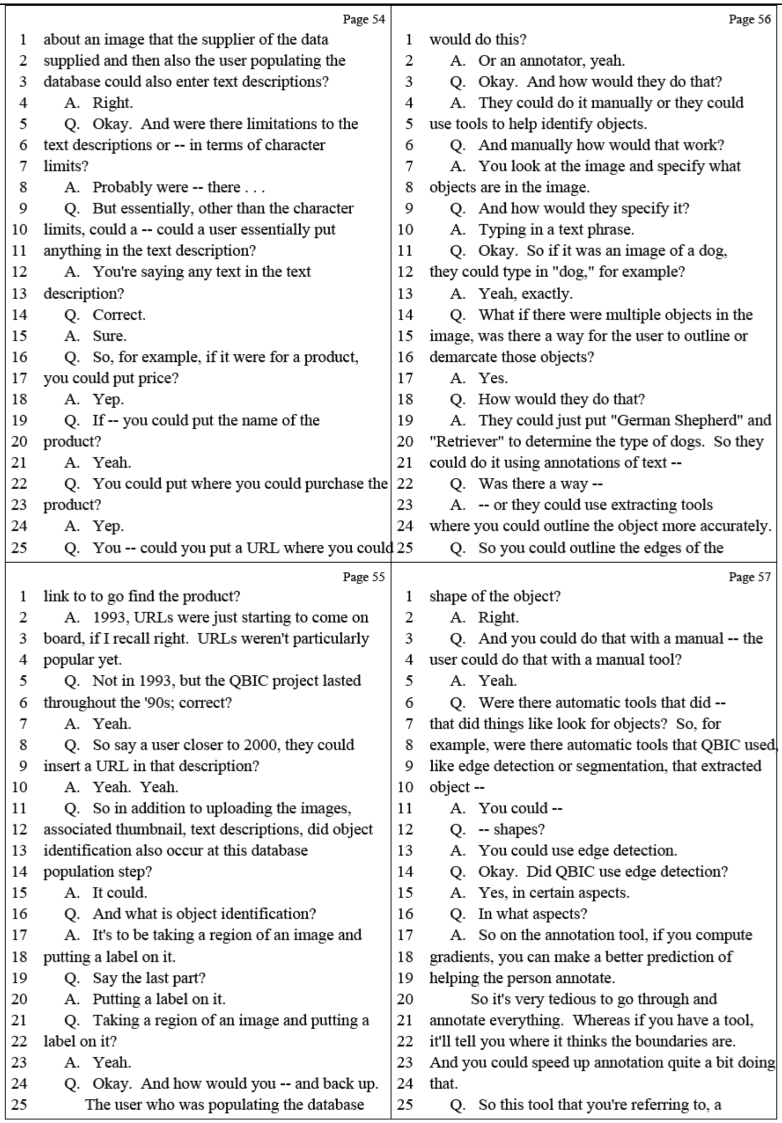

Page 54

about an image that the supplier of the data supplied and then also the user populating the database could also enter text descriptions?
A. Right.
Q. Okay. And were there limitations to the text descriptions or -- in terms of character limits?
A. Probably were -- there . . .
Q. But essentially, other than the character limits, could a -- could a user essentially put anything in the text description?
A. You're saying any text in the text description?
Q. Correct.
A. Sure.
Q. So, for example, if it were for a product, you could put price?
A. Yep.
Q. If -- you could put the name of the product?
A. Yeah.
Q. You could put where you could purchase the product?
A. Yep.
Q. You -- could you put a URL where you could

Page 55

link to to go find the product?
A. 1993, URLs were just starting to come on board, if I recall right. URLs weren't particularly popular yet.
Q. Not in 1993, but the QBIC project lasted throughout the '90s; correct?
A. Yeah.
Q. So say a user closer to 2000, they could insert a URL in that description?
A. Yeah. Yeah.
Q. So in addition to uploading the images, associated thumbnail, text descriptions, did object identification also occur at this database population step?
A. It could.
Q. And what is object identification?
A. It's to be taking a region of an image and putting a label on it.
Q. Say the last part?
A. Putting a label on it.
Q. Taking a region of an image and putting a label on it?
A. Yeah.
Q. Okay. And how would you -- and back up.
The user who was populating the database

Page 56

would do this?
A. Or an annotator, yeah.
Q. Okay. And how would they do that?
A. They could do it manually or they could use tools to help identify objects.
Q. And manually how would that work?
A. You look at the image and specify what objects are in the image.
Q. And how would they specify it?
A. Typing in a text phrase.
Q. Okay. So if it was an image of a dog, they could type in "dog," for example?
A. Yeah, exactly.
Q. What if there were multiple objects in the image, was there a way for the user to outline or demarcate those objects?
A. Yes.
Q. How would they do that?
A. They could just put "German Shepherd" and "Retriever" to determine the type of dogs. So they could do it using annotations of text --
Q. Was there a way --
A. -- or they could use extracting tools where you could outline the object more accurately.
Q. So you could outline the edges of the

Page 57

shape of the object?
A. Right.
Q. And you could do that with a manual -- the user could do that with a manual tool?
A. Yeah.
Q. Were there automatic tools that did -- that did things like look for objects? So, for example, were there automatic tools that QBIC used, like edge detection or segmentation, that extracted object --
A. You could --
Q. -- shapes?
A. You could use edge detection.
Q. Okay. Did QBIC use edge detection?
A. Yes, in certain aspects.
Q. In what aspects?
A. So on the annotation tool, if you compute gradients, you can make a better prediction of helping the person annotate.
So it's very tedious to go through and annotate everything. Whereas if you have a tool, it'll tell you where it thinks the boundaries are. And you could speed up annotation quite a bit doing that.
Q. So this tool that you're referring to, a

		To the extent it is found that QBIC does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that QBIC does not anticipate this claim, QBIC renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine QBIC with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

4. Claim 5

Claim	'278 Claim Recitation	Exemplary Citations in Reference
5	The method of claim 3, further comprising detecting whether there is a match within a specified tolerance and returning a message indicating the absence of a match within a specified tolerance.	QBIC discloses the method of claim 3, further comprising detecting whether there is a match within a specified tolerance and returning a message indicating the absence of a match within a specified tolerance. The QBIC Project: Querying Images by Content Using Color, Texture, and Shape (IBM 0002390–404) at 2397-98 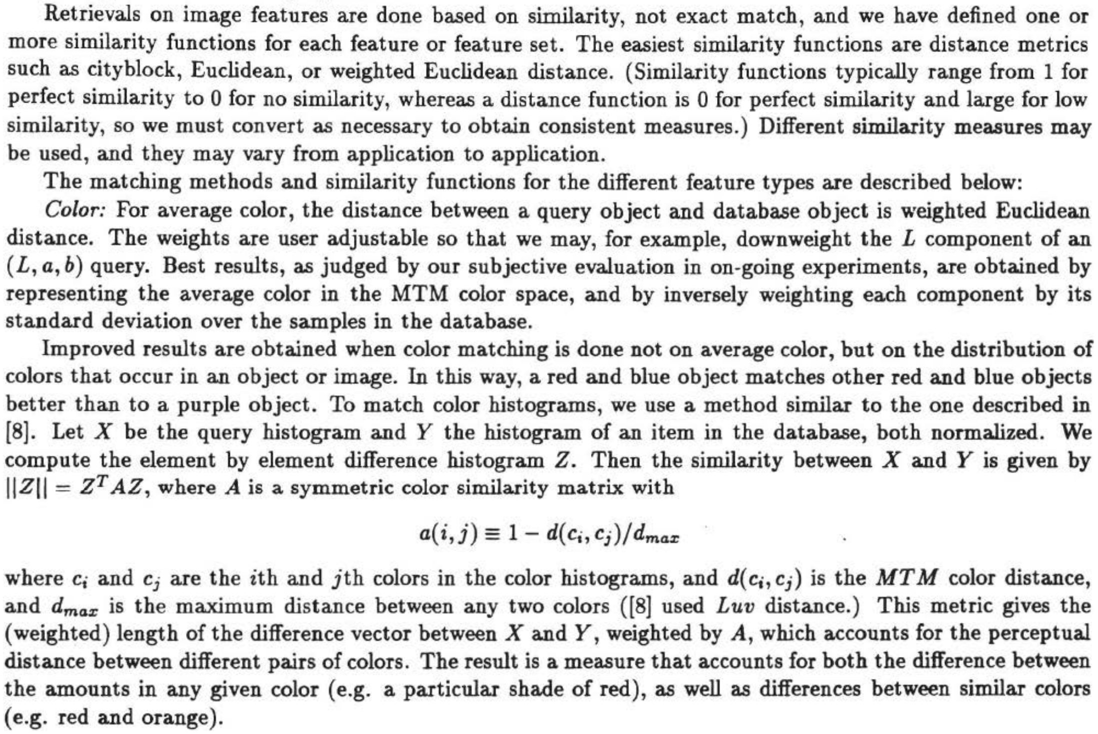 Retrievals on image features are done based on similarity, not exact match, and we have defined one or more similarity functions for each feature or feature set. The easiest similarity functions are distance metrics such as cityblock, Euclidean, or weighted Euclidean distance. (Similarity functions typically range from 1 for perfect similarity to 0 for no similarity, whereas a distance function is 0 for perfect similarity and large for low similarity, so we must convert as necessary to obtain consistent measures.) Different similarity measures may be used, and they may vary from application to application. The matching methods and similarity functions for the different feature types are described below: *Color:* For average color, the distance between a query object and database object is weighted Euclidean distance. The weights are user adjustable so that we may, for example, downweight the L component of an (L, a, b) query. Best results, as judged by our subjective evaluation in on-going experiments, are obtained by representing the average color in the MTM color space, and by inversely weighting each component by its standard deviation over the samples in the database. Improved results are obtained when color matching is done not on average color, but on the distribution of colors that occur in an object or image. In this way, a red and blue object matches other red and blue objects better than to a purple object. To match color histograms, we use a method similar to the one described in [8]. Let X be the query histogram and Y the histogram of an item in the database, both normalized. We compute the element by element difference histogram Z. Then the similarity between X and Y is given by $\|Z\| = Z^T A Z$, where A is a symmetric color similarity matrix with $a(i,j) \equiv 1 - d(c_i, c_j)/d_{max}$ where c_i and c_j are the ith and jth colors in the color histograms, and $d(c_i, c_j)$ is the MTM color distance, and d_{max} is the maximum distance between any two colors ([8] used Luv distance.) This metric gives the (weighted) length of the difference vector between X and Y, weighted by A, which accounts for the perceptual distance between different pairs of colors. The result is a measure that accounts for both the difference between the amounts in any given color (e.g. a particular shade of red), as well as differences between similar colors (e.g. red and orange). Query by Image and Video Content: The QBIC System (IBM 0000893–902): "QBIC has two main components: database population (the process of creating an image database) and database query. During the population, images and

		videos are processed to extract features describing their content-colors, textures, shapes, and camera and object motion-and the features are stored in a database. During the query, the user composes a query graphically. Features are generated from the graphical query and then input to a matching engine that finds images or videos from the database with similar features. Figure 2 shows the system architecture." *Id.* at p. 3. "INTEGRATION WITH TEXT AND PARAMETRIC ANNOTATION. Query by visual content complements and extends existing query methods. Systems must be able to integrate queries combining date, subject matter, price, and availability with content properties such as color, texture, and shape." *Id.* at p. 8. Image, Audio, and Video Extenders Administration and Programming (IBM 000001–617): "To make an image available for searching by content, you catalog the image. When you catalog an image, the Image Extender analyzes the image, by computing the feature values for the image, and stores the values in a QBIC catalog. When you search for an image by content, your query identifies one or more features for the search (such as average color), a source for each feature (such as an example image), and a target set of cataloged images. The Image Extender computes the feature value of the source and compares it to the cataloged feature values for the target images. It then computes a score that indicates how similar the feature values of the target images are to the source. You can have the Image Extender return the images whose features are most similar to the source. The Image Extender will return the handle of each image and the image score. You can also have the Image Extender return only the score of a single image." *Id.* at pp. 20–21. Managing Enterprise Information Portal (IBM 000777–880): "Many paper-intensive enterprises, such as insurance companies and financial

institutions, administer large volumes of business-related content. The need for an enterprise solution for managing and accessing business information spans many industries. A content server is a software system that stores multimedia, business forms, documents, and related data, along with metadata that allows employees to process and work with the content." *Id.* at p. 3.

"The image search server uses IBM's QBIC (query by image content) technology to help you search for objects by certain visual properties, such as color and texture. The image search server analyzes images and stores the image information in a database. Then users can run image queries, which use the visual properties of images, to match colors, textures, and their positions without describing them in words. You can combine content-based queries with text and keyword searches for more powerful retrieval of image and multimedia data." *Id.* at p. 63.

QBIC Demo (IBM000747)

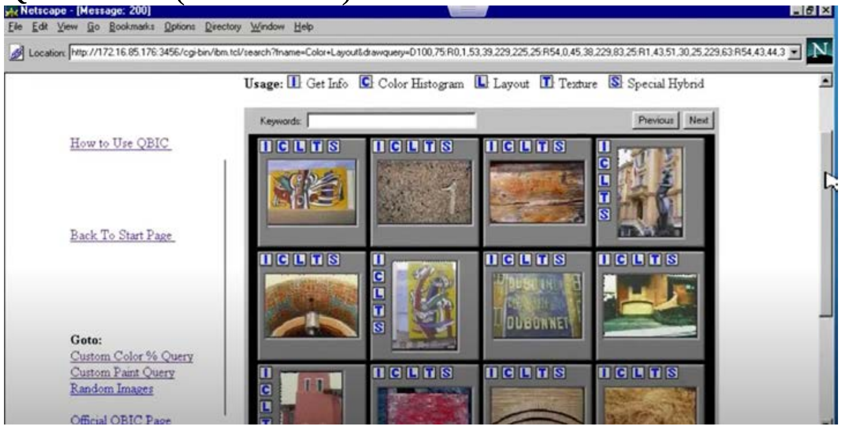

Location: http://172.16.85.176:3456/cgi-bin/ibm.tcl/search?fname=Color+Layout&drawquery=D100,75:R0,1,53,39,229,225,25:R54,0,45,38,229,83,25:R1,43,51,30,25,229,63:R54,43,44,3

Flickner Depo at 28-29

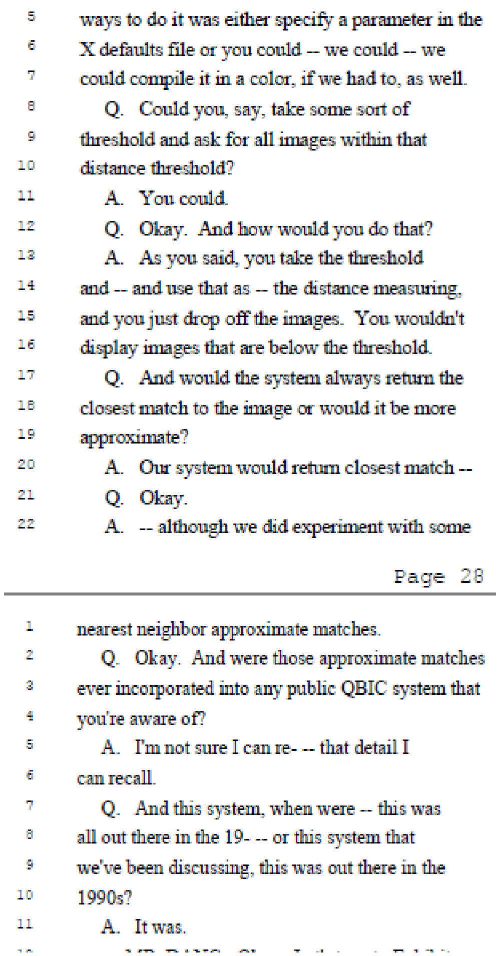

Q. Okay. And how would you alter it then?

A. I don't remember exactly how we -- we ended up doing it in the actual system. But the ways to do it was either specify a parameter in the X defaults file or you could -- we could -- we could compile it in a color, if we had to, as well.

Q. Could you, say, take some sort of threshold and ask for all images within that distance threshold?

A. You could.

Q. Okay. And how would you do that?

A. As you said, you take the threshold and -- and use that as -- the distance measuring, and you just drop off the images. You wouldn't display images that are below the threshold.

Q. And would the system always return the closest match to the image or would it be more approximate?

A. Our system would return closest match --

Q. Okay.

A. -- although we did experiment with some

Page 28

nearest neighbor approximate matches.

Q. Okay. And were those approximate matches ever incorporated into any public QBIC system that you're aware of?

A. I'm not sure I can re- -- that detail I can recall.

Q. And this system, when were -- this was all out there in the 19- -- or this system that we've been discussing, this was out there in the 1990s?

A. It was.

		Flickner Depo at 41-42 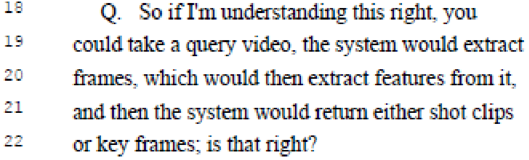 Q. So if I'm understanding this right, you could take a query video, the system would extract frames, which would then extract features from it, and then the system would return either shot clips or key frames; is that right? 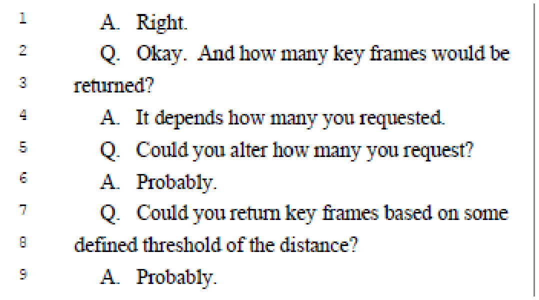 A. Right. Q. Okay. And how many key frames would be returned? A. It depends how many you requested. Q. Could you alter how many you request? A. Probably. Q. Could you return key frames based on some defined threshold of the distance? A. Probably. Flickner 8/29/23 Depo at 119–128

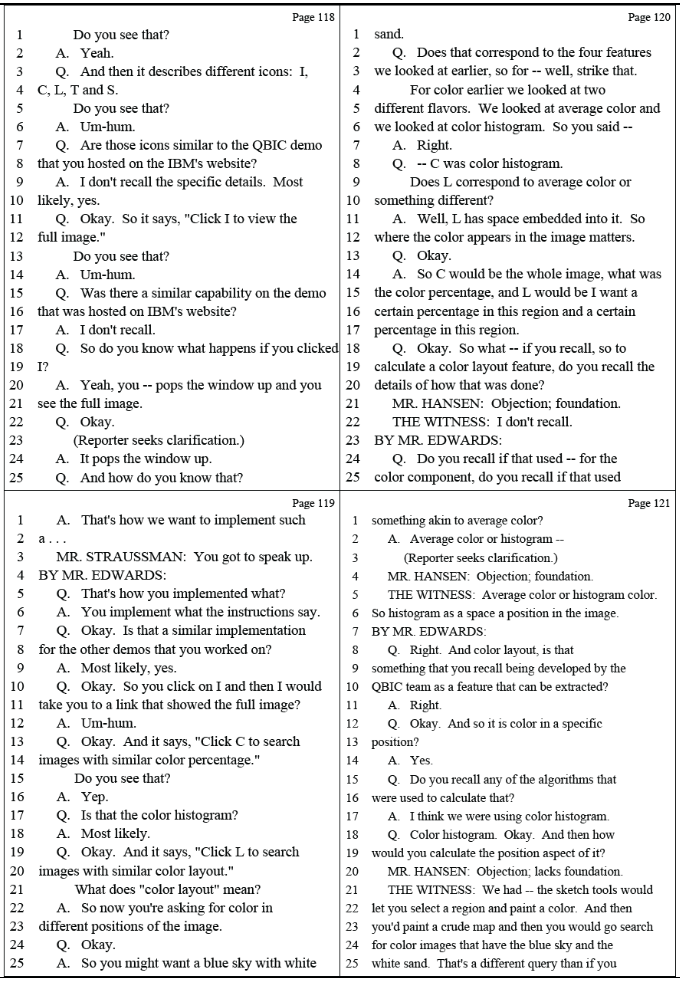

Page 118

Do you see that?
A. Yeah.
Q. And then it describes different icons: I, C, L, T and S.
Do you see that?
A. Um-hum.
Q. Are those icons similar to the QBIC demo that you hosted on the IBM's website?
A. I don't recall the specific details. Most likely, yes.
Q. Okay. So it says, "Click I to view the full image."
Do you see that?
A. Um-hum.
Q. Was there a similar capability on the demo that was hosted on IBM's website?
A. I don't recall.
Q. So do you know what happens if you clicked I?
A. Yeah, you -- pops the window up and you see the full image.
Q. Okay.
(Reporter seeks clarification.)
A. It pops the window up.
Q. And how do you know that?

Page 119

A. That's how we want to implement such a . . .
MR. STRAUSSMAN: You got to speak up.
BY MR. EDWARDS:
Q. That's how you implemented what?
A. You implement what the instructions say.
Q. Okay. Is that a similar implementation for the other demos that you worked on?
A. Most likely, yes.
Q. Okay. So you click on I and then I would take you to a link that showed the full image?
A. Um-hum.
Q. Okay. And it says, "Click C to search images with similar color percentage."
Do you see that?
A. Yep.
Q. Is that the color histogram?
A. Most likely.
Q. Okay. And it says, "Click L to search images with similar color layout."
What does "color layout" mean?
A. So now you're asking for color in different positions of the image.
Q. Okay.
A. So you might want a blue sky with white

Page 120

sand.
Q. Does that correspond to the four features we looked at earlier, so for -- well, strike that.
For color earlier we looked at two different flavors. We looked at average color and we looked at color histogram. So you said --
A. Right.
Q. -- C was color histogram.
Does L correspond to average color or something different?
A. Well, L has space embedded into it. So where the color appears in the image matters.
Q. Okay.
A. So C would be the whole image, what was the color percentage, and L would be I want a certain percentage in this region and a certain percentage in this region.
Q. Okay. So what -- if you recall, so to calculate a color layout feature, do you recall the details of how that was done?
MR. HANSEN: Objection; foundation.
THE WITNESS: I don't recall.
BY MR. EDWARDS:
Q. Do you recall if that used -- for the color component, do you recall if that used

Page 121

something akin to average color?
A. Average color or histogram --
(Reporter seeks clarification.)
MR. HANSEN: Objection; foundation.
THE WITNESS: Average color or histogram color. So histogram as a space a position in the image.
BY MR. EDWARDS:
Q. Right. And color layout, is that something that you recall being developed by the QBIC team as a feature that can be extracted?
A. Right.
Q. Okay. And so it is color in a specific position?
A. Yes.
Q. Do you recall any of the algorithms that were used to calculate that?
A. I think we were using color histogram.
Q. Color histogram. Okay. And then how would you calculate the position aspect of it?
MR. HANSEN: Objection; lacks foundation.
THE WITNESS: We had -- the sketch tools would let you select a region and paint a color. And then you'd paint a crude map and then you would go search for color images that have the blue sky and the white sand. That's a different query than if you

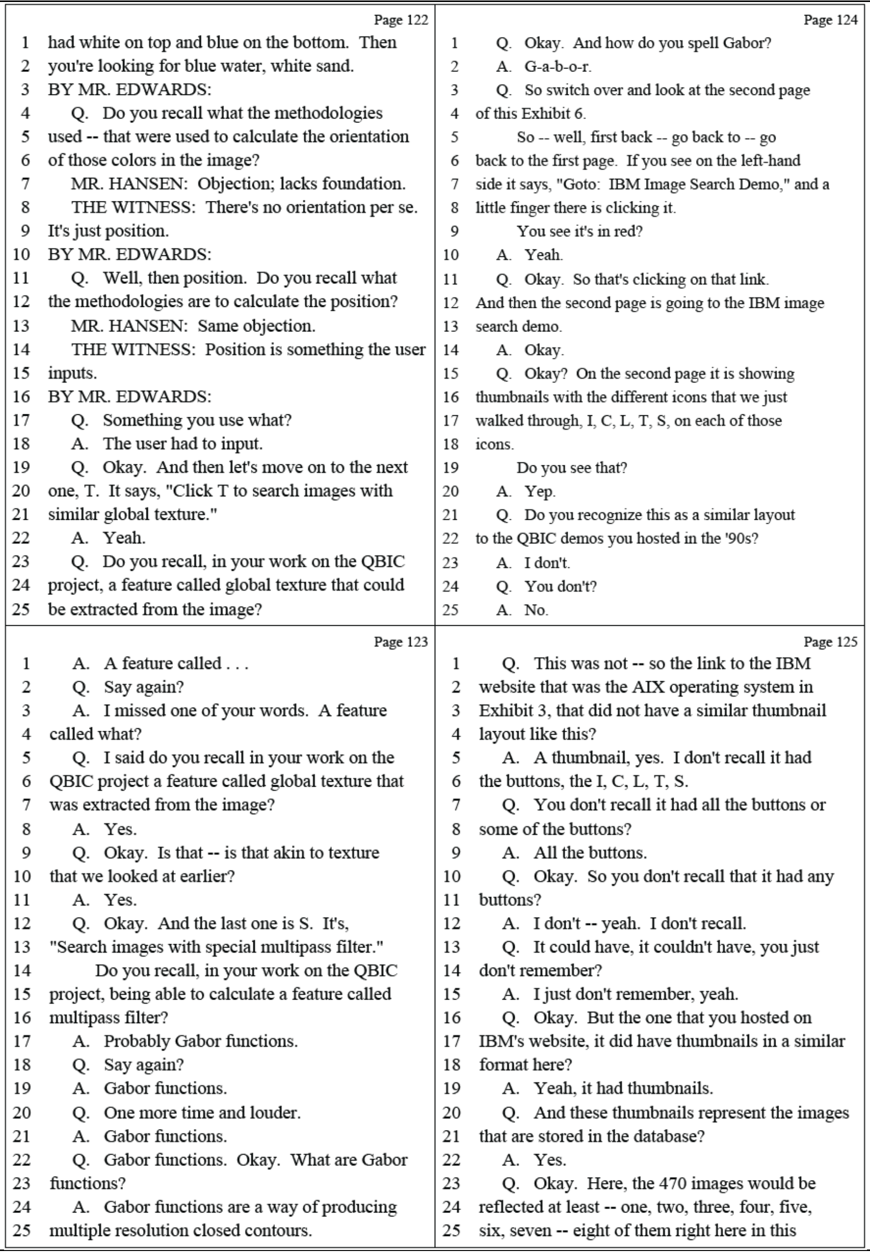

Page 122

had white on top and blue on the bottom. Then you're looking for blue water, white sand.
BY MR. EDWARDS:
Q. Do you recall what the methodologies used -- that were used to calculate the orientation of those colors in the image?
MR. HANSEN: Objection; lacks foundation.
THE WITNESS: There's no orientation per se. It's just position.
BY MR. EDWARDS:
Q. Well, then position. Do you recall what the methodologies are to calculate the position?
MR. HANSEN: Same objection.
THE WITNESS: Position is something the user inputs.
BY MR. EDWARDS:
Q. Something you use what?
A. The user had to input.
Q. Okay. And then let's move on to the next one, T. It says, "Click T to search images with similar global texture."
A. Yeah.
Q. Do you recall, in your work on the QBIC project, a feature called global texture that could be extracted from the image?

Page 123

A. A feature called . . .
Q. Say again?
A. I missed one of your words. A feature called what?
Q. I said do you recall in your work on the QBIC project a feature called global texture that was extracted from the image?
A. Yes.
Q. Okay. Is that -- is that akin to texture that we looked at earlier?
A. Yes.
Q. Okay. And the last one is S. It's, "Search images with special multipass filter."
Do you recall, in your work on the QBIC project, being able to calculate a feature called multipass filter?
A. Probably Gabor functions.
Q. Say again?
A. Gabor functions.
Q. One more time and louder.
A. Gabor functions.
Q. Gabor functions. Okay. What are Gabor functions?
A. Gabor functions are a way of producing multiple resolution closed contours.

Page 124

Q. Okay. And how do you spell Gabor?
A. G-a-b-o-r.
Q. So switch over and look at the second page of this Exhibit 6.
So -- well, first back -- go back to -- go back to the first page. If you see on the left-hand side it says, "Goto: IBM Image Search Demo," and a little finger there is clicking it.
You see it's in red?
A. Yeah.
Q. Okay. So that's clicking on that link. And then the second page is going to the IBM image search demo.
A. Okay.
Q. Okay? On the second page it is showing thumbnails with the different icons that we just walked through, I, C, L, T, S, on each of those icons.
Do you see that?
A. Yep.
Q. Do you recognize this as a similar layout to the QBIC demos you hosted in the '90s?
A. I don't.
Q. You don't?
A. No.

Page 125

Q. This was not -- so the link to the IBM website that was the AIX operating system in Exhibit 3, that did not have a similar thumbnail layout like this?
A. A thumbnail, yes. I don't recall it had the buttons, the I, C, L, T, S.
Q. You don't recall it had all the buttons or some of the buttons?
A. All the buttons.
Q. Okay. So you don't recall that it had any buttons?
A. I don't -- yeah. I don't recall.
Q. It could have, it couldn't have, you just don't remember?
A. I just don't remember, yeah.
Q. Okay. But the one that you hosted on IBM's website, it did have thumbnails in a similar format here?
A. Yeah, it had thumbnails.
Q. And these thumbnails represent the images that are stored in the database?
A. Yes.
Q. Okay. Here, the 470 images would be reflected at least -- one, two, three, four, five, six, seven -- eight of them right here in this

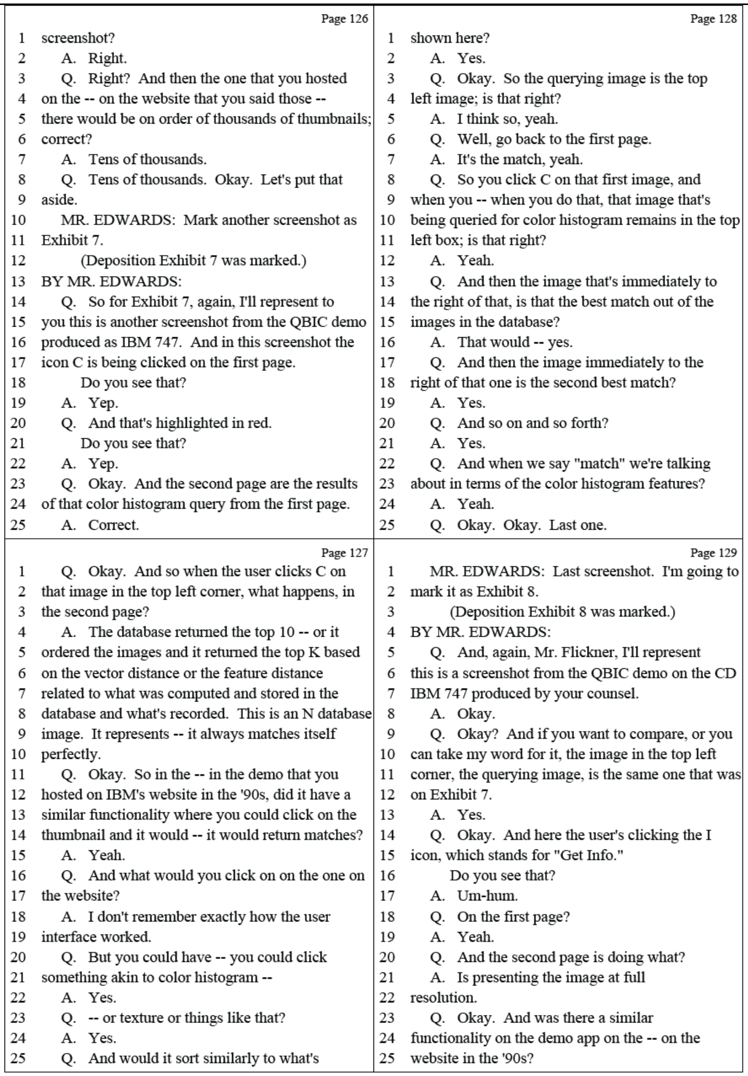

Page 126

screenshot?

A. Right.

Q. Right? And then the one that you hosted on the -- on the website that you said those -- there would be on order of thousands of thumbnails; correct?

A. Tens of thousands.

Q. Tens of thousands. Okay. Let's put that aside.

MR. EDWARDS: Mark another screenshot as Exhibit 7.

(Deposition Exhibit 7 was marked.)

BY MR. EDWARDS:

Q. So for Exhibit 7, again, I'll represent to you this is another screenshot from the QBIC demo produced as IBM 747. And in this screenshot the icon C is being clicked on the first page.

Do you see that?

A. Yep.

Q. And that's highlighted in red.

Do you see that?

A. Yep.

Q. Okay. And the second page are the results of that color histogram query from the first page.

A. Correct.

Page 127

Q. Okay. And so when the user clicks C on that image in the top left corner, what happens, in the second page?

A. The database returned the top 10 -- or it ordered the images and it returned the top K based on the vector distance or the feature distance related to what was computed and stored in the database and what's recorded. This is an N database image. It represents -- it always matches itself perfectly.

Q. Okay. So in the -- in the demo that you hosted on IBM's website in the '90s, did it have a similar functionality where you could click on the thumbnail and it would -- it would return matches?

A. Yeah.

Q. And what would you click on on the one on the website?

A. I don't remember exactly how the user interface worked.

Q. But you could have -- you could click something akin to color histogram --

A. Yes.

Q. -- or texture or things like that?

A. Yes.

Q. And would it sort similarly to what's

Page 128

shown here?

A. Yes.

Q. Okay. So the querying image is the top left image; is that right?

A. I think so, yeah.

Q. Well, go back to the first page.

A. It's the match, yeah.

Q. So you click C on that first image, and when you -- when you do that, that image that's being queried for color histogram remains in the top left box; is that right?

A. Yeah.

Q. And then the image that's immediately to the right of that, is that the best match out of the images in the database?

A. That would -- yes.

Q. And then the image immediately to the right of that one is the second best match?

A. Yes.

Q. And so on and so forth?

A. Yes.

Q. And when we say "match" we're talking about in terms of the color histogram features?

A. Yeah.

Q. Okay. Okay. Last one.

Page 129

MR. EDWARDS: Last screenshot. I'm going to mark it as Exhibit 8.

(Deposition Exhibit 8 was marked.)

BY MR. EDWARDS:

Q. And, again, Mr. Flickner, I'll represent this is a screenshot from the QBIC demo on the CD IBM 747 produced by your counsel.

A. Okay.

Q. Okay? And if you want to compare, or you can take my word for it, the image in the top left corner, the querying image, is the same one that was on Exhibit 7.

A. Yes.

Q. Okay. And here the user's clicking the I icon, which stands for "Get Info."

Do you see that?

A. Um-hum.

Q. On the first page?

A. Yeah.

Q. And the second page is doing what?

A. Is presenting the image at full resolution.

Q. Okay. And was there a similar functionality on the demo app on the -- on the website in the '90s?

Id. at 135 – 136

		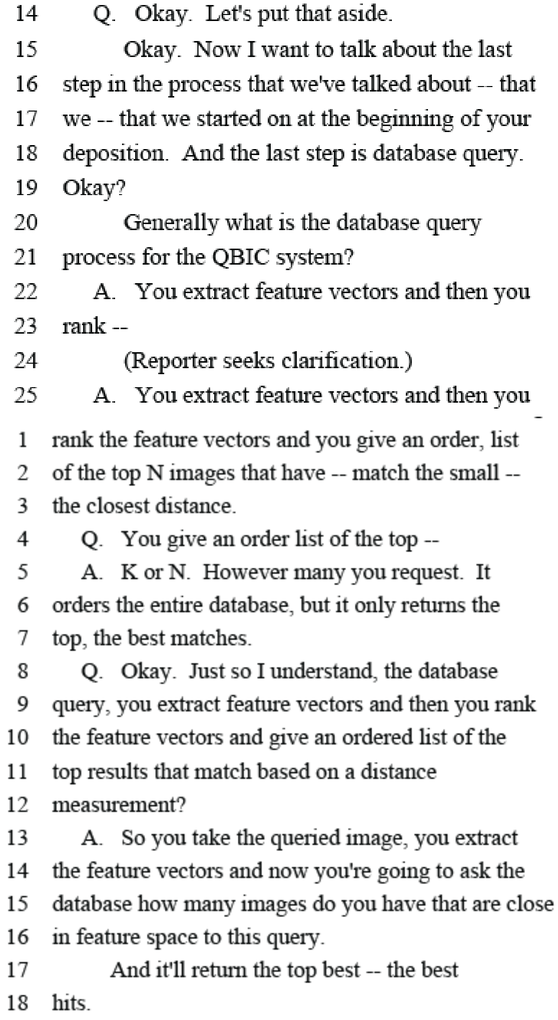 Q. Okay. Let's put that aside. Okay. Now I want to talk about the last step in the process that we've talked about -- that we -- that we started on at the beginning of your deposition. And the last step is database query. Okay? Generally what is the database query process for the QBIC system? A. You extract feature vectors and then you rank -- (Reporter seeks clarification.) A. You extract feature vectors and then you rank the feature vectors and you give an order, list of the top N images that have -- match the small -- the closest distance. Q. You give an order list of the top -- A. K or N. However many you request. It orders the entire database, but it only returns the top, the best matches. Q. Okay. Just so I understand, the database query, you extract feature vectors and then you rank the feature vectors and give an ordered list of the top results that match based on a distance measurement? A. So you take the queried image, you extract the feature vectors and now you're going to ask the database how many images do you have that are close in feature space to this query. And it'll return the top best -- the best hits. *Id.* at 144

		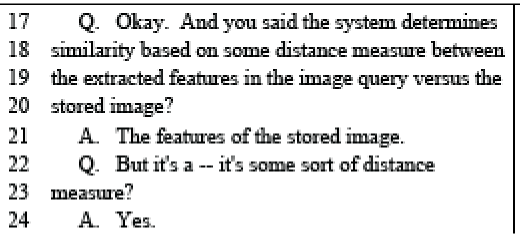 Q. Okay. And you said the system determines similarity based on some distance measure between the extracted features in the image query versus the stored image? A. The features of the stored image. Q. But it's a -- it's some sort of distance measure? A. Yes. *Id.* at 317–322 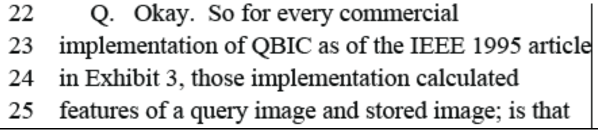 Q. Okay. So for every commercial implementation of QBIC as of the IEEE 1995 article in Exhibit 3, those implementation calculated features of a query image and stored image; is that

		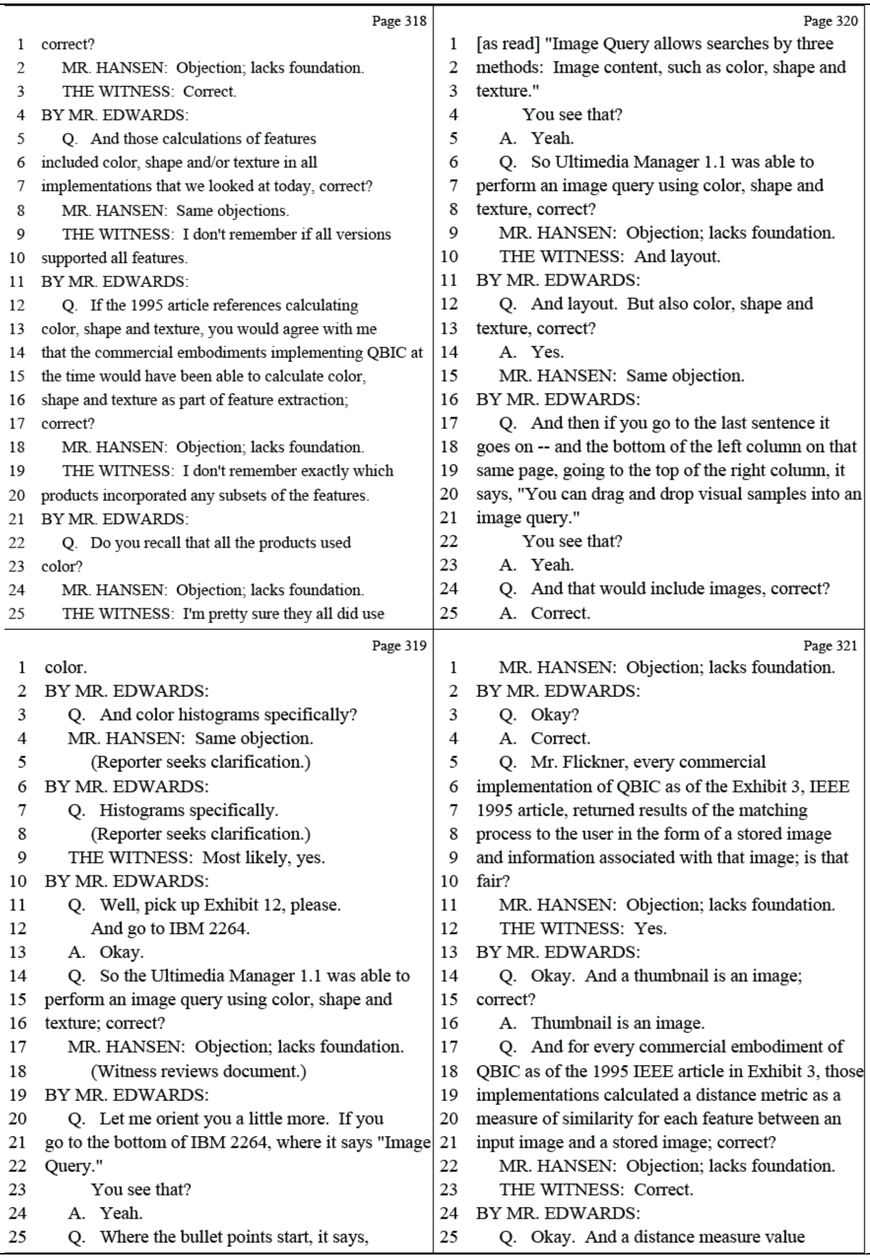

Page 318

correct?
MR. HANSEN: Objection; lacks foundation.
THE WITNESS: Correct.
BY MR. EDWARDS:
Q. And those calculations of features included color, shape and/or texture in all implementations that we looked at today, correct?
MR. HANSEN: Same objections.
THE WITNESS: I don't remember if all versions supported all features.
BY MR. EDWARDS:
Q. If the 1995 article references calculating color, shape and texture, you would agree with me that the commercial embodiments implementing QBIC at the time would have been able to calculate color, shape and texture as part of feature extraction; correct?
MR. HANSEN: Objection; lacks foundation.
THE WITNESS: I don't remember exactly which products incorporated any subsets of the features.
BY MR. EDWARDS:
Q. Do you recall that all the products used color?
MR. HANSEN: Objection; lacks foundation.
THE WITNESS: I'm pretty sure they all did use

Page 319

color.
BY MR. EDWARDS:
Q. And color histograms specifically?
MR. HANSEN: Same objection.
(Reporter seeks clarification.)
BY MR. EDWARDS:
Q. Histograms specifically.
(Reporter seeks clarification.)
THE WITNESS: Most likely, yes.
BY MR. EDWARDS:
Q. Well, pick up Exhibit 12, please.
And go to IBM 2264.
A. Okay.
Q. So the Ultimedia Manager 1.1 was able to perform an image query using color, shape and texture; correct?
MR. HANSEN: Objection; lacks foundation.
(Witness reviews document.)
BY MR. EDWARDS:
Q. Let me orient you a little more. If you go to the bottom of IBM 2264, where it says "Image Query."
You see that?
A. Yeah.
Q. Where the bullet points start, it says,

Page 320

[as read] "Image Query allows searches by three methods: Image content, such as color, shape and texture."
You see that?
A. Yeah.
Q. So Ultimedia Manager 1.1 was able to perform an image query using color, shape and texture, correct?
MR. HANSEN: Objection; lacks foundation.
THE WITNESS: And layout.
BY MR. EDWARDS:
Q. And layout. But also color, shape and texture, correct?
A. Yes.
MR. HANSEN: Same objection.
BY MR. EDWARDS:
Q. And then if you go to the last sentence it goes on -- and the bottom of the left column on that same page, going to the top of the right column, it says, "You can drag and drop visual samples into an image query."
You see that?
A. Yeah.
Q. And that would include images, correct?
A. Correct.

Page 321

MR. HANSEN: Objection; lacks foundation.
BY MR. EDWARDS:
Q. Okay?
A. Correct.
Q. Mr. Flickner, every commercial implementation of QBIC as of the Exhibit 3, IEEE 1995 article, returned results of the matching process to the user in the form of a stored image and information associated with that image; is that fair?
MR. HANSEN: Objection; lacks foundation.
THE WITNESS: Yes.
BY MR. EDWARDS:
Q. Okay. And a thumbnail is an image; correct?
A. Thumbnail is an image.
Q. And for every commercial embodiment of QBIC as of the 1995 IEEE article in Exhibit 3, those implementations calculated a distance metric as a measure of similarity for each feature between an input image and a stored image; correct?
MR. HANSEN: Objection; lacks foundation.
THE WITNESS: Correct.
BY MR. EDWARDS:
Q. Okay. And a distance measure value

		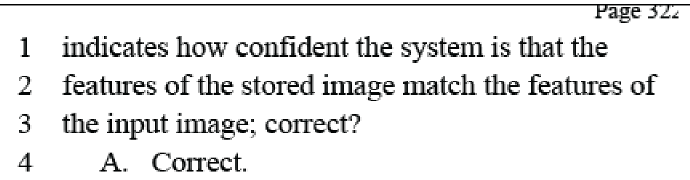 Page 322 indicates how confident the system is that the features of the stored image match the features of the input image; correct? A. Correct. *See also* Flickner Depo at 25-26, 131-132, 162-163 To the extent it is found that QBIC does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that QBIC does not anticipate this claim, QBIC renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine QBIC with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

Exhibit F-31

Bank of America Corporation ("BofA" or "Defendant") provides this chart subject to all reservations, objections, statements, and disclaimers set forth herein and in BofA's Final Invalidity Contentions Cover Pleading, as well as any amendment, supplement, or modification thereof, which are incorporated herein by reference in their entirety.

On information and belief, CoolTown was made, known or used by others in the United States no later than July 10, 2000. *See generally* T. Kindberg et al., A Web-Based Nomadic Computing System; J. Barton et al., The Cooltown User Experience; cooltown developer's network, coolbase overview; cooltown developer's network, people, places, things: web presence for the real world. CoolTown is available as prior art at least under 35 U.S.C. § 102(a). On information and belief, HP was using CoolTown and made it available to the public in the United States by releasing the website on the Internet (and accessible to U.S. users) by at least March 2000, making the software available on the HP website by March 2000, and making available on the Internet (and accessible to U.S. users) the open source code, which reflects the functionality of CoolTown by at least March 2000. *See e.g.*, HP000043–044; HP000045–046; HP000060–061; *see also* https://www.zdnet.com/article/hp-brings-open-source-to-cooltown/. Therefore, HP Cooltown was at least known, used, and/or sold to the public by at least 2000. On information and belief, CoolTown is also available as prior art under 102(g). BofA is diligently investigating CoolTown and will supplement with additional information, if necessary.

As shown in the chart below, CoolTown anticipates asserted claims 1, 4, and 5 of the '278 Patent. To the extent CoolTown is found not to expressly disclose certain limitations in the asserted claims, such limitations are inherent. To the extent CoolTown is found not to anticipate any asserted claims or claim elements of the '278 Patent, the reference nevertheless renders those claims or claim elements obvious under 35 U.S.C. § 103, either alone or in combination with other art identified herein or in the cover pleading of BofA's Invalidity Contentions, for the reasons identified in the cover pleading. These Final Invalidity Contentions are not an admission by BofA that the accused products, including any current or past versions of the accused products, are covered by, or infringe the asserted claims, but are based instead on the recognition that if the claims are interpreted to be broad enough to encompass the accused products, the claims must also be construed to have that same scope when considering whether they are invalid.

The following chart is based on the Court's claim construction rulings (Dkt. Nos. 145, 153, 236), and to the extent not contradictory, Plaintiffs' apparent interpretations of the claims in Plaintiffs' "Preliminary" Final Infringement Contentions. BofA notes that in many instances, Plaintiffs' "Preliminary" Final Infringement Contentions ignore and/or contradict the Court's claim construction rulings. BofA does not accept the assumptions concerning the scope and meaning implicit in Plaintiffs' Preliminary Final Infringement Contentions, to the extent those assumptions are discernible, and reserves the right to challenge Plaintiffs' proposed (or implied) constructions. Furthermore, Plaintiffs' "Preliminary" Final Contentions fail to comply with the level of specificity required by SPR

2.1.2 and 2.1.3, and include improper reservation of rights to supplement with "additional ways in which BoA Accused Products infringe." Any alleged deficient discovery is of Plaintiffs own making. Most if not all of the deficiencies in Plaintiffs' Preliminary Final Contentions have been documented by Defendants through multiple letters since at least May 18, 2021 and up to Bank of America's most recent letter of August 19, 2023. Plaintiffs' shifting infringement contentions, deficiencies in its infringement contentions, and changing claim scope have prejudiced BofA from being able to reasonably prepare its defenses, including its invalidity contentions. BofA also reserves the right to revise and supplement these charts if and when Plaintiffs are permitted to provide revised or supplemental Infringement Contentions. Citations given in the chart below are merely representative of the respective elements to provide notice to Plaintiffs of BofA's invalidity theories and are not meant to be exhaustive.

1. Claim 1

Element	'278 Claim Recitation	Exemplary Citations in Reference
[1Pre]	A computer-assisted method, comprising:	CoolTown discloses a computer-assisted method. The Cooltown User Experience (HP000088 – 093) at 1: 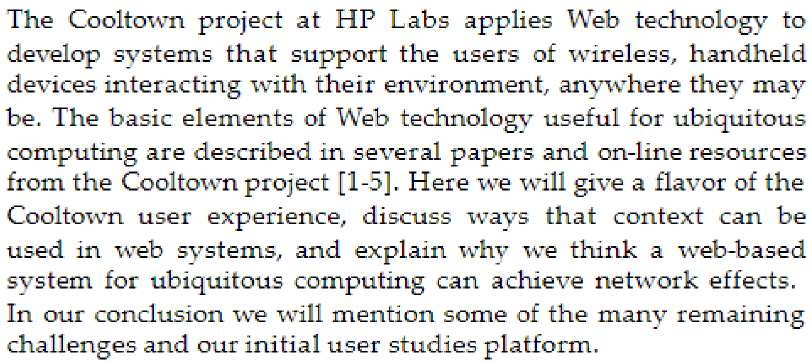 The Cooltown project at HP Labs applies Web technology to develop systems that support the users of wireless, handheld devices interacting with their environment, anywhere they may be. The basic elements of Web technology useful for ubiquitous computing are described in several papers and on-line resources from the Cooltown project [1-5]. Here we will give a flavor of the Cooltown user experience, discuss ways that context can be used in web systems, and explain why we think a web-based system for ubiquitous computing can achieve network effects. In our conclusion we will mention some of the many remaining challenges and our initial user studies platform. A Web-Based Nomadic Computing System (HP000074 – 87) at 1 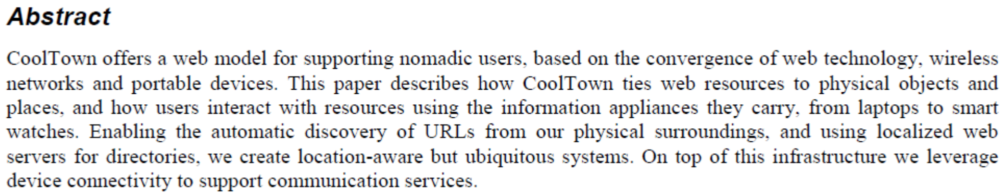 ***Abstract*** CoolTown offers a web model for supporting nomadic users, based on the convergence of web technology, wireless networks and portable devices. This paper describes how CoolTown ties web resources to physical objects and places, and how users interact with resources using the information appliances they carry, from laptops to smart watches. Enabling the automatic discovery of URLs from our physical surroundings, and using localized web servers for directories, we create location-aware but ubiquitous systems. On top of this infrastructure we leverage device connectivity to support communication services. Coolbase[1] Overview at 1

---

[1] *See* https://web.archive.org/web/20011030090650/http://cooltown.hp.com/dev/coolbase-overview.asp

Element	'278 Claim Recitation	Exemplary Citations in Reference
		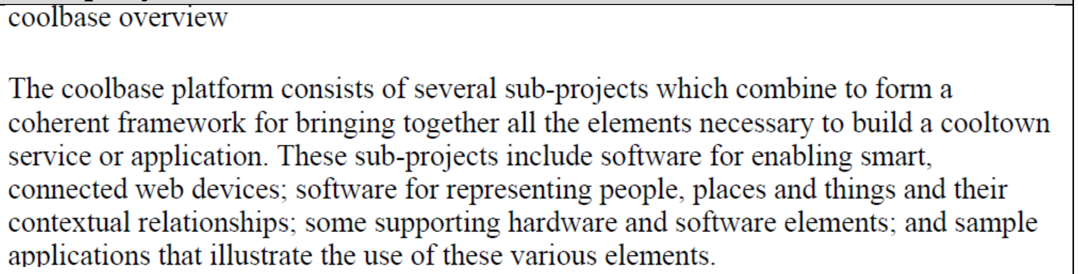 coolbase overview The coolbase platform consists of several sub-projects which combine to form a coherent framework for bringing together all the elements necessary to build a cooltown service or application. These sub-projects include software for enabling smart, connected web devices; software for representing people, places and things and their contextual relationships; some supporting hardware and software elements; and sample applications that illustrate the use of these various elements.
[1A]	receiving, via a mobile device, an image comprising a representation of at least a portion of a document;	CoolTown discloses receiving, via a mobile device, an image comprising a representation of at least a portion of a document. The Cooltown: People, Places, Things Web Presence for the Real World (HP000025 – 34) at 20:

Element	'278 Claim Recitation	Exemplary Citations in Reference
		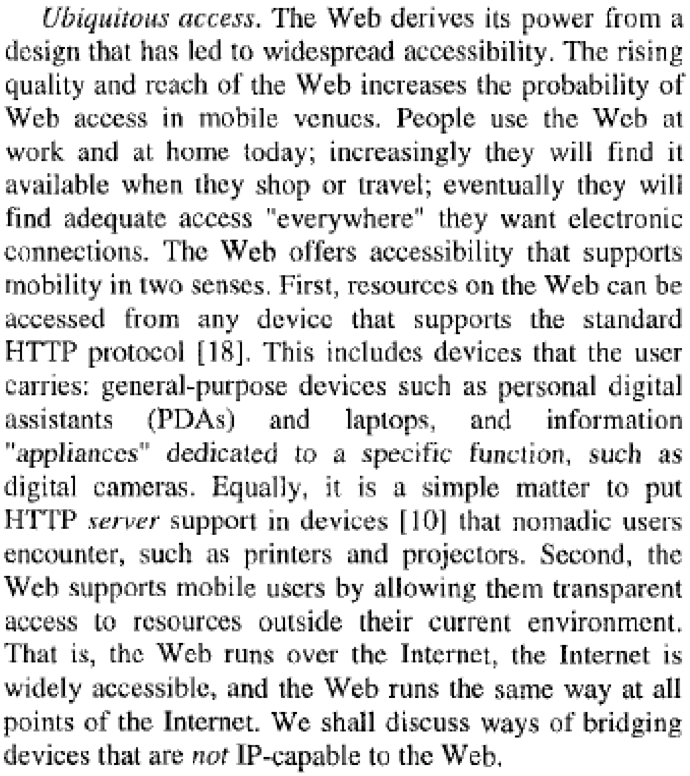 *Ubiquitous access*. The Web derives its power from a design that has led to widespread accessibility. The rising quality and reach of the Web increases the probability of Web access in mobile venues. People use the Web at work and at home today; increasingly they will find it available when they shop or travel; eventually they will find adequate access "everywhere" they want electronic connections. The Web offers accessibility that supports mobility in two senses. First, resources on the Web can be accessed from any device that supports the standard HTTP protocol [18]. This includes devices that the user carries: general-purpose devices such as personal digital assistants (PDAs) and laptops, and information "appliances" dedicated to a specific function, such as digital cameras. Equally, it is a simple matter to put HTTP *server* support in devices [10] that nomadic users encounter, such as printers and projectors. Second, the Web supports mobile users by allowing them transparent access to resources outside their current environment. That is, the Web runs over the Internet, the Internet is widely accessible, and the Web runs the same way at all points of the Internet. We shall discuss ways of bridging devices that are *not* IP-capable to the Web. The Cooltown: People, Places, Things Web Presence for the Real World (HP000025 – 34) at 24:

Element	’278 Claim Recitation	Exemplary Citations in Reference
		3.4. Infrastructure for things; ChaiServer and eSquirt Given a URL from one of the lower-layer services, we 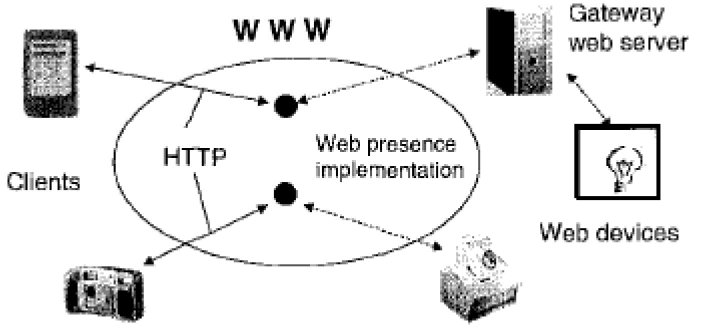 **Figure 5. Devices (printer, light), with and without embedded web servers.** Figure 5 shows interactions between two clients and their respective web devices. It shows a digital camera that interacts with a printer via HTTP operations on the printer's point of web presence. The URL for the printer resolves to its network address. The camera first interrogates the printer to find out the formats it accepts. The printer can supply this as an HTML page or as an XML document that describes it. Then the camera sends the image as, for example, JPEG encoded data. Figure 5 also shows lights connected to a home PC by a controller unit. The home PC acts as a gateway to the lighting unit. A remote user issues HTTP operations via a web browser on a PDA; the programs on the PC that handle HTTP operations directed to the light's URL manipulate the lighting unit accordingly.

Element	'278 Claim Recitation	Exemplary Citations in Reference
		The Cooltown User Experience (HP000088 – 093) at 1: 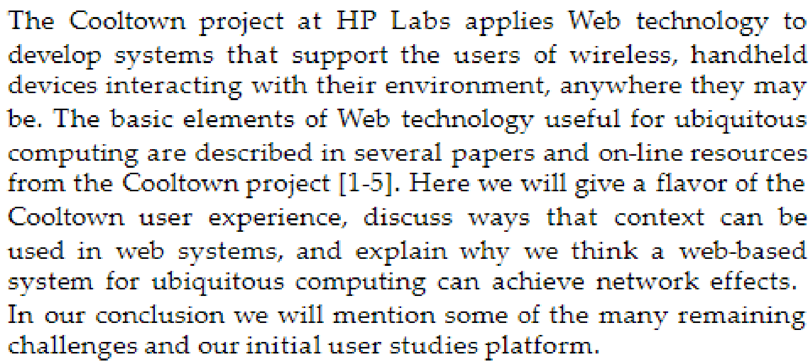 The Cooltown project at HP Labs applies Web technology to develop systems that support the users of wireless, handheld devices interacting with their environment, anywhere they may be. The basic elements of Web technology useful for ubiquitous computing are described in several papers and on-line resources from the Cooltown project [1-5]. Here we will give a flavor of the Cooltown user experience, discuss ways that context can be used in web systems, and explain why we think a web-based system for ubiquitous computing can achieve network effects. In our conclusion we will mention some of the many remaining challenges and our initial user studies platform. A Web-Based Nomadic Computing System (HP000074 – 87) at 1 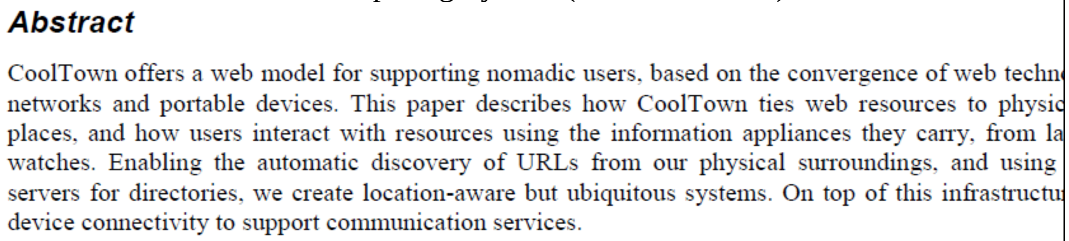 ***Abstract*** CoolTown offers a web model for supporting nomadic users, based on the convergence of web techn networks and portable devices. This paper describes how CoolTown ties web resources to physic places, and how users interact with resources using the information appliances they carry, from la watches. Enabling the automatic discovery of URLs from our physical surroundings, and using servers for directories, we create location-aware but ubiquitous systems. On top of this infrastructu device connectivity to support communication services. A Web-Based Nomadic Computing System (HP000074 – 87) at 2

Element	'278 Claim Recitation	Exemplary Citations in Reference
		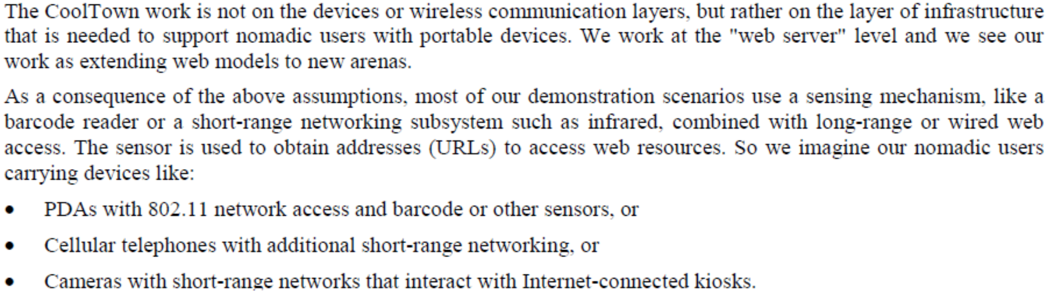 The CoolTown work is not on the devices or wireless communication layers, but rather on the layer of infrastructure that is needed to support nomadic users with portable devices. We work at the "web server" level and we see our work as extending web models to new arenas. As a consequence of the above assumptions, most of our demonstration scenarios use a sensing mechanism, like a barcode reader or a short-range networking subsystem such as infrared, combined with long-range or wired web access. The sensor is used to obtain addresses (URLs) to access web resources. So we imagine our nomadic users carrying devices like: • PDAs with 802.11 network access and barcode or other sensors, or • Cellular telephones with additional short-range networking, or • Cameras with short-range networks that interact with Internet-connected kiosks. A Web-Based Nomadic Computing System (HP000074 – 87) at 3 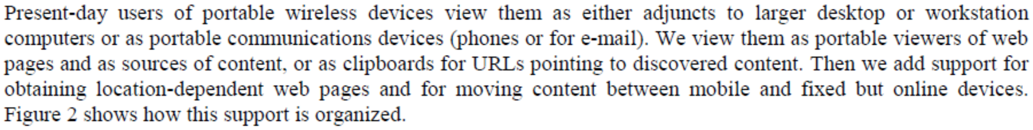 Present-day users of portable wireless devices view them as either adjuncts to larger desktop or workstation computers or as portable communications devices (phones or for e-mail). We view them as portable viewers of web pages and as sources of content, or as clipboards for URLs pointing to discovered content. Then we add support for obtaining location-dependent web pages and for moving content between mobile and fixed but online devices. Figure 2 shows how this support is organized. A Web-Based Nomadic Computing System (HP000074 – 87) at 3

Element	'278 Claim Recitation	Exemplary Citations in Reference
		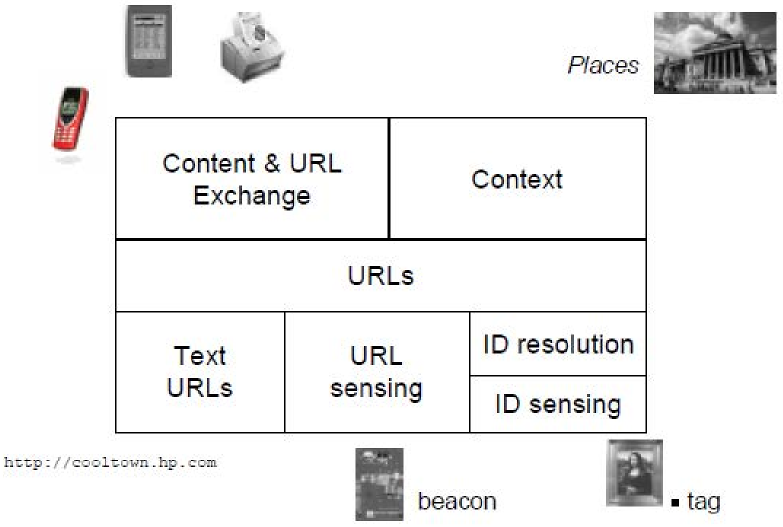 Figure 2. Component relationships. Sources of URLs, including keyboards, beacons, and resolved tags point to context or content. A Web-Based Nomadic Computing System (HP000074 – 87) at 4 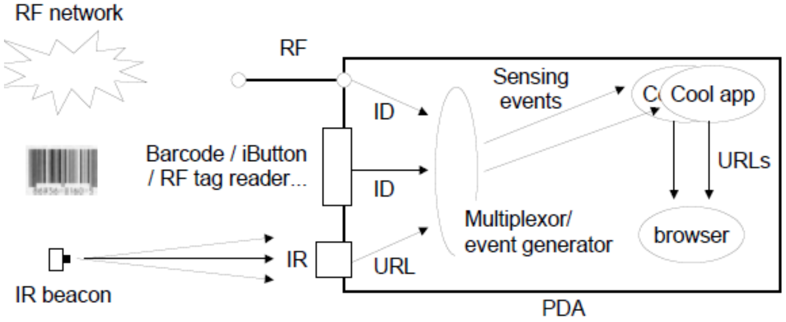 Figure 3. A sketch of a PDA equipped for sensing URLs and identifiers.

Element	'278 Claim Recitation	Exemplary Citations in Reference
		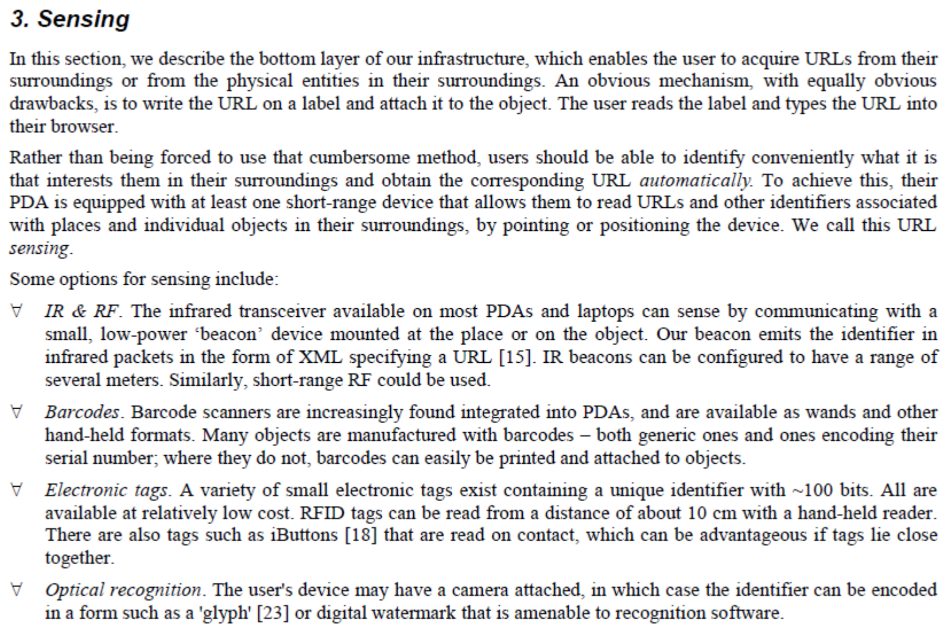 ***3. Sensing*** In this section, we describe the bottom layer of our infrastructure, which enables the user to acquire URLs from their surroundings or from the physical entities in their surroundings. An obvious mechanism, with equally obvious drawbacks, is to write the URL on a label and attach it to the object. The user reads the label and types the URL into their browser. Rather than being forced to use that cumbersome method, users should be able to identify conveniently what it is that interests them in their surroundings and obtain the corresponding URL *automatically*. To achieve this, their PDA is equipped with at least one short-range device that allows them to read URLs and other identifiers associated with places and individual objects in their surroundings, by pointing or positioning the device. We call this URL *sensing*. Some options for sensing include: ∀ *IR & RF*. The infrared transceiver available on most PDAs and laptops can sense by communicating with a small, low-power 'beacon' device mounted at the place or on the object. Our beacon emits the identifier in infrared packets in the form of XML specifying a URL [15]. IR beacons can be configured to have a range of several meters. Similarly, short-range RF could be used. ∀ *Barcodes*. Barcode scanners are increasingly found integrated into PDAs, and are available as wands and other hand-held formats. Many objects are manufactured with barcodes – both generic ones and ones encoding their serial number; where they do not, barcodes can easily be printed and attached to objects. ∀ *Electronic tags*. A variety of small electronic tags exist containing a unique identifier with ~100 bits. All are available at relatively low cost. RFID tags can be read from a distance of about 10 cm with a hand-held reader. There are also tags such as iButtons [18] that are read on contact, which can be advantageous if tags lie close together. ∀ *Optical recognition*. The user's device may have a camera attached, in which case the identifier can be encoded in a form such as a 'glyph' [23] or digital watermark that is amenable to recognition software. A Web-Based Nomadic Computing System (HP000074 – 87) at 4 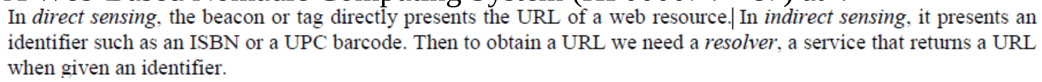 In *direct sensing*, the beacon or tag directly presents the URL of a web resource. In *indirect sensing*, it presents an identifier such as an ISBN or a UPC barcode. Then to obtain a URL we need a *resolver*, a service that returns a URL when given an identifier. A Web-Based Nomadic Computing System (HP000074 – 87) at 8 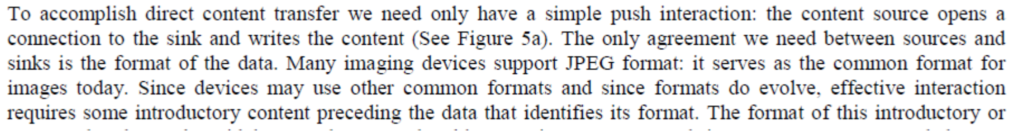 To accomplish direct content transfer we need only have a simple push interaction: the content source opens a connection to the sink and writes the content (See Figure 5a). The only agreement we need between sources and sinks is the format of the data. Many imaging devices support JPEG format: it serves as the common format for images today. Since devices may use other common formats and since formats do evolve, effective interaction requires some introductory content preceding the data that identifies its format. The format of this introductory or The Cooltown: People, Places, Things Web Presence for the Real World (HP000025 – 34) at 20:

Element	'278 Claim Recitation	Exemplary Citations in Reference
		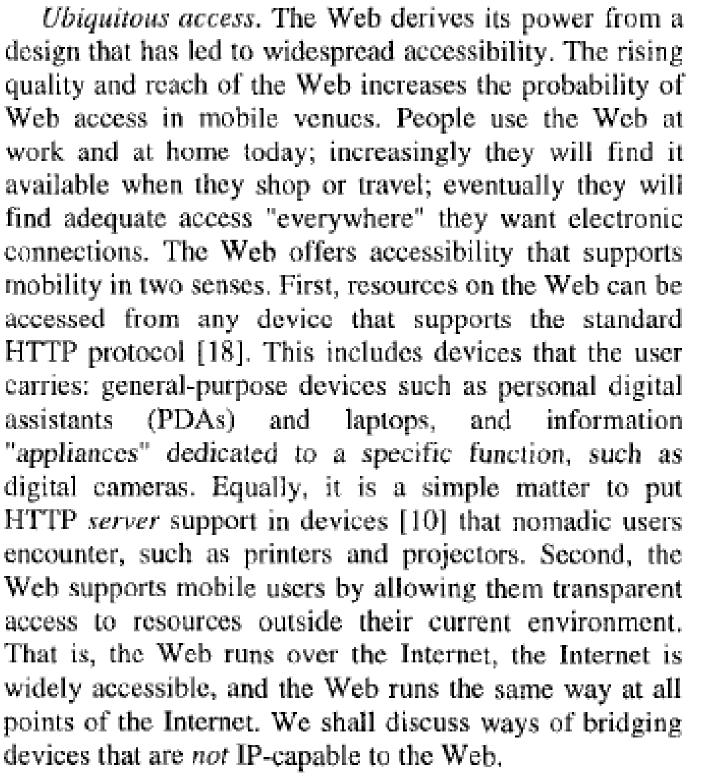 *Ubiquitous access*. The Web derives its power from a design that has led to widespread accessibility. The rising quality and reach of the Web increases the probability of Web access in mobile venues. People use the Web at work and at home today; increasingly they will find it available when they shop or travel; eventually they will find adequate access "everywhere" they want electronic connections. The Web offers accessibility that supports mobility in two senses. First, resources on the Web can be accessed from any device that supports the standard HTTP protocol [18]. This includes devices that the user carries: general-purpose devices such as personal digital assistants (PDAs) and laptops, and information "appliances" dedicated to a specific function, such as digital cameras. Equally, it is a simple matter to put HTTP *server* support in devices [10] that nomadic users encounter, such as printers and projectors. Second, the Web supports mobile users by allowing them transparent access to resources outside their current environment. That is, the Web runs over the Internet, the Internet is widely accessible, and the Web runs the same way at all points of the Internet. We shall discuss ways of bridging devices that are *not* IP-capable to the Web. To the extent it is found that CoolTown does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that CoolTown does not anticipate this claim, CoolTown renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine CoolTown with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these

Element	'278 Claim Recitation	Exemplary Citations in Reference
		references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1B]	determining that symbolic content is on the at least the portion of the document based on the image;	CoolTown discloses determining that symbolic content is on the at least the portion of the document based on the image. A Web-Based Nomadic Computing System (HP000074 – 87) at 3 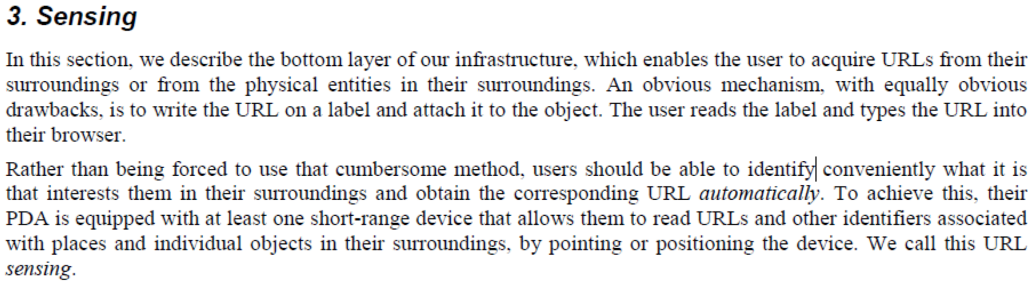 ***3. Sensing*** In this section, we describe the bottom layer of our infrastructure, which enables the user to acquire URLs from their surroundings or from the physical entities in their surroundings. An obvious mechanism, with equally obvious drawbacks, is to write the URL on a label and attach it to the object. The user reads the label and types the URL into their browser. Rather than being forced to use that cumbersome method, users should be able to identify conveniently what it is that interests them in their surroundings and obtain the corresponding URL *automatically*. To achieve this, their PDA is equipped with at least one short-range device that allows them to read URLs and other identifiers associated with places and individual objects in their surroundings, by pointing or positioning the device. We call this URL *sensing*. A Web-Based Nomadic Computing System (HP000074 – 87) at 4 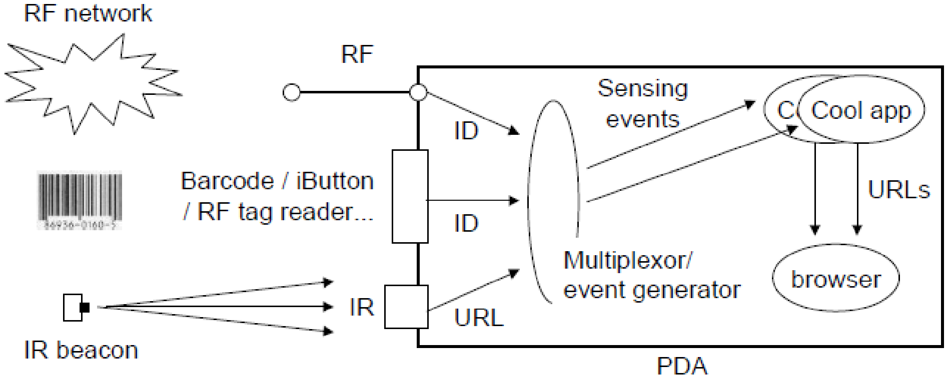 Figure 3. A sketch of a PDA equipped for sensing URLs and identifiers.

Element	'278 Claim Recitation	Exemplary Citations in Reference
		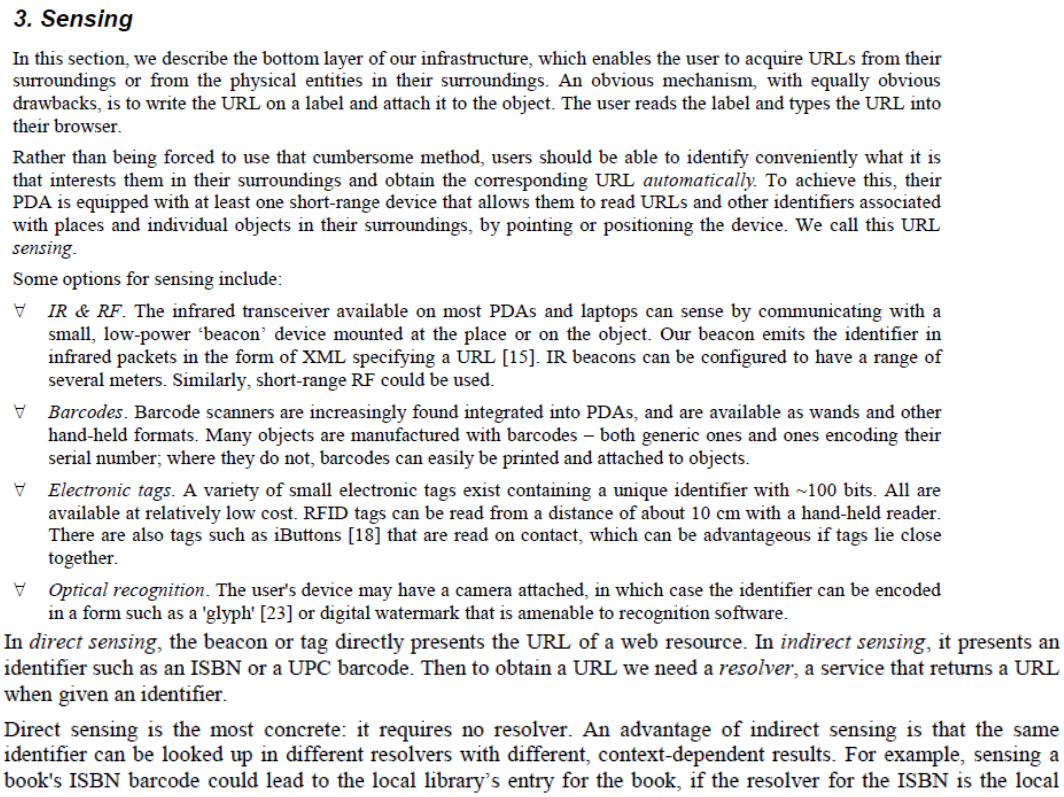 ***3. Sensing*** In this section, we describe the bottom layer of our infrastructure, which enables the user to acquire URLs from their surroundings or from the physical entities in their surroundings. An obvious mechanism, with equally obvious drawbacks, is to write the URL on a label and attach it to the object. The user reads the label and types the URL into their browser. Rather than being forced to use that cumbersome method, users should be able to identify conveniently what it is that interests them in their surroundings and obtain the corresponding URL *automatically*. To achieve this, their PDA is equipped with at least one short-range device that allows them to read URLs and other identifiers associated with places and individual objects in their surroundings, by pointing or positioning the device. We call this URL *sensing*. Some options for sensing include: ∀ *IR & RF*. The infrared transceiver available on most PDAs and laptops can sense by communicating with a small, low-power 'beacon' device mounted at the place or on the object. Our beacon emits the identifier in infrared packets in the form of XML specifying a URL [15]. IR beacons can be configured to have a range of several meters. Similarly, short-range RF could be used. ∀ *Barcodes*. Barcode scanners are increasingly found integrated into PDAs, and are available as wands and other hand-held formats. Many objects are manufactured with barcodes – both generic ones and ones encoding their serial number; where they do not, barcodes can easily be printed and attached to objects. ∀ *Electronic tags*. A variety of small electronic tags exist containing a unique identifier with ~100 bits. All are available at relatively low cost. RFID tags can be read from a distance of about 10 cm with a hand-held reader. There are also tags such as iButtons [18] that are read on contact, which can be advantageous if tags lie close together. ∀ *Optical recognition*. The user's device may have a camera attached, in which case the identifier can be encoded in a form such as a 'glyph' [23] or digital watermark that is amenable to recognition software. In *direct sensing*, the beacon or tag directly presents the URL of a web resource. In *indirect sensing*, it presents an identifier such as an ISBN or a UPC barcode. Then to obtain a URL we need a *resolver*, a service that returns a URL when given an identifier. Direct sensing is the most concrete: it requires no resolver. An advantage of indirect sensing is that the same identifier can be looked up in different resolvers with different, context-dependent results. For example, sensing a book's ISBN barcode could lead to the local library's entry for the book, if the resolver for the ISBN is the local Cooltown Developer's Network, People, Place, and Things: web presence for the real world (HP000025–34) (HP000025–34) at 20 – 21

Element	'278 Claim Recitation	Exemplary Citations in Reference
		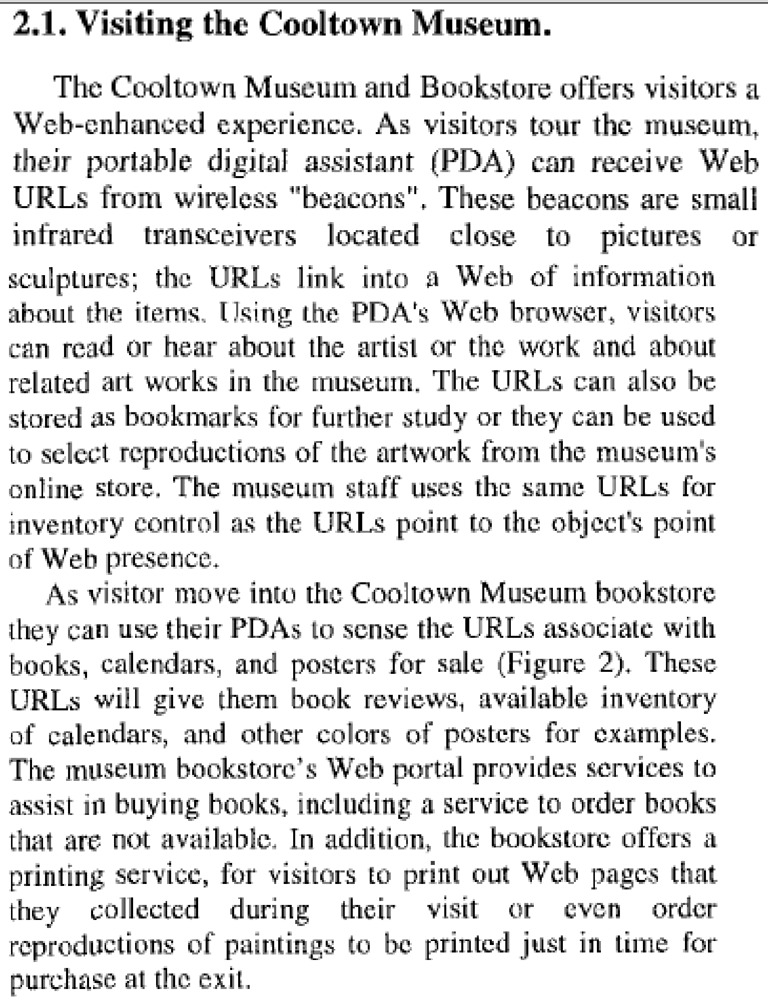 **2.1. Visiting the Cooltown Museum.** The Cooltown Museum and Bookstore offers visitors a Web-enhanced experience. As visitors tour the museum, their portable digital assistant (PDA) can receive Web URLs from wireless "beacons". These beacons are small infrared transceivers located close to pictures or sculptures; the URLs link into a Web of information about the items. Using the PDA's Web browser, visitors can read or hear about the artist or the work and about related art works in the museum. The URLs can also be stored as bookmarks for further study or they can be used to select reproductions of the artwork from the museum's online store. The museum staff uses the same URLs for inventory control as the URLs point to the object's point of Web presence. As visitor move into the Cooltown Museum bookstore they can use their PDAs to sense the URLs associate with books, calendars, and posters for sale (Figure 2). These URLs will give them book reviews, available inventory of calendars, and other colors of posters for examples. The museum bookstore's Web portal provides services to assist in buying books, including a service to order books that are not available. In addition, the bookstore offers a printing service, for visitors to print out Web pages that they collected during their visit or even order reproductions of paintings to be printed just in time for purchase at the exit. Cooltown Developer's Network, People, Place, and Things: web presence for the real world (HP000025–34) (HP000025–34) at 23

Element	'278 Claim Recitation	Exemplary Citations in Reference
		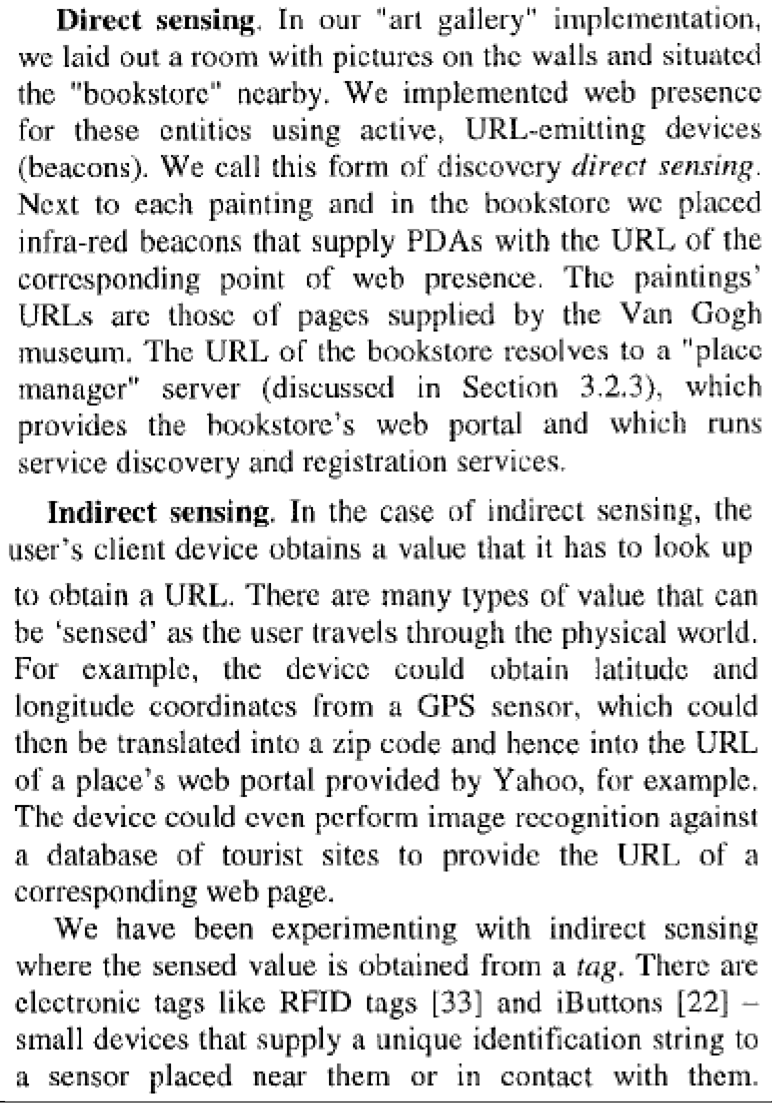 **Direct sensing**. In our "art gallery" implementation, we laid out a room with pictures on the walls and situated the "bookstore" nearby. We implemented web presence for these entities using active, URL-emitting devices (beacons). We call this form of discovery *direct sensing*. Next to each painting and in the bookstore we placed infra-red beacons that supply PDAs with the URL of the corresponding point of web presence. The paintings' URLs are those of pages supplied by the Van Gogh museum. The URL of the bookstore resolves to a "place manager" server (discussed in Section 3.2.3), which provides the bookstore's web portal and which runs service discovery and registration services. **Indirect sensing**. In the case of indirect sensing, the user's client device obtains a value that it has to look up to obtain a URL. There are many types of value that can be 'sensed' as the user travels through the physical world. For example, the device could obtain latitude and longitude coordinates from a GPS sensor, which could then be translated into a zip code and hence into the URL of a place's web portal provided by Yahoo, for example. The device could even perform image recognition against a database of tourist sites to provide the URL of a corresponding web page. We have been experimenting with indirect sensing where the sensed value is obtained from a *tag*. There are electronic tags like RFID tags [33] and iButtons [22] – small devices that supply a unique identification string to a sensor placed near them or in contact with them.

Element	'278 Claim Recitation	Exemplary Citations in Reference
		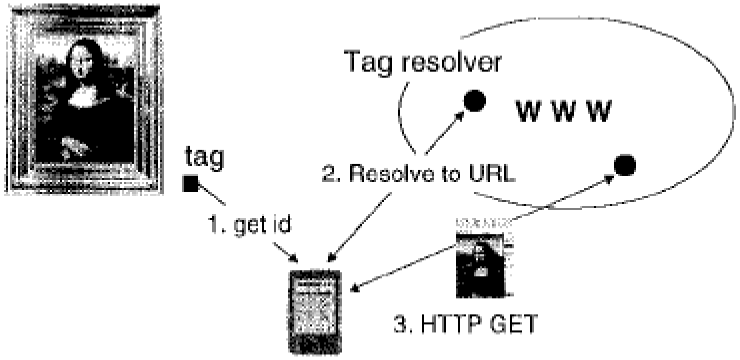 **Figure 4. Indirect sensing: the web presence for a painting.** Barcodes can be used instead of an electronic tag. Figure 4 shows a different art gallery demonstrator in which a PDA client senses a painting's tag. A correlated web page is resolved from the tag identifier and the client presents it to the user. We have produced a prototype implementation for sensing tagged things and displaying a corresponding context-dependent web page. The sensing device knows the URL of a *resolver*. When the client senses a tag, it sends the identifier to the resolver. If the resolver has an entry for the identifier, it sends back the corresponding URL. The user may choose a resolver that reflects his or her personal requirements, or those of a group of users with interests or tasks in common. For example, visitors to an art gallery could be offered resolvers that provide the URLs of pages in their native language, or pages provided by the art school at which they study.

Element	'278 Claim Recitation	Exemplary Citations in Reference
		The Cooltown User Experience (HP000088 – 093) at 4 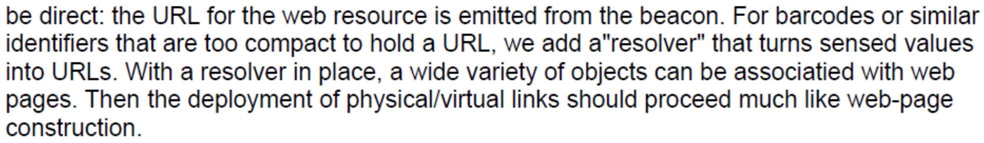 be direct: the URL for the web resource is emitted from the beacon. For barcodes or similar identifiers that are too compact to hold a URL, we add a"resolver" that turns sensed values into URLs. With a resolver in place, a wide variety of objects can be associatied with web pages. Then the deployment of physical/virtual links should proceed much like web-page construction. To the extent it is found that CoolTown does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that CoolTown does not anticipate this claim, CoolTown renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine CoolTown with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1C]	extracting symbol information based on the symbolic content according to symbol type;	CoolTown discloses extracting symbol information based on the symbolic content according to symbol type. A Web-Based Nomadic Computing System (HP000074 – 87) at 3

Element	'278 Claim Recitation	Exemplary Citations in Reference
		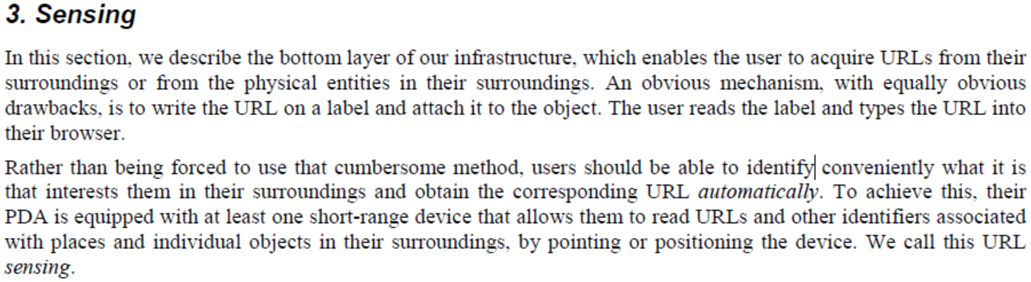 ***3. Sensing*** In this section, we describe the bottom layer of our infrastructure, which enables the user to acquire URLs from their surroundings or from the physical entities in their surroundings. An obvious mechanism, with equally obvious drawbacks, is to write the URL on a label and attach it to the object. The user reads the label and types the URL into their browser. Rather than being forced to use that cumbersome method, users should be able to identify conveniently what it is that interests them in their surroundings and obtain the corresponding URL *automatically*. To achieve this, their PDA is equipped with at least one short-range device that allows them to read URLs and other identifiers associated with places and individual objects in their surroundings, by pointing or positioning the device. We call this URL *sensing*. A Web-Based Nomadic Computing System (HP000074 – 87) at 4 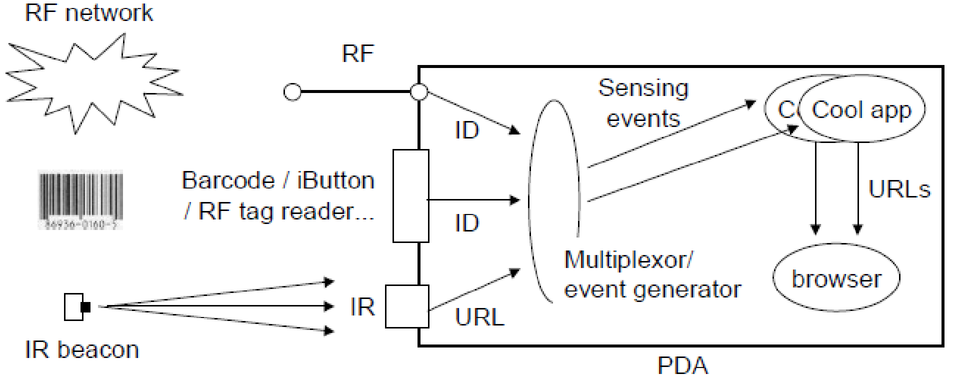 Figure 3. A sketch of a PDA equipped for sensing URLs and identifiers.

Element	'278 Claim Recitation	Exemplary Citations in Reference
		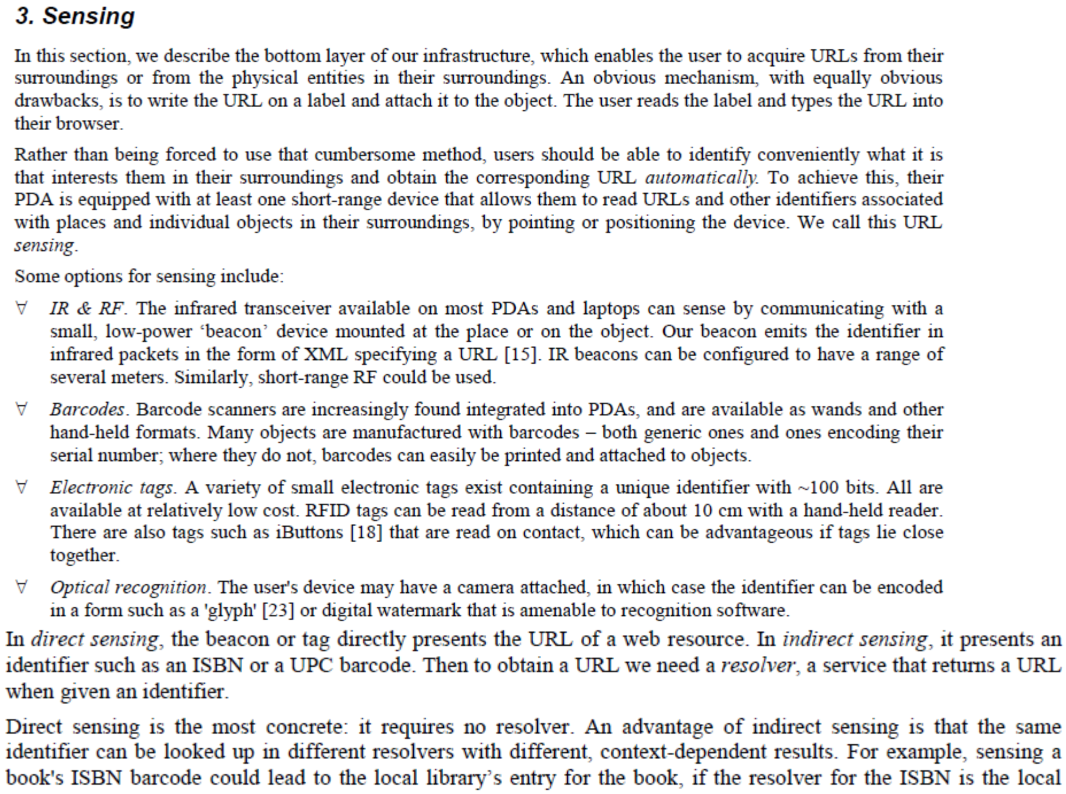 ***3. Sensing*** In this section, we describe the bottom layer of our infrastructure, which enables the user to acquire URLs from their surroundings or from the physical entities in their surroundings. An obvious mechanism, with equally obvious drawbacks, is to write the URL on a label and attach it to the object. The user reads the label and types the URL into their browser. Rather than being forced to use that cumbersome method, users should be able to identify conveniently what it is that interests them in their surroundings and obtain the corresponding URL *automatically*. To achieve this, their PDA is equipped with at least one short-range device that allows them to read URLs and other identifiers associated with places and individual objects in their surroundings, by pointing or positioning the device. We call this URL *sensing*. Some options for sensing include: ∀ *IR & RF*. The infrared transceiver available on most PDAs and laptops can sense by communicating with a small, low-power 'beacon' device mounted at the place or on the object. Our beacon emits the identifier in infrared packets in the form of XML specifying a URL [15]. IR beacons can be configured to have a range of several meters. Similarly, short-range RF could be used. ∀ *Barcodes*. Barcode scanners are increasingly found integrated into PDAs, and are available as wands and other hand-held formats. Many objects are manufactured with barcodes – both generic ones and ones encoding their serial number; where they do not, barcodes can easily be printed and attached to objects. ∀ *Electronic tags*. A variety of small electronic tags exist containing a unique identifier with ~100 bits. All are available at relatively low cost. RFID tags can be read from a distance of about 10 cm with a hand-held reader. There are also tags such as iButtons [18] that are read on contact, which can be advantageous if tags lie close together. ∀ *Optical recognition*. The user's device may have a camera attached, in which case the identifier can be encoded in a form such as a 'glyph' [23] or digital watermark that is amenable to recognition software. In *direct sensing*, the beacon or tag directly presents the URL of a web resource. In *indirect sensing*, it presents an identifier such as an ISBN or a UPC barcode. Then to obtain a URL we need a *resolver*, a service that returns a URL when given an identifier. Direct sensing is the most concrete: it requires no resolver. An advantage of indirect sensing is that the same identifier can be looked up in different resolvers with different, context-dependent results. For example, sensing a book's ISBN barcode could lead to the local library's entry for the book, if the resolver for the ISBN is the local Cooltown Developer's Network, People, Place, and Things: web presence for the real world (HP000025–34) (HP000025–34) at 23

Element	'278 Claim Recitation	Exemplary Citations in Reference
		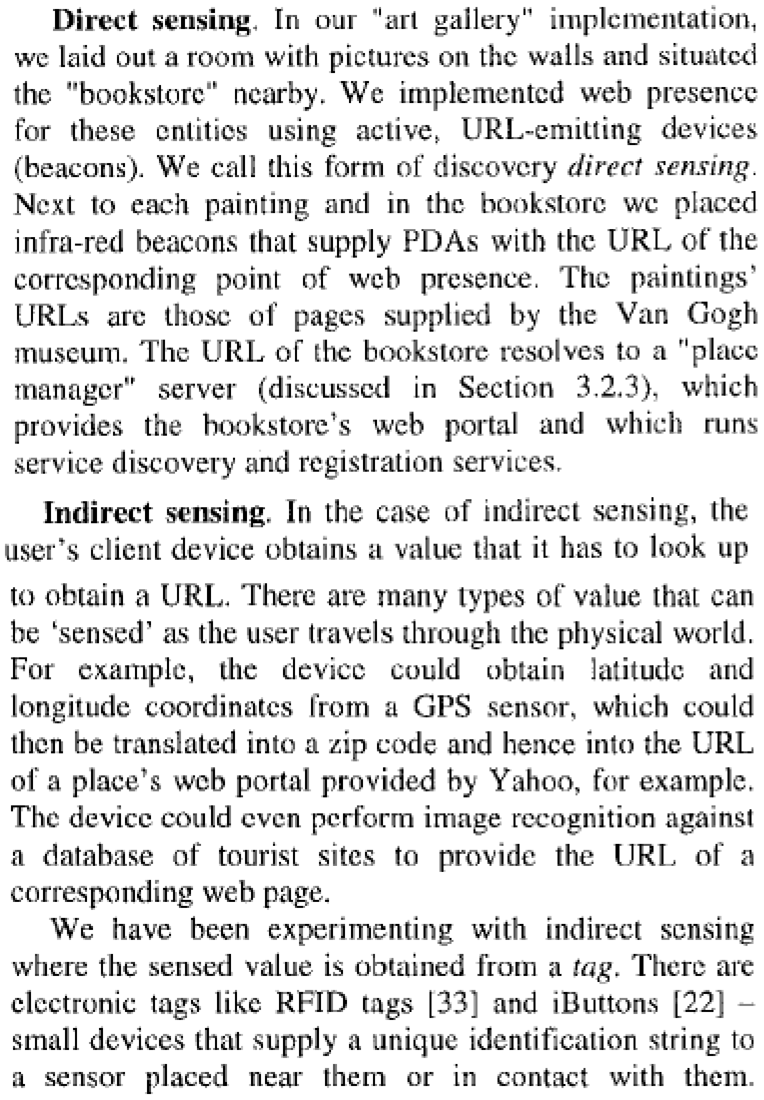 **Direct sensing**. In our "art gallery" implementation, we laid out a room with pictures on the walls and situated the "bookstore" nearby. We implemented web presence for these entities using active, URL-emitting devices (beacons). We call this form of discovery *direct sensing*. Next to each painting and in the bookstore we placed infra-red beacons that supply PDAs with the URL of the corresponding point of web presence. The paintings' URLs are those of pages supplied by the Van Gogh museum. The URL of the bookstore resolves to a "place manager" server (discussed in Section 3.2.3), which provides the bookstore's web portal and which runs service discovery and registration services. **Indirect sensing**. In the case of indirect sensing, the user's client device obtains a value that it has to look up to obtain a URL. There are many types of value that can be 'sensed' as the user travels through the physical world. For example, the device could obtain latitude and longitude coordinates from a GPS sensor, which could then be translated into a zip code and hence into the URL of a place's web portal provided by Yahoo, for example. The device could even perform image recognition against a database of tourist sites to provide the URL of a corresponding web page. We have been experimenting with indirect sensing where the sensed value is obtained from a *tag*. There are electronic tags like RFID tags [33] and iButtons [22] – small devices that supply a unique identification string to a sensor placed near them or in contact with them.

Element	'278 Claim Recitation	Exemplary Citations in Reference
		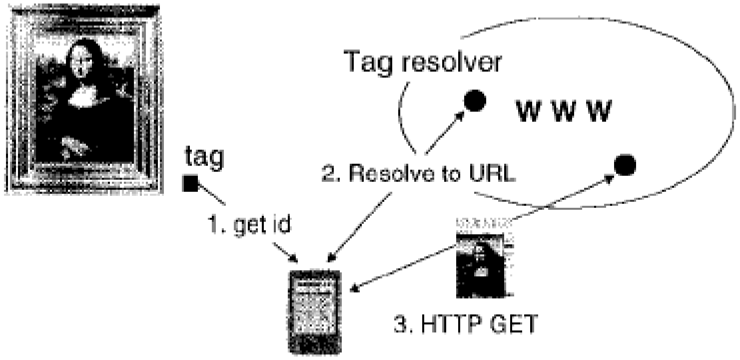 **Figure 4. Indirect sensing: the web presence for a painting.** Barcodes can be used instead of an electronic tag. Figure 4 shows a different art gallery demonstrator in which a PDA client senses a painting's tag. A correlated web page is resolved from the tag identifier and the client presents it to the user. We have produced a prototype implementation for sensing tagged things and displaying a corresponding context-dependent web page. The sensing device knows the URL of a *resolver*. When the client senses a tag, it sends the identifier to the resolver. If the resolver has an entry for the identifier, it sends back the corresponding URL. The user may choose a resolver that reflects his or her personal requirements, or those of a group of users with interests or tasks in common. For example, visitors to an art gallery could be offered resolvers that provide the URLs of pages in their native language, or pages provided by the art school at which they study.

Element	'278 Claim Recitation	Exemplary Citations in Reference
		The Cooltown User Experience (HP000088 – 093) at 4 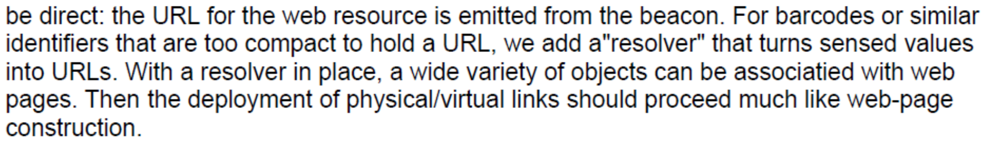 be direct: the URL for the web resource is emitted from the beacon. For barcodes or similar identifiers that are too compact to hold a URL, we add a"resolver" that turns sensed values into URLs. With a resolver in place, a wide variety of objects can be associatied with web pages. Then the deployment of physical/virtual links should proceed much like web-page construction. A Web-Based Nomadic Computing System (HP000074 – 87) at 6 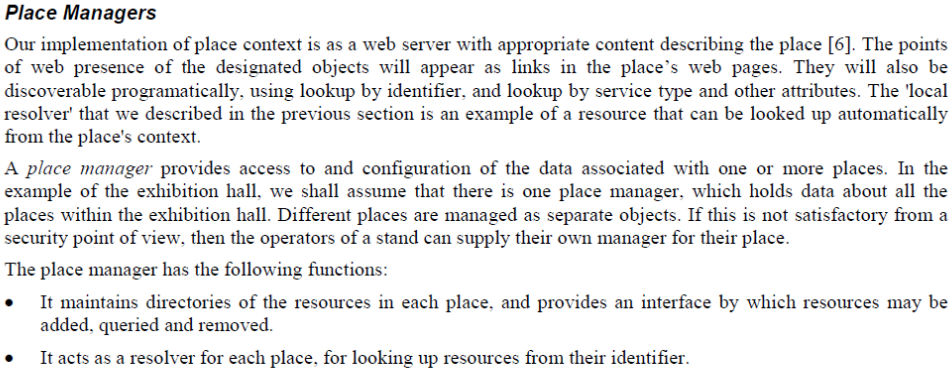 ***Place Managers*** Our implementation of place context is as a web server with appropriate content describing the place [6]. The points of web presence of the designated objects will appear as links in the place's web pages. They will also be discoverable programatically, using lookup by identifier, and lookup by service type and other attributes. The 'local resolver' that we described in the previous section is an example of a resource that can be looked up automatically from the place's context. A *place manager* provides access to and configuration of the data associated with one or more places. In the example of the exhibition hall, we shall assume that there is one place manager, which holds data about all the places within the exhibition hall. Different places are managed as separate objects. If this is not satisfactory from a security point of view, then the operators of a stand can supply their own manager for their place. The place manager has the following functions: • It maintains directories of the resources in each place, and provides an interface by which resources may be added, queried and removed. • It acts as a resolver for each place, for looking up resources from their identifier. The Cooltown User Experience (HP000074 – 87) at 4 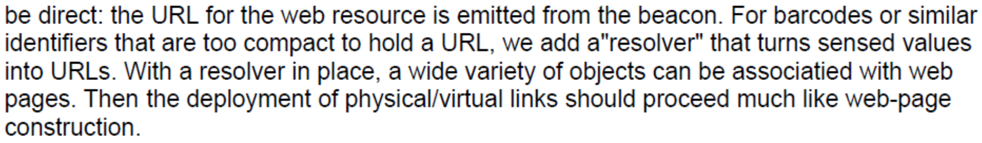 be direct: the URL for the web resource is emitted from the beacon. For barcodes or similar identifiers that are too compact to hold a URL, we add a"resolver" that turns sensed values into URLs. With a resolver in place, a wide variety of objects can be associatied with web pages. Then the deployment of physical/virtual links should proceed much like web-page construction. To the extent it is found that CoolTown does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that CoolTown does not anticipate this claim, CoolTown renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination

Element	'278 Claim Recitation	Exemplary Citations in Reference
		with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine CoolTown with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1D]	determining a validity of the document based at least in part on the image and the symbol information; and	CoolTown teaches determining a validity of the document based at least in part on the image and the symbol information. Cooltown Developer's Network, People, Place, and Things: web presence for the real world (HP000025–34) (HP000025–34) at 20 – 21 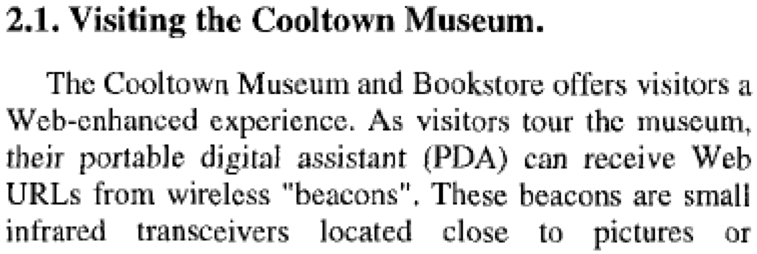 **2.1. Visiting the Cooltown Museum.** The Cooltown Museum and Bookstore offers visitors a Web-enhanced experience. As visitors tour the museum, their portable digital assistant (PDA) can receive Web URLs from wireless "beacons". These beacons are small infrared transceivers located close to pictures or

Element	'278 Claim Recitation	Exemplary Citations in Reference
		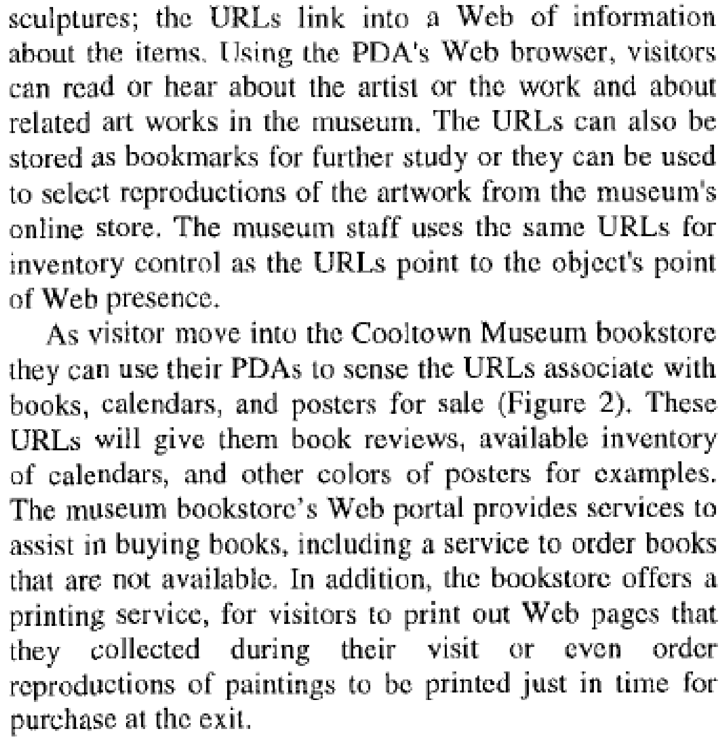 sculptures; the URLs link into a Web of information about the items. Using the PDA's Web browser, visitors can read or hear about the artist or the work and about related art works in the museum. The URLs can also be stored as bookmarks for further study or they can be used to select reproductions of the artwork from the museum's online store. The museum staff uses the same URLs for inventory control as the URLs point to the object's point of Web presence. As visitor move into the Cooltown Museum bookstore they can use their PDAs to sense the URLs associate with books, calendars, and posters for sale (Figure 2). These URLs will give them book reviews, available inventory of calendars, and other colors of posters for examples. The museum bookstore's Web portal provides services to assist in buying books, including a service to order books that are not available. In addition, the bookstore offers a printing service, for visitors to print out Web pages that they collected during their visit or even order reproductions of paintings to be printed just in time for purchase at the exit. Cooltown Developer's Network, People, Place, and Things: web presence for the real world (HP000025–34) (HP000025–34) at 23

Element	'278 Claim Recitation	Exemplary Citations in Reference
		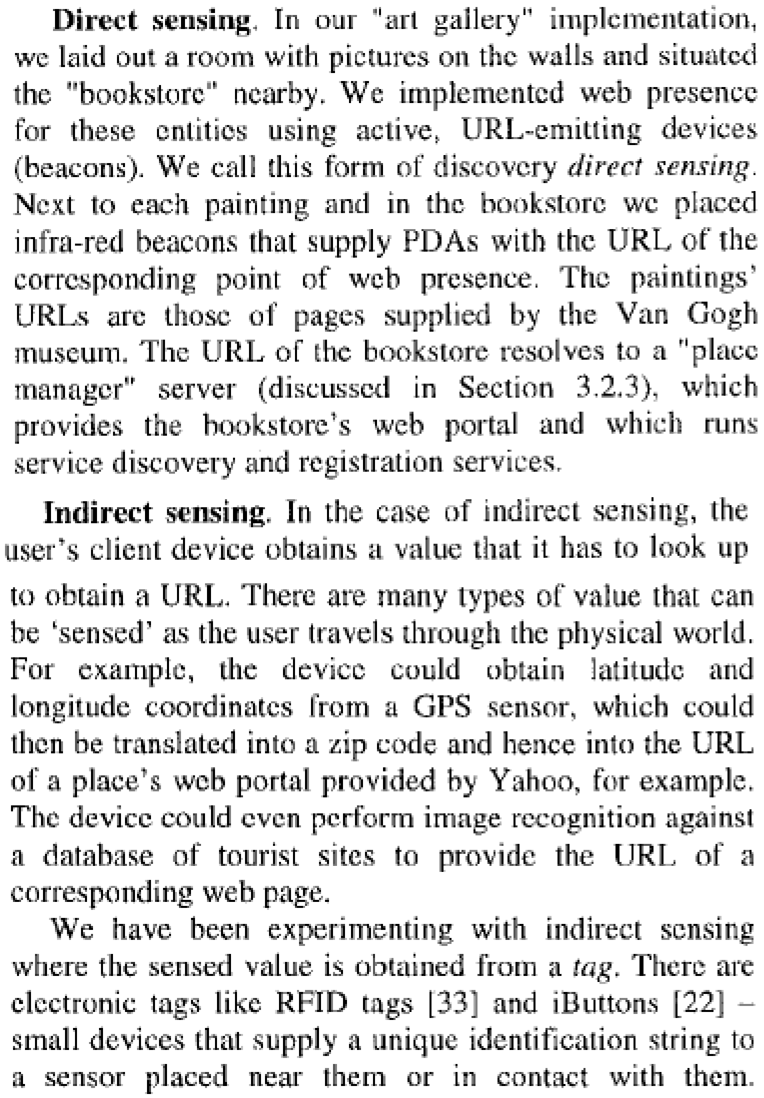 **Direct sensing**. In our "art gallery" implementation, we laid out a room with pictures on the walls and situated the "bookstore" nearby. We implemented web presence for these entities using active, URL-emitting devices (beacons). We call this form of discovery *direct sensing*. Next to each painting and in the bookstore we placed infra-red beacons that supply PDAs with the URL of the corresponding point of web presence. The paintings' URLs are those of pages supplied by the Van Gogh museum. The URL of the bookstore resolves to a "place manager" server (discussed in Section 3.2.3), which provides the bookstore's web portal and which runs service discovery and registration services. **Indirect sensing**. In the case of indirect sensing, the user's client device obtains a value that it has to look up to obtain a URL. There are many types of value that can be 'sensed' as the user travels through the physical world. For example, the device could obtain latitude and longitude coordinates from a GPS sensor, which could then be translated into a zip code and hence into the URL of a place's web portal provided by Yahoo, for example. The device could even perform image recognition against a database of tourist sites to provide the URL of a corresponding web page. We have been experimenting with indirect sensing where the sensed value is obtained from a *tag*. There are electronic tags like RFID tags [33] and iButtons [22] – small devices that supply a unique identification string to a sensor placed near them or in contact with them.

Element	'278 Claim Recitation	Exemplary Citations in Reference
		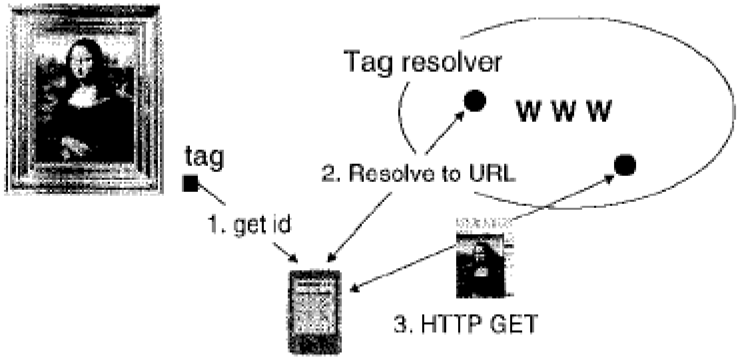 **Figure 4. Indirect sensing: the web presence for a painting.** Barcodes can be used instead of an electronic tag. Figure 4 shows a different art gallery demonstrator in which a PDA client senses a painting's tag. A correlated web page is resolved from the tag identifier and the client presents it to the user. We have produced a prototype implementation for sensing tagged things and displaying a corresponding context-dependent web page. The sensing device knows the URL of a *resolver*. When the client senses a tag, it sends the identifier to the resolver. If the resolver has an entry for the identifier, it sends back the corresponding URL. The user may choose a resolver that reflects his or her personal requirements, or those of a group of users with interests or tasks in common. For example, visitors to an art gallery could be offered resolvers that provide the URLs of pages in their native language, or pages provided by the art school at which they study.

Element	'278 Claim Recitation	Exemplary Citations in Reference
		A Web-Based Nomadic Computing System (HP000074 – 87) at 10 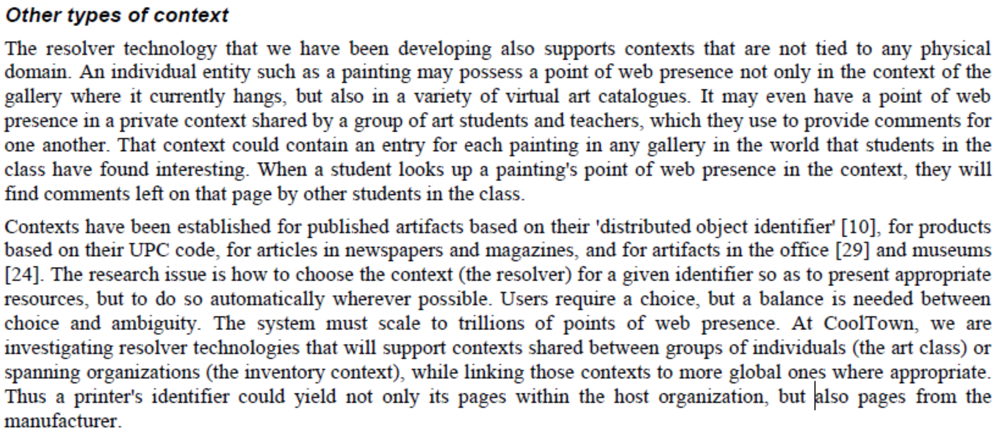 ***Other types of context*** The resolver technology that we have been developing also supports contexts that are not tied to any physical domain. An individual entity such as a painting may possess a point of web presence not only in the context of the gallery where it currently hangs, but also in a variety of virtual art catalogues. It may even have a point of web presence in a private context shared by a group of art students and teachers, which they use to provide comments for one another. That context could contain an entry for each painting in any gallery in the world that students in the class have found interesting. When a student looks up a painting's point of web presence in the context, they will find comments left on that page by other students in the class. Contexts have been established for published artifacts based on their 'distributed object identifier' [10], for products based on their UPC code, for articles in newspapers and magazines, and for artifacts in the office [29] and museums [24]. The research issue is how to choose the context (the resolver) for a given identifier so as to present appropriate resources, but to do so automatically wherever possible. Users require a choice, but a balance is needed between choice and ambiguity. The system must scale to trillions of points of web presence. At CoolTown, we are investigating resolver technologies that will support contexts shared between groups of individuals (the art class) or spanning organizations (the inventory context), while linking those contexts to more global ones where appropriate. Thus a printer's identifier could yield not only its pages within the host organization, but also pages from the manufacturer. A Web-Based Nomadic Computing System (HP000074 – 87) at Fig. 2

Element	'278 Claim Recitation	Exemplary Citations in Reference
		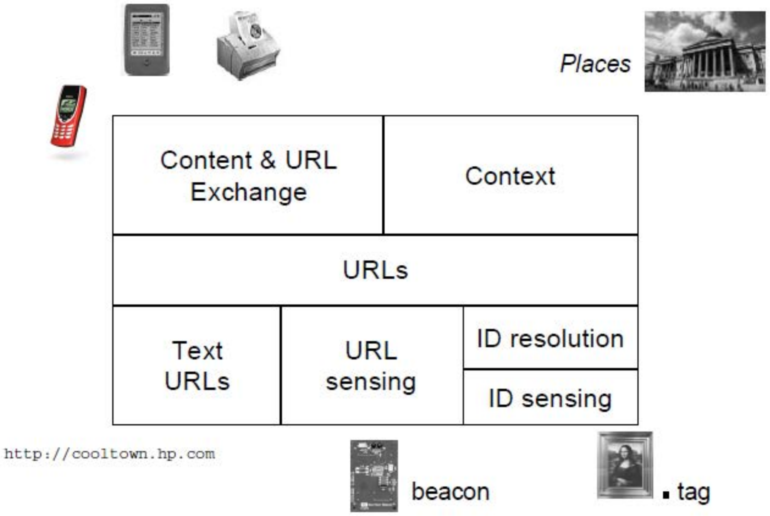 Figure 2. Component relationships. Sources of URLs, including keyboards, beacons, and resolved tags point to context or content. A Web-Based Nomadic Computing System (HP000074 – 87) at 4

Element	'278 Claim Recitation	Exemplary Citations in Reference
		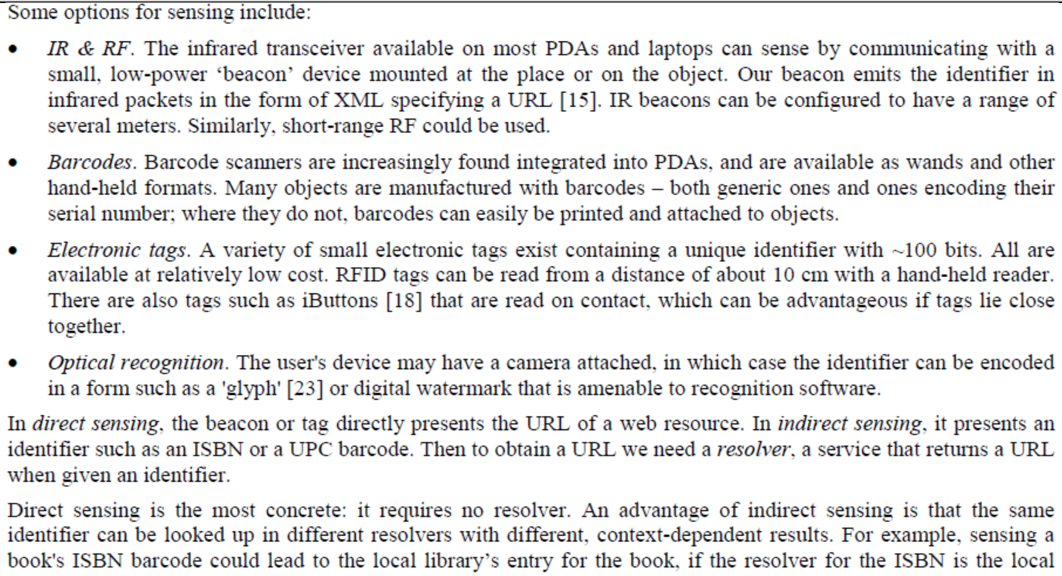 Some options for sensing include: • *IR & RF*. The infrared transceiver available on most PDAs and laptops can sense by communicating with a small, low-power 'beacon' device mounted at the place or on the object. Our beacon emits the identifier in infrared packets in the form of XML specifying a URL [15]. IR beacons can be configured to have a range of several meters. Similarly, short-range RF could be used. • *Barcodes*. Barcode scanners are increasingly found integrated into PDAs, and are available as wands and other hand-held formats. Many objects are manufactured with barcodes – both generic ones and ones encoding their serial number; where they do not, barcodes can easily be printed and attached to objects. • *Electronic tags*. A variety of small electronic tags exist containing a unique identifier with ~100 bits. All are available at relatively low cost. RFID tags can be read from a distance of about 10 cm with a hand-held reader. There are also tags such as iButtons [18] that are read on contact, which can be advantageous if tags lie close together. • *Optical recognition*. The user's device may have a camera attached, in which case the identifier can be encoded in a form such as a 'glyph' [23] or digital watermark that is amenable to recognition software. In *direct sensing*, the beacon or tag directly presents the URL of a web resource. In *indirect sensing*, it presents an identifier such as an ISBN or a UPC barcode. Then to obtain a URL we need a *resolver*, a service that returns a URL when given an identifier. Direct sensing is the most concrete: it requires no resolver. An advantage of indirect sensing is that the same identifier can be looked up in different resolvers with different, context-dependent results. For example, sensing a book's ISBN barcode could lead to the local library's entry for the book, if the resolver for the ISBN is the local To the extent it is found that CoolTown does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that CoolTown does not anticipate this claim, CoolTown renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine CoolTown with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

Element	'278 Claim Recitation	Exemplary Citations in Reference
[1E]	recognizing the document as a first target object based at least in part on	CoolTown discloses this limitation. A Web-Based Nomadic Computing System (HP000074 – 87) at Fig. 2 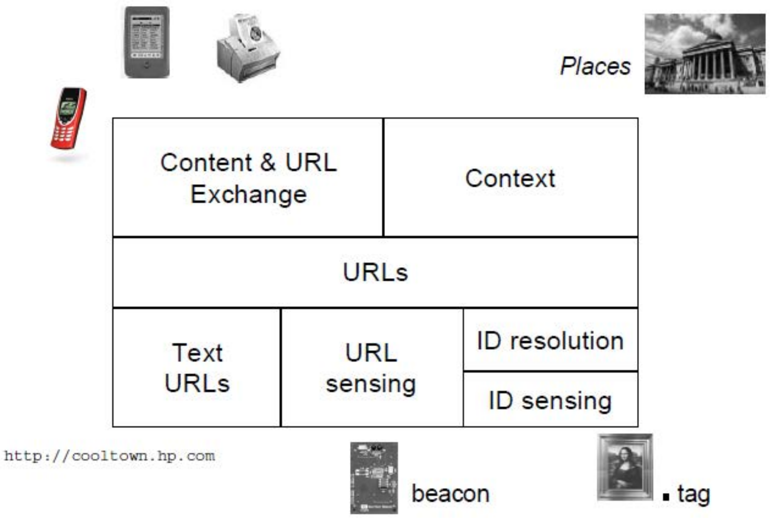 Figure 2. Component relationships. Sources of URLs, including keyboards, beacons, and resolved tags point to context or content. A Web-Based Nomadic Computing System (HP000074 – 87) at 3 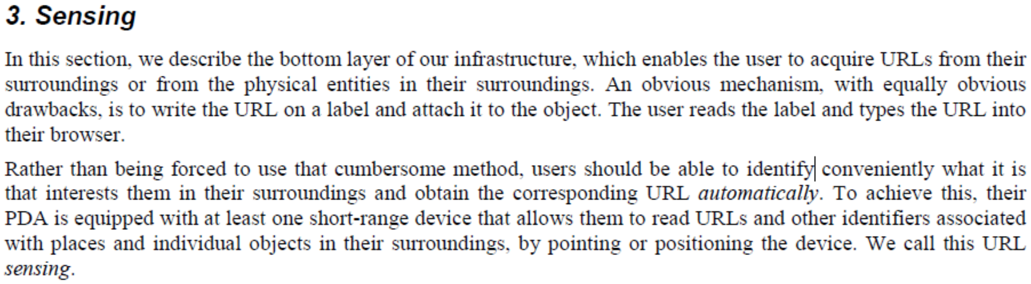 ***3. Sensing*** In this section, we describe the bottom layer of our infrastructure, which enables the user to acquire URLs from their surroundings or from the physical entities in their surroundings. An obvious mechanism, with equally obvious drawbacks, is to write the URL on a label and attach it to the object. The user reads the label and types the URL into their browser. Rather than being forced to use that cumbersome method, users should be able to identify conveniently what it is that interests them in their surroundings and obtain the corresponding URL *automatically*. To achieve this, their PDA is equipped with at least one short-range device that allows them to read URLs and other identifiers associated with places and individual objects in their surroundings, by pointing or positioning the device. We call this URL *sensing*.

Element	'278 Claim Recitation	Exemplary Citations in Reference
		A Web-Based Nomadic Computing System (HP000074 – 87) at 4 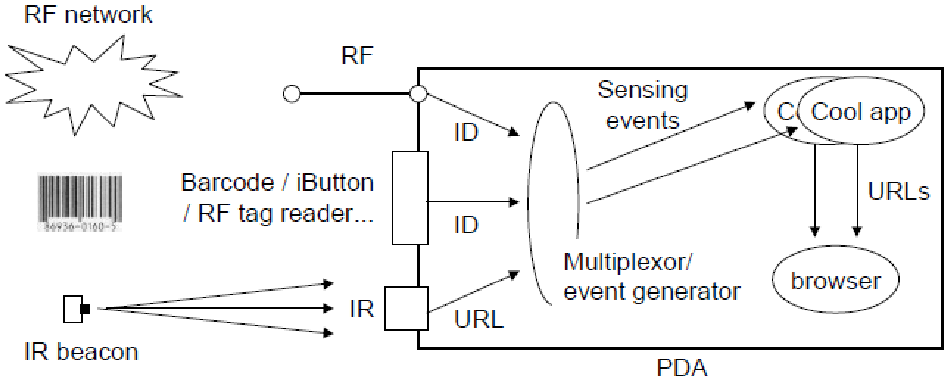 Figure 3. A sketch of a PDA equipped for sensing URLs and identifiers. A Web-Based Nomadic Computing System (HP000074 – 87) at 4

Element	'278 Claim Recitation	Exemplary Citations in Reference
		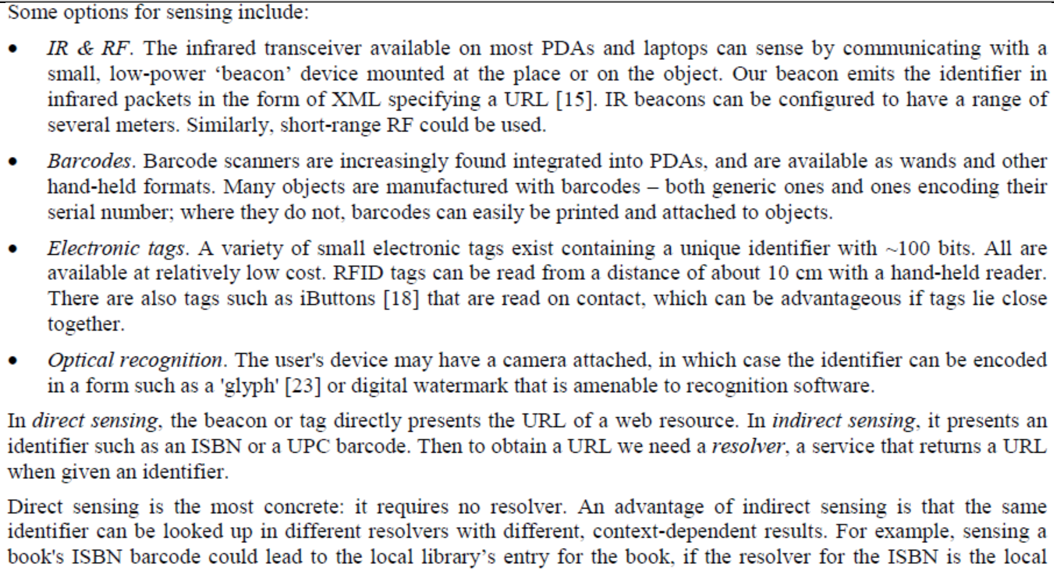 Some options for sensing include: • *IR & RF*. The infrared transceiver available on most PDAs and laptops can sense by communicating with a small, low-power 'beacon' device mounted at the place or on the object. Our beacon emits the identifier in infrared packets in the form of XML specifying a URL [15]. IR beacons can be configured to have a range of several meters. Similarly, short-range RF could be used. • *Barcodes*. Barcode scanners are increasingly found integrated into PDAs, and are available as wands and other hand-held formats. Many objects are manufactured with barcodes – both generic ones and ones encoding their serial number; where they do not, barcodes can easily be printed and attached to objects. • *Electronic tags*. A variety of small electronic tags exist containing a unique identifier with ~100 bits. All are available at relatively low cost. RFID tags can be read from a distance of about 10 cm with a hand-held reader. There are also tags such as iButtons [18] that are read on contact, which can be advantageous if tags lie close together. • *Optical recognition*. The user's device may have a camera attached, in which case the identifier can be encoded in a form such as a 'glyph' [23] or digital watermark that is amenable to recognition software. In *direct sensing*, the beacon or tag directly presents the URL of a web resource. In *indirect sensing*, it presents an identifier such as an ISBN or a UPC barcode. Then to obtain a URL we need a *resolver*, a service that returns a URL when given an identifier. Direct sensing is the most concrete: it requires no resolver. An advantage of indirect sensing is that the same identifier can be looked up in different resolvers with different, context-dependent results. For example, sensing a book's ISBN barcode could lead to the local library's entry for the book, if the resolver for the ISBN is the local A Web-Based Nomadic Computing System (HP000074 – 87) at 5 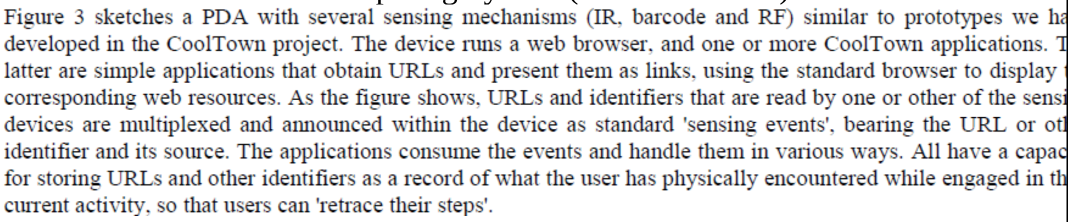 Figure 3 sketches a PDA with several sensing mechanisms (IR, barcode and RF) similar to prototypes we ha developed in the CoolTown project. The device runs a web browser, and one or more CoolTown applications. T latter are simple applications that obtain URLs and present them as links, using the standard browser to display corresponding web resources. As the figure shows, URLs and identifiers that are read by one or other of the sensi devices are multiplexed and announced within the device as standard 'sensing events', bearing the URL or ot identifier and its source. The applications consume the events and handle them in various ways. All have a capac for storing URLs and other identifiers as a record of what the user has physically encountered while engaged in th current activity, so that users can 'retrace their steps'. A Web-Based Nomadic Computing System (HP000074 – 87) at 6 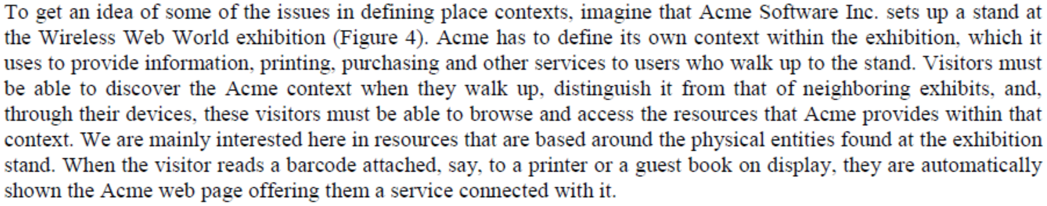 To get an idea of some of the issues in defining place contexts, imagine that Acme Software Inc. sets up a stand at the Wireless Web World exhibition (Figure 4). Acme has to define its own context within the exhibition, which it uses to provide information, printing, purchasing and other services to users who walk up to the stand. Visitors must be able to discover the Acme context when they walk up, distinguish it from that of neighboring exhibits, and, through their devices, these visitors must be able to browse and access the resources that Acme provides within that context. We are mainly interested here in resources that are based around the physical entities found at the exhibition stand. When the visitor reads a barcode attached, say, to a printer or a guest book on display, they are automatically shown the Acme web page offering them a service connected with it.

Element	'278 Claim Recitation	Exemplary Citations in Reference
		A Web-Based Nomadic Computing System (HP000074 – 87) at 6 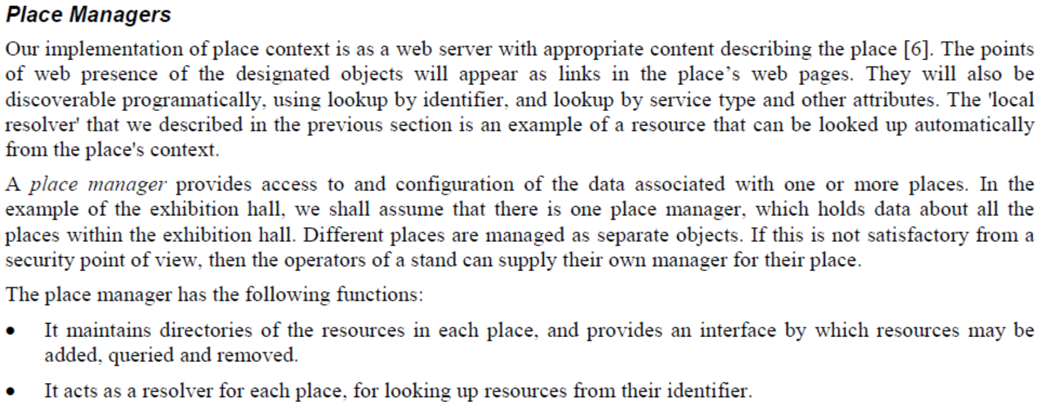 ***Place Managers*** Our implementation of place context is as a web server with appropriate content describing the place [6]. The points of web presence of the designated objects will appear as links in the place's web pages. They will also be discoverable programatically, using lookup by identifier, and lookup by service type and other attributes. The 'local resolver' that we described in the previous section is an example of a resource that can be looked up automatically from the place's context. A *place manager* provides access to and configuration of the data associated with one or more places. In the example of the exhibition hall, we shall assume that there is one place manager, which holds data about all the places within the exhibition hall. Different places are managed as separate objects. If this is not satisfactory from a security point of view, then the operators of a stand can supply their own manager for their place. The place manager has the following functions: • It maintains directories of the resources in each place, and provides an interface by which resources may be added, queried and removed. • It acts as a resolver for each place, for looking up resources from their identifier. Cooltown Developer's Network, People, Place, and Things: web presence for the real world (HP000025–34) (HP000025–34) at 8 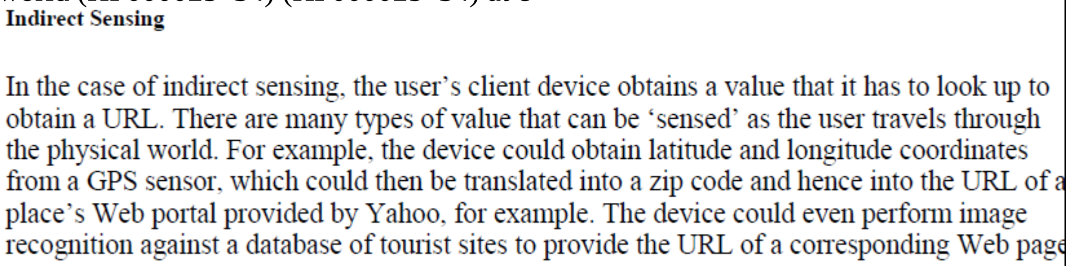 **Indirect Sensing** In the case of indirect sensing, the user's client device obtains a value that it has to look up to obtain a URL. There are many types of value that can be 'sensed' as the user travels through the physical world. For example, the device could obtain latitude and longitude coordinates from a GPS sensor, which could then be translated into a zip code and hence into the URL of a place's Web portal provided by Yahoo, for example. The device could even perform image recognition against a database of tourist sites to provide the URL of a corresponding Web pag Cooltown Developer's Network, People, Place, and Things: web presence for the real world (HP000025–34) (HP000025–34) at 23

Element	'278 Claim Recitation	Exemplary Citations in Reference
		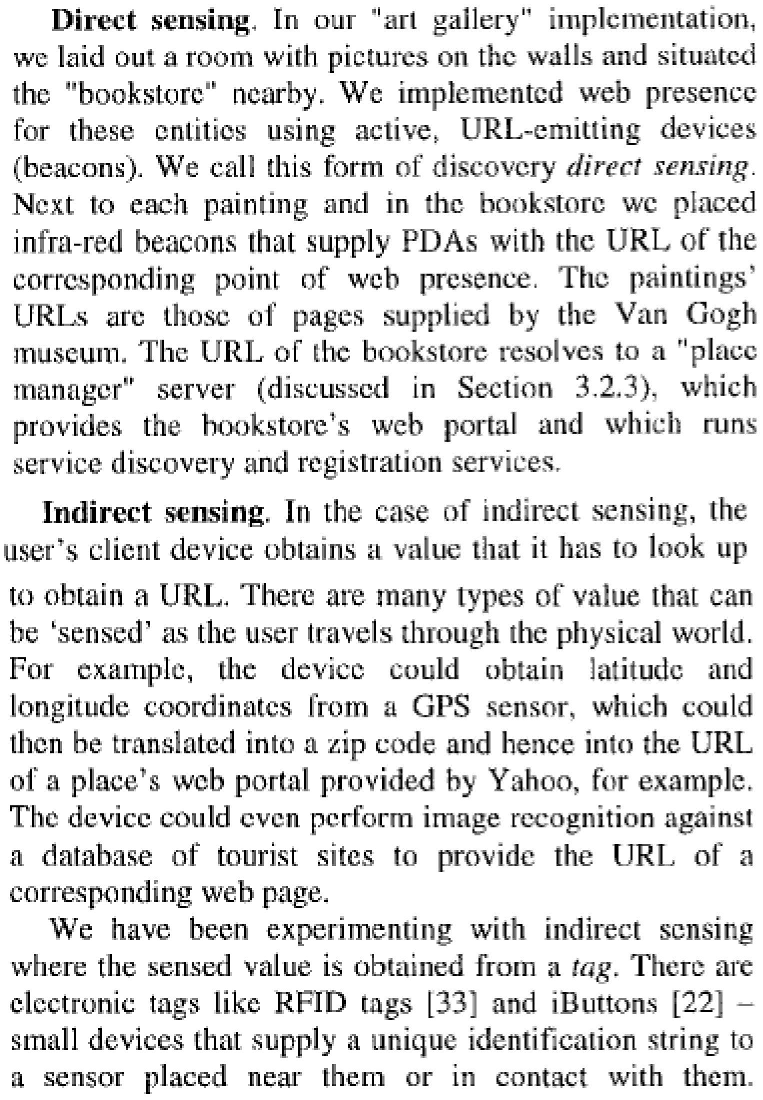 **Direct sensing**. In our "art gallery" implementation, we laid out a room with pictures on the walls and situated the "bookstore" nearby. We implemented web presence for these entities using active, URL-emitting devices (beacons). We call this form of discovery *direct sensing*. Next to each painting and in the bookstore we placed infra-red beacons that supply PDAs with the URL of the corresponding point of web presence. The paintings' URLs are those of pages supplied by the Van Gogh museum. The URL of the bookstore resolves to a "place manager" server (discussed in Section 3.2.3), which provides the bookstore's web portal and which runs service discovery and registration services. **Indirect sensing**. In the case of indirect sensing, the user's client device obtains a value that it has to look up to obtain a URL. There are many types of value that can be 'sensed' as the user travels through the physical world. For example, the device could obtain latitude and longitude coordinates from a GPS sensor, which could then be translated into a zip code and hence into the URL of a place's web portal provided by Yahoo, for example. The device could even perform image recognition against a database of tourist sites to provide the URL of a corresponding web page. We have been experimenting with indirect sensing where the sensed value is obtained from a *tag*. There are electronic tags like RFID tags [33] and iButtons [22] – small devices that supply a unique identification string to a sensor placed near them or in contact with them.

Element	'278 Claim Recitation	Exemplary Citations in Reference
		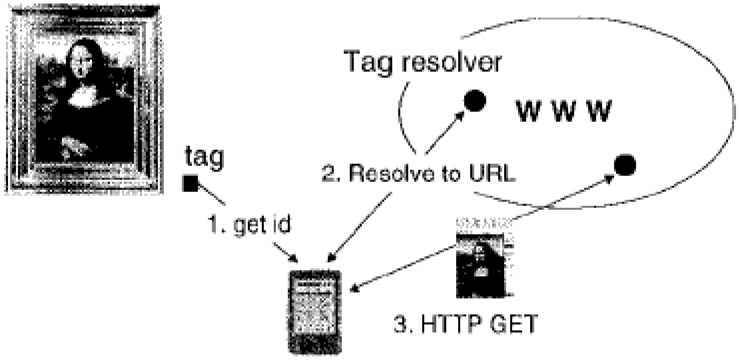 **Figure 4. Indirect sensing: the web presence for a painting.** Barcodes can be used instead of an electronic tag. Figure 4 shows a different art gallery demonstrator in which a PDA client senses a painting's tag. A correlated web page is resolved from the tag identifier and the client presents it to the user. We have produced a prototype implementation for sensing tagged things and displaying a corresponding context-dependent web page. The sensing device knows the URL of a *resolver*. When the client senses a tag, it sends the identifier to the resolver. If the resolver has an entry for the identifier, it sends back the corresponding URL. The user may choose a resolver that reflects his or her personal requirements, or those of a group of users with interests or tasks in common. For example, visitors to an art gallery could be offered resolvers that provide the URLs of pages in their native language, or pages provided by the art school at which they study.

Element	'278 Claim Recitation	Exemplary Citations in Reference
		Cooltown Developer's Network, People, Place, and Things: web presence for the real world (HP000025 – 34) at 8 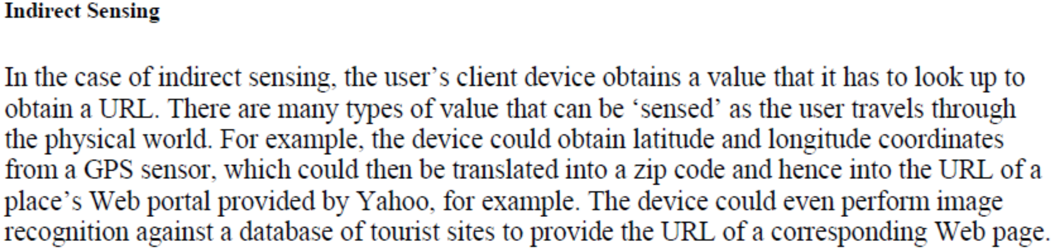 **Indirect Sensing** In the case of indirect sensing, the user's client device obtains a value that it has to look up to obtain a URL. There are many types of value that can be 'sensed' as the user travels through the physical world. For example, the device could obtain latitude and longitude coordinates from a GPS sensor, which could then be translated into a zip code and hence into the URL of a place's Web portal provided by Yahoo, for example. The device could even perform image recognition against a database of tourist sites to provide the URL of a corresponding Web page. To the extent it is found that CoolTown does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that CoolTown does not anticipate this claim, CoolTown renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine CoolTown with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1E.i]	the image,	CoolTown discloses this limitation. *See* limitation [1E] above. To the extent it is found that CoolTown does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that CoolTown does not anticipate this claim, CoolTown renders it obvious, either alone, in combination

Element	'278 Claim Recitation	Exemplary Citations in Reference
		with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine CoolTown with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1E.ii]	the symbol information,	CoolTown discloses this limitation. *See* limitation [1E] above. To the extent it is found that CoolTown does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that CoolTown does not anticipate this claim, CoolTown renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine CoolTown with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1E.iii]	and a query of a database storing target object information associated with a plurality of target objects	CoolTown discloses this limitation. *See* limitation [1E] above.

Element	'278 Claim Recitation	Exemplary Citations in Reference
	including the first target object;	To the extent it is found that CoolTown does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that CoolTown does not anticipate this claim, CoolTown renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine CoolTown with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1F]	receiving, via an address, first target object information associated with the first target object,	CoolTown discloses receiving, via an address, first target object information associated with the first target object. A Web-Based Nomadic Computing System (HP000074 – 87) at Fig. 2

Element	'278 Claim Recitation	Exemplary Citations in Reference
		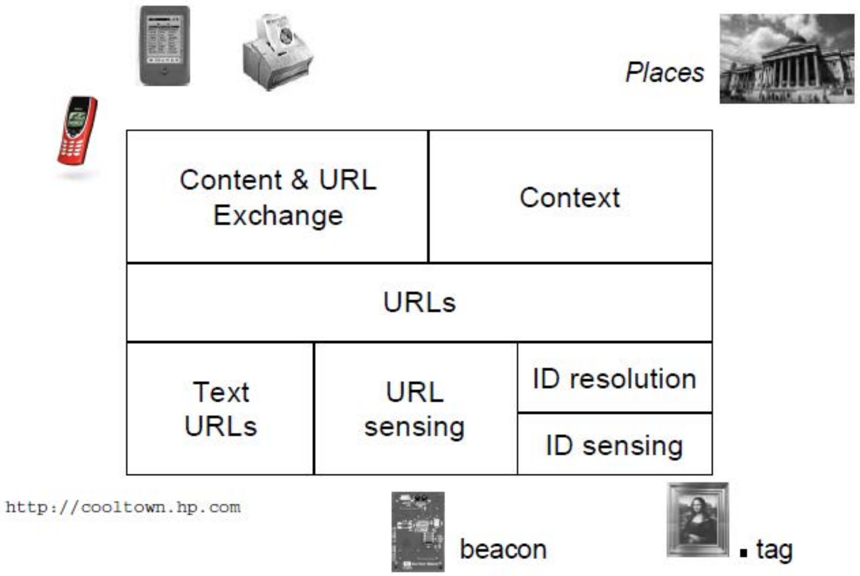 Figure 2. Component relationships. Sources of URLs, including keyboards, beacons, and resolv tags point to context or content. A Web-Based Nomadic Computing System (HP000074 – 87) at 4

Element	'278 Claim Recitation	Exemplary Citations in Reference
		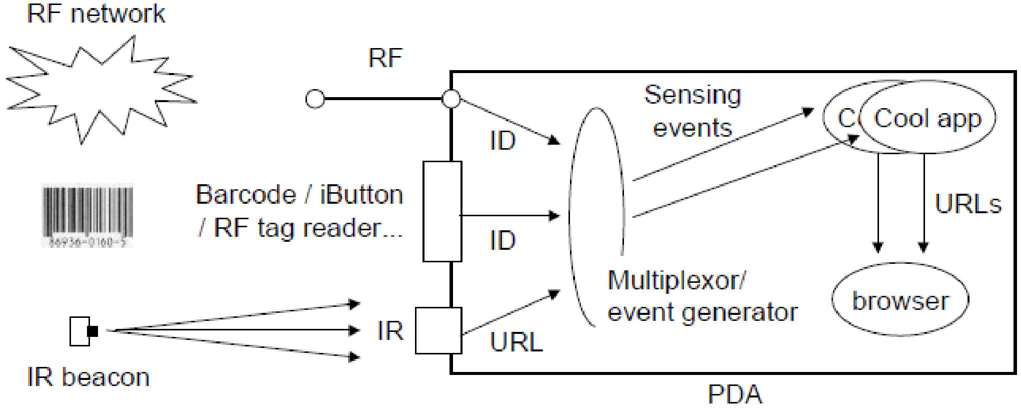 Figure 3. A sketch of a PDA equipped for sensing URLs and identifiers. A Web-Based Nomadic Computing System (HP000074 – 87) at 5 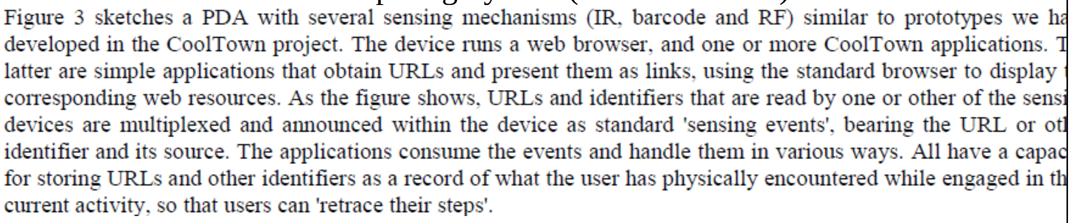Figure 3 sketches a PDA with several sensing mechanisms (IR, barcode and RF) similar to prototypes we ha developed in the CoolTown project. The device runs a web browser, and one or more CoolTown applications. T latter are simple applications that obtain URLs and present them as links, using the standard browser to display t corresponding web resources. As the figure shows, URLs and identifiers that are read by one or other of the sensi devices are multiplexed and announced within the device as standard 'sensing events', bearing the URL or ot identifier and its source. The applications consume the events and handle them in various ways. All have a capac for storing URLs and other identifiers as a record of what the user has physically encountered while engaged in th current activity, so that users can 'retrace their steps'. A Web-Based Nomadic Computing System (HP000074 – 87) at 5 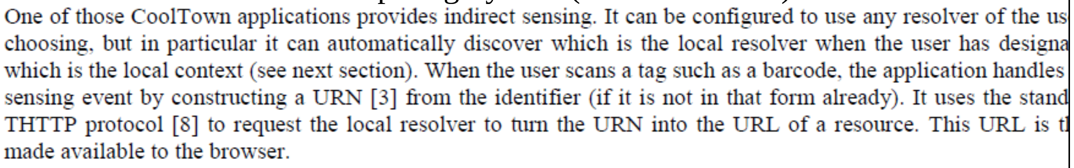One of those CoolTown applications provides indirect sensing. It can be configured to use any resolver of the us choosing, but in particular it can automatically discover which is the local resolver when the user has designa which is the local context (see next section). When the user scans a tag such as a barcode, the application handles sensing event by constructing a URN [3] from the identifier (if it is not in that form already). It uses the stand THTTP protocol [8] to request the local resolver to turn the URN into the URL of a resource. This URL is t made available to the browser. A Web-Based Nomadic Computing System (HP000074 – 87) at 6

Element	'278 Claim Recitation	Exemplary Citations in Reference
		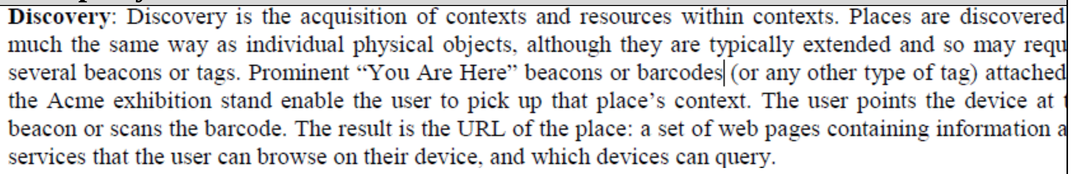 **Discovery**: Discovery is the acquisition of contexts and resources within contexts. Places are discovered much the same way as individual physical objects, although they are typically extended and so may requ several beacons or tags. Prominent "You Are Here" beacons or barcodes (or any other type of tag) attached the Acme exhibition stand enable the user to pick up that place's context. The user points the device at beacon or scans the barcode. The result is the URL of the place: a set of web pages containing information a services that the user can browse on their device, and which devices can query. Cooltown Developer's Network, People, Place, and Things: web presence for the real world (HP000025–34) at 5 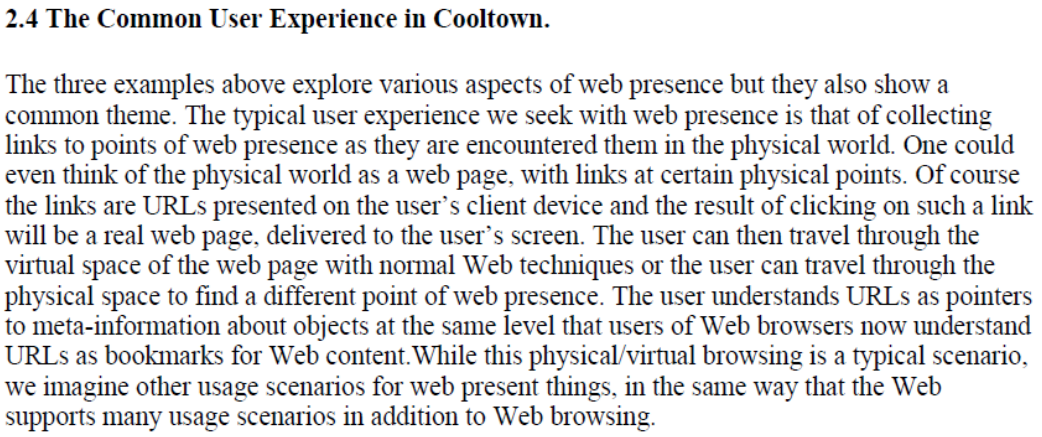 **2.4 The Common User Experience in Cooltown.** The three examples above explore various aspects of web presence but they also show a common theme. The typical user experience we seek with web presence is that of collecting links to points of web presence as they are encountered them in the physical world. One could even think of the physical world as a web page, with links at certain physical points. Of course the links are URLs presented on the user's client device and the result of clicking on such a link will be a real web page, delivered to the user's screen. The user can then travel through the virtual space of the web page with normal Web techniques or the user can travel through the physical space to find a different point of web presence. The user understands URLs as pointers to meta-information about objects at the same level that users of Web browsers now understand URLs as bookmarks for Web content.While this physical/virtual browsing is a typical scenario, we imagine other usage scenarios for web present things, in the same way that the Web supports many usage scenarios in addition to Web browsing. Cooltown Developer's Network, People, Place, and Things: web presence for the real world (HP000025–34) (HP000025–34) at 23

Element	'278 Claim Recitation	Exemplary Citations in Reference
		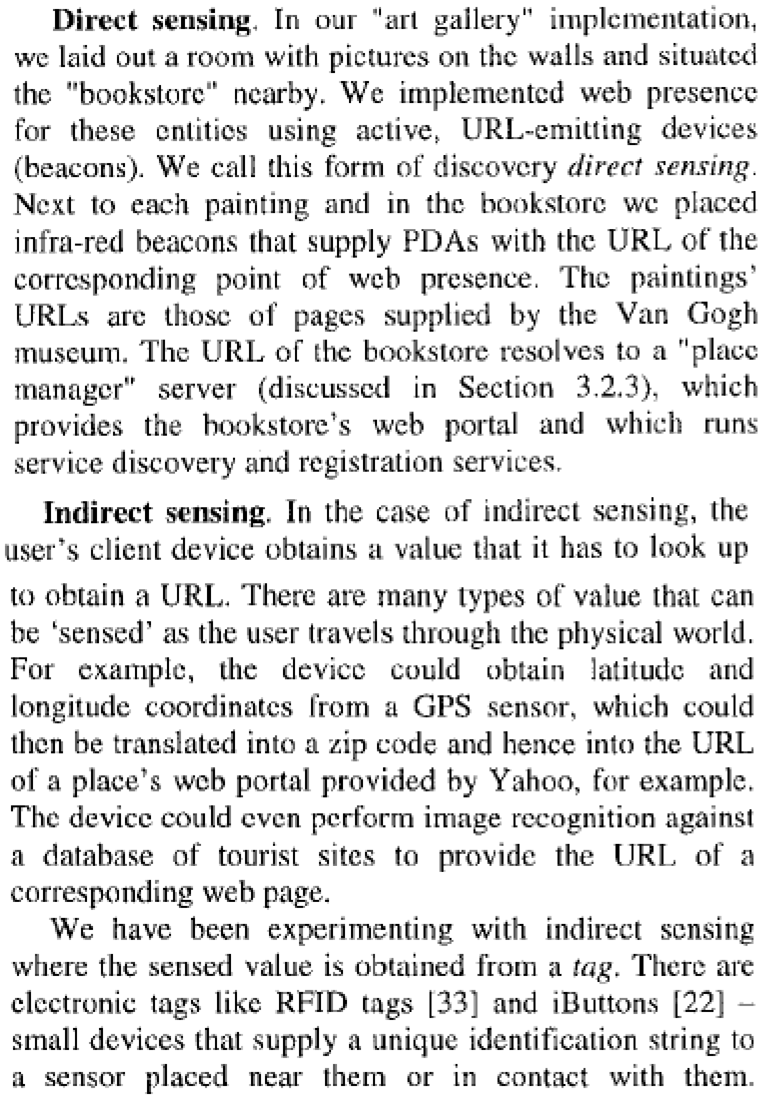 **Direct sensing**. In our "art gallery" implementation, we laid out a room with pictures on the walls and situated the "bookstore" nearby. We implemented web presence for these entities using active, URL-emitting devices (beacons). We call this form of discovery *direct sensing*. Next to each painting and in the bookstore we placed infra-red beacons that supply PDAs with the URL of the corresponding point of web presence. The paintings' URLs are those of pages supplied by the Van Gogh museum. The URL of the bookstore resolves to a "place manager" server (discussed in Section 3.2.3), which provides the bookstore's web portal and which runs service discovery and registration services. **Indirect sensing**. In the case of indirect sensing, the user's client device obtains a value that it has to look up to obtain a URL. There are many types of value that can be 'sensed' as the user travels through the physical world. For example, the device could obtain latitude and longitude coordinates from a GPS sensor, which could then be translated into a zip code and hence into the URL of a place's web portal provided by Yahoo, for example. The device could even perform image recognition against a database of tourist sites to provide the URL of a corresponding web page. We have been experimenting with indirect sensing where the sensed value is obtained from a *tag*. There are electronic tags like RFID tags [33] and iButtons [22] – small devices that supply a unique identification string to a sensor placed near them or in contact with them.

Element	'278 Claim Recitation	Exemplary Citations in Reference
		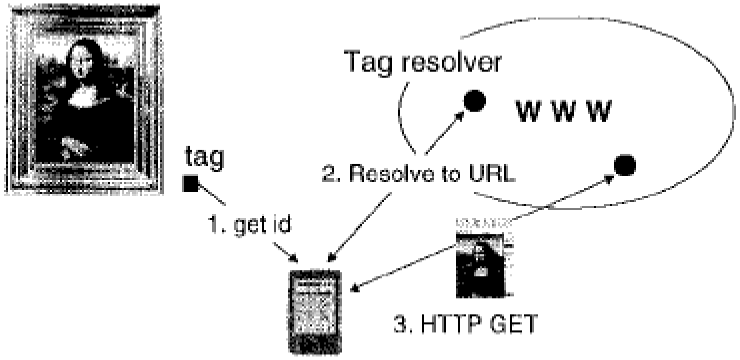 **Figure 4. Indirect sensing: the web presence for a painting.** Barcodes can be used instead of an electronic tag. Figure 4 shows a different art gallery demonstrator in which a PDA client senses a painting's tag. A correlated web page is resolved from the tag identifier and the client presents it to the user. We have produced a prototype implementation for sensing tagged things and displaying a corresponding context-dependent web page. The sensing device knows the URL of a *resolver*. When the client senses a tag, it sends the identifier to the resolver. If the resolver has an entry for the identifier, it sends back the corresponding URL. The user may choose a resolver that reflects his or her personal requirements, or those of a group of users with interests or tasks in common. For example, visitors to an art gallery could be offered resolvers that provide the URLs of pages in their native language, or pages provided by the art school at which they study.

Element	'278 Claim Recitation	Exemplary Citations in Reference
		A Web-Based Nomadic Computing System (HP000074 – 87) at 6 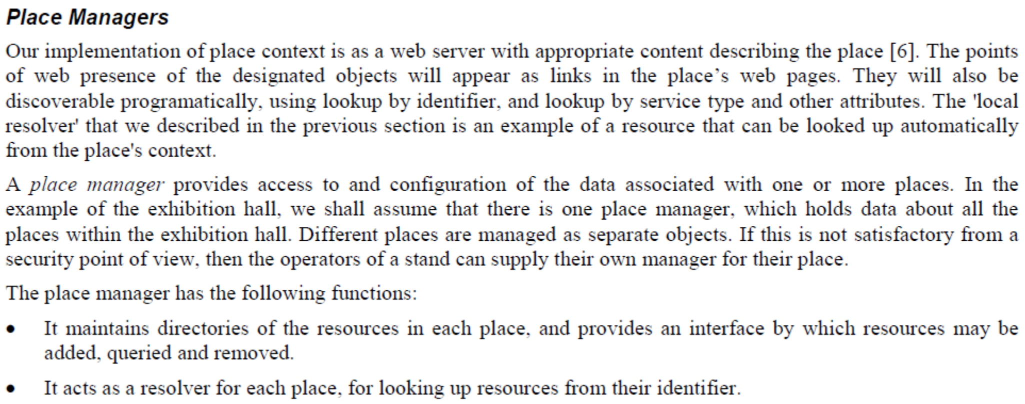 ***Place Managers*** Our implementation of place context is as a web server with appropriate content describing the place [6]. The points of web presence of the designated objects will appear as links in the place's web pages. They will also be discoverable programatically, using lookup by identifier, and lookup by service type and other attributes. The 'local resolver' that we described in the previous section is an example of a resource that can be looked up automatically from the place's context. A *place manager* provides access to and configuration of the data associated with one or more places. In the example of the exhibition hall, we shall assume that there is one place manager, which holds data about all the places within the exhibition hall. Different places are managed as separate objects. If this is not satisfactory from a security point of view, then the operators of a stand can supply their own manager for their place. The place manager has the following functions: • It maintains directories of the resources in each place, and provides an interface by which resources may be added, queried and removed. • It acts as a resolver for each place, for looking up resources from their identifier. To the extent it is found that CoolTown does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that CoolTown does not anticipate this claim, CoolTown renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine CoolTown with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1F.i]	wherein the first target object information comprises a response regarding the validity of the	Cooltown Developer's Network, People, Place, and Things: web presence for the real world (HP000025–34) (HP000025–34) at 20 – 21

Element	'278 Claim Recitation	Exemplary Citations in Reference
	document.	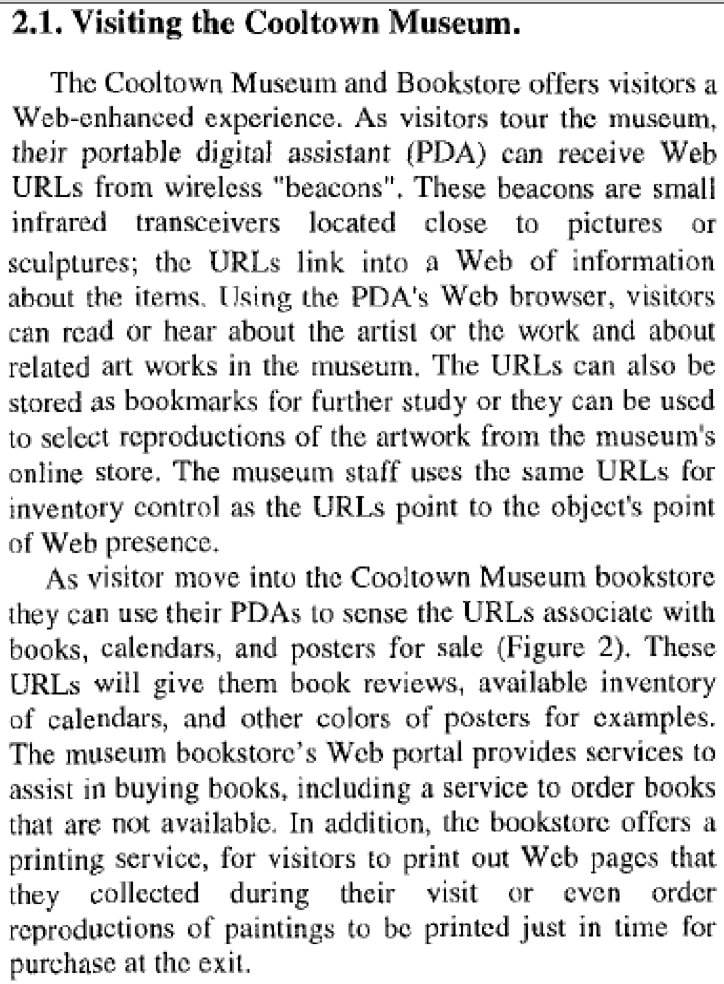 **2.1. Visiting the Cooltown Museum.** The Cooltown Museum and Bookstore offers visitors a Web-enhanced experience. As visitors tour the museum, their portable digital assistant (PDA) can receive Web URLs from wireless "beacons". These beacons are small infrared transceivers located close to pictures or sculptures; the URLs link into a Web of information about the items. Using the PDA's Web browser, visitors can read or hear about the artist or the work and about related art works in the museum. The URLs can also be stored as bookmarks for further study or they can be used to select reproductions of the artwork from the museum's online store. The museum staff uses the same URLs for inventory control as the URLs point to the object's point of Web presence. As visitor move into the Cooltown Museum bookstore they can use their PDAs to sense the URLs associate with books, calendars, and posters for sale (Figure 2). These URLs will give them book reviews, available inventory of calendars, and other colors of posters for examples. The museum bookstore's Web portal provides services to assist in buying books, including a service to order books that are not available. In addition, the bookstore offers a printing service, for visitors to print out Web pages that they collected during their visit or even order reproductions of paintings to be printed just in time for purchase at the exit. A Web-Based Nomadic Computing System (HP000074 – 87) at 10

Element	'278 Claim Recitation	Exemplary Citations in Reference
		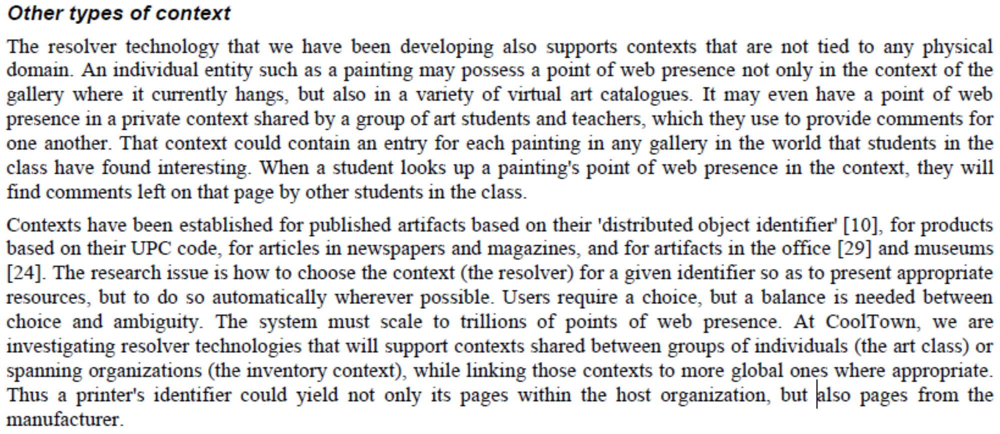 ***Other types of context*** The resolver technology that we have been developing also supports contexts that are not tied to any physical domain. An individual entity such as a painting may possess a point of web presence not only in the context of the gallery where it currently hangs, but also in a variety of virtual art catalogues. It may even have a point of web presence in a private context shared by a group of art students and teachers, which they use to provide comments for one another. That context could contain an entry for each painting in any gallery in the world that students in the class have found interesting. When a student looks up a painting's point of web presence in the context, they will find comments left on that page by other students in the class. Contexts have been established for published artifacts based on their 'distributed object identifier' [10], for products based on their UPC code, for articles in newspapers and magazines, and for artifacts in the office [29] and museums [24]. The research issue is how to choose the context (the resolver) for a given identifier so as to present appropriate resources, but to do so automatically wherever possible. Users require a choice, but a balance is needed between choice and ambiguity. The system must scale to trillions of points of web presence. At CoolTown, we are investigating resolver technologies that will support contexts shared between groups of individuals (the art class) or spanning organizations (the inventory context), while linking those contexts to more global ones where appropriate. Thus a printer's identifier could yield not only its pages within the host organization, but also pages from the manufacturer. A Web-Based Nomadic Computing System (HP000074 – 87) at Fig. 2

Element	'278 Claim Recitation	Exemplary Citations in Reference
		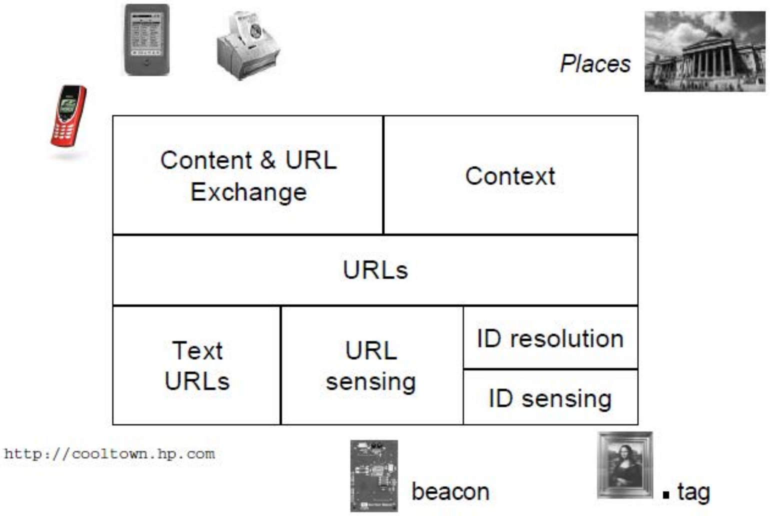 Figure 2. Component relationships. Sources of URLs, including keyboards, beacons, and resolved tags point to context or content. A Web-Based Nomadic Computing System (HP000074 – 87) at 4

Element	'278 Claim Recitation	Exemplary Citations in Reference
		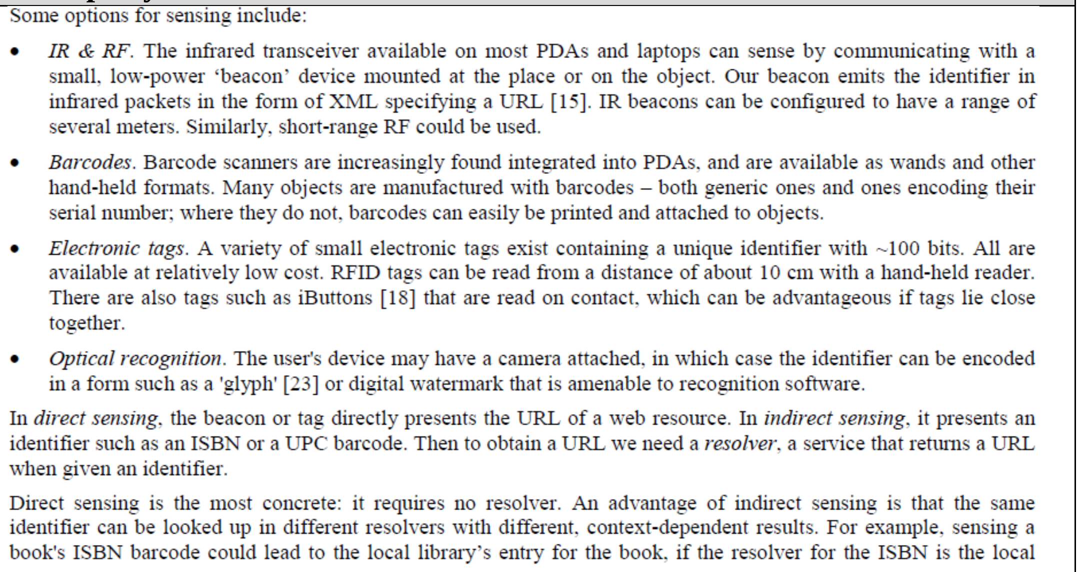 Some options for sensing include: • *IR & RF*. The infrared transceiver available on most PDAs and laptops can sense by communicating with a small, low-power 'beacon' device mounted at the place or on the object. Our beacon emits the identifier in infrared packets in the form of XML specifying a URL [15]. IR beacons can be configured to have a range of several meters. Similarly, short-range RF could be used. • *Barcodes*. Barcode scanners are increasingly found integrated into PDAs, and are available as wands and other hand-held formats. Many objects are manufactured with barcodes – both generic ones and ones encoding their serial number; where they do not, barcodes can easily be printed and attached to objects. • *Electronic tags*. A variety of small electronic tags exist containing a unique identifier with ~100 bits. All are available at relatively low cost. RFID tags can be read from a distance of about 10 cm with a hand-held reader. There are also tags such as iButtons [18] that are read on contact, which can be advantageous if tags lie close together. • *Optical recognition*. The user's device may have a camera attached, in which case the identifier can be encoded in a form such as a 'glyph' [23] or digital watermark that is amenable to recognition software. In *direct sensing*, the beacon or tag directly presents the URL of a web resource. In *indirect sensing*, it presents an identifier such as an ISBN or a UPC barcode. Then to obtain a URL we need a *resolver*, a service that returns a URL when given an identifier. Direct sensing is the most concrete: it requires no resolver. An advantage of indirect sensing is that the same identifier can be looked up in different resolvers with different, context-dependent results. For example, sensing a book's ISBN barcode could lead to the local library's entry for the book, if the resolver for the ISBN is the local To the extent it is found that CoolTown does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that CoolTown does not anticipate this claim, CoolTown renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine CoolTown with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

2. Claim 4

Claim	'278 Claim Recitation	Exemplary Citations in Reference
4	The method of claim 3, wherein the image includes a representation of the code.	CoolTown discloses the method of claim 3, wherein the image includes a representation of the code. A Web-Based Nomadic Computing System at (HP000074 – 87) 3 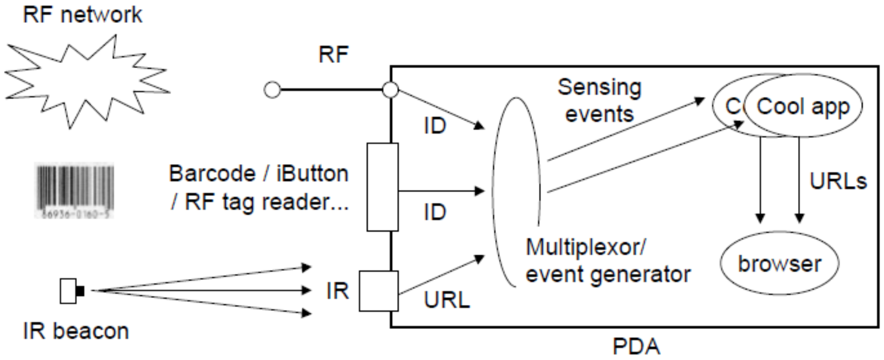 Figure 3. A sketch of a PDA equipped for sensing URLs and identifiers. A Web-Based Nomadic Computing System (HP000074 – 87) at 4

		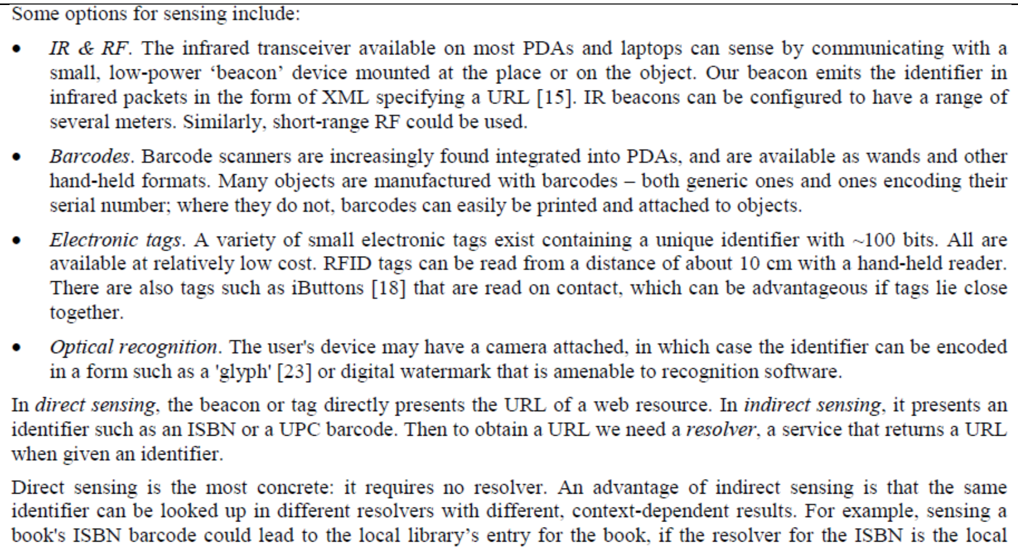 Some options for sensing include: • *IR & RF*. The infrared transceiver available on most PDAs and laptops can sense by communicating with a small, low-power 'beacon' device mounted at the place or on the object. Our beacon emits the identifier in infrared packets in the form of XML specifying a URL [15]. IR beacons can be configured to have a range of several meters. Similarly, short-range RF could be used. • *Barcodes*. Barcode scanners are increasingly found integrated into PDAs, and are available as wands and other hand-held formats. Many objects are manufactured with barcodes – both generic ones and ones encoding their serial number; where they do not, barcodes can easily be printed and attached to objects. • *Electronic tags*. A variety of small electronic tags exist containing a unique identifier with ~100 bits. All are available at relatively low cost. RFID tags can be read from a distance of about 10 cm with a hand-held reader. There are also tags such as iButtons [18] that are read on contact, which can be advantageous if tags lie close together. • *Optical recognition*. The user's device may have a camera attached, in which case the identifier can be encoded in a form such as a 'glyph' [23] or digital watermark that is amenable to recognition software. In *direct sensing*, the beacon or tag directly presents the URL of a web resource. In *indirect sensing*, it presents an identifier such as an ISBN or a UPC barcode. Then to obtain a URL we need a *resolver*, a service that returns a URL when given an identifier. Direct sensing is the most concrete: it requires no resolver. An advantage of indirect sensing is that the same identifier can be looked up in different resolvers with different, context-dependent results. For example, sensing a book's ISBN barcode could lead to the local library's entry for the book, if the resolver for the ISBN is the local Cooltown Developer's Network, People, Place, and Things: web presence for the real world (HP000025–34) (HP000025–34) at 23 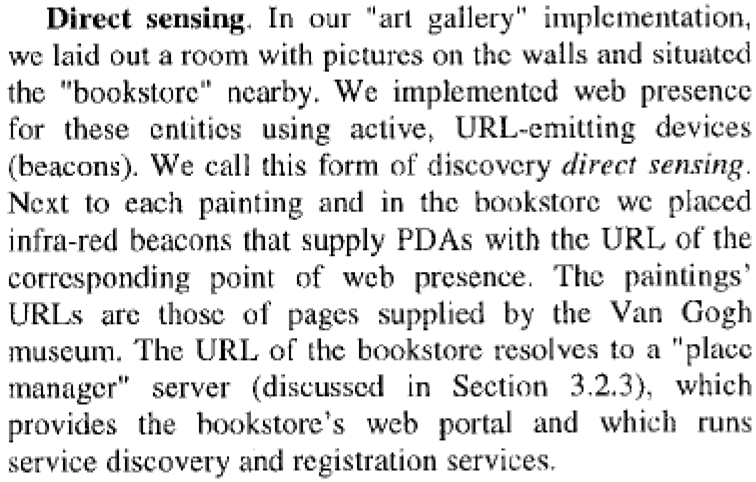 **Direct sensing**. In our "art gallery" implementation, we laid out a room with pictures on the walls and situated the "bookstore" nearby. We implemented web presence for these entities using active, URL-emitting devices (beacons). We call this form of discovery *direct sensing*. Next to each painting and in the bookstore we placed infra-red beacons that supply PDAs with the URL of the corresponding point of web presence. The paintings' URLs are those of pages supplied by the Van Gogh museum. The URL of the bookstore resolves to a "place manager" server (discussed in Section 3.2.3), which provides the bookstore's web portal and which runs service discovery and registration services.

		Indirect sensing. In the case of indirect sensing, the user's client device obtains a value that it has to look up to obtain a URL. There are many types of value that can be 'sensed' as the user travels through the physical world. For example, the device could obtain latitude and longitude coordinates from a GPS sensor, which could then be translated into a zip code and hence into the URL of a place's web portal provided by Yahoo, for example. The device could even perform image recognition against a database of tourist sites to provide the URL of a corresponding web page. We have been experimenting with indirect sensing where the sensed value is obtained from a *tag*. There are electronic tags like RFID tags [33] and iButtons [22] – small devices that supply a unique identification string to a sensor placed near them or in contact with them. 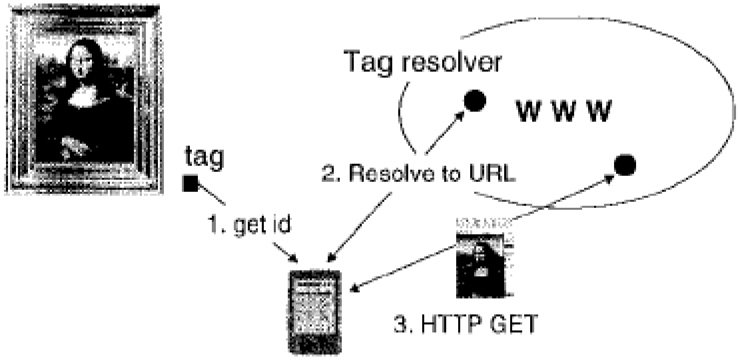 **Figure 4. Indirect sensing: the web presence for a painting.**

		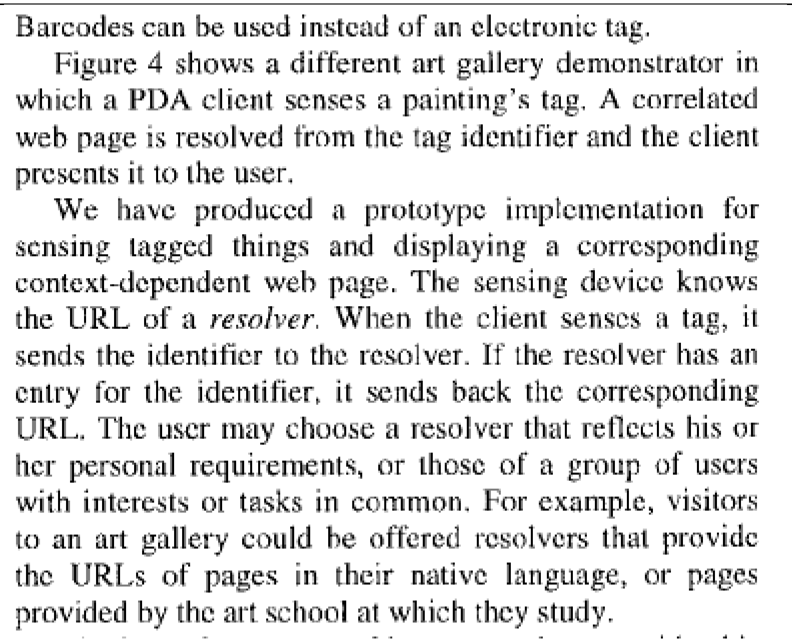 Barcodes can be used instead of an electronic tag. Figure 4 shows a different art gallery demonstrator in which a PDA client senses a painting's tag. A correlated web page is resolved from the tag identifier and the client presents it to the user. We have produced a prototype implementation for sensing tagged things and displaying a corresponding context-dependent web page. The sensing device knows the URL of a *resolver*. When the client senses a tag, it sends the identifier to the resolver. If the resolver has an entry for the identifier, it sends back the corresponding URL. The user may choose a resolver that reflects his or her personal requirements, or those of a group of users with interests or tasks in common. For example, visitors to an art gallery could be offered resolvers that provide the URLs of pages in their native language, or pages provided by the art school at which they study. To the extent it is found that CoolTown does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that CoolTown does not anticipate this claim, CoolTown renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine CoolTown with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

3. Claim 5

Claim	'278 Claim Recitation	Exemplary Citations in Reference
5	The method of claim 3, further comprising detecting whether there is a match within a specified tolerance and returning a message indicating the absence of a match within a specified tolerance.	To the extent it is found that CoolTown does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that CoolTown does not anticipate this claim, CoolTown renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine CoolTown with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

Nantworks, LLC v. Bank of America Corporation
Central District of California, No. 2:20-cv-7872
Invalidity Chart for U.S. Patent 7,881,529 based on
U.S. Patent No. 6,947,517 to Rhoads ("Rhoads")

Exhibit G-3

Bank of America Corporation ("BofA" or "Defendant") provides this chart subject to all reservations, objections, statements, and disclaimers set forth herein and in BofA's Final Invalidity Contentions Cover Pleading, as well as any amendment, supplement, or modification thereof, which are incorporated herein by reference in their entirety.

On information and belief, the application that matured into Rhoads was filed on May 15, 2000, claiming the priority date of May 19, 1999. Rhoads published at least as of September 20, 2005, and is available as prior art at least under 35 U.S.C. § 102(e).

As shown in the chart below, Rhoads anticipates asserted claims 1, 4 and 20 of the '529 Patent. To the extent Rhoads is found not to expressly disclose certain limitations in the asserted claims, such limitations are inherent. To the extent Rhoads is found not to anticipate any asserted claims or claim elements of the '529 Patent, the reference nevertheless renders those claims or claim elements obvious under 35 U.S.C. § 103, either alone or in combination with other art identified herein or in the cover pleading of BofA's Final Invalidity Contentions, for the reasons identified in the cover pleading. These Final Invalidity Contentions are not an admission by BofA that the accused products, including any current or past versions of the accused products, are covered by, or infringe the asserted claims, but are based instead on the recognition that if the claims are interpreted to be broad enough to encompass the accused products, the claims must also be construed to have that same scope when considering whether they are invalid.

The following chart is based on the Court's claim construction rulings (Dkt. Nos. 145, 153, 236), and to the extent not contradictory, Plaintiffs' apparent interpretations of the claims in Plaintiff's "Preliminary" Final Infringement Contentions. BofA notes that in many instances, Plaintiff's "Preliminary" Final Infringement Contentions ignore and/or contradict the Court's claim construction rulings. BofA does not accept the assumptions concerning the scope and meaning implicit in Plaintiff's Preliminary Final Infringement Contentions, to the extent those assumptions are discernible, and reserves the right to challenge Plaintiff's proposed (or implied) constructions. Furthermore, Plaintiffs' "Preliminary" Final Contentions fail to comply with the level of specificity required by SPR 2.1.2 and 2.1.3, and include improper reservation of rights to supplement with "additional ways in which BoA Accused Products infringe." Any alleged deficient discovery is of Plaintiffs own making. Most if not all of the deficiencies in Plaintiffs' Preliminary Final Contentions have been documented by Defendants through multiple letters since at least May 18, 2021 and up to Bank of America's most recent letter of August 19, 2023. Plaintiffs' shifting infringement contentions, deficiencies in its infringement contentions, and changing claim scope have prejudiced BofA from being able to reasonably prepare its defenses, including its invalidity contentions. BofA also reserves the right to revise and supplement these charts if and when Plaintiff is permitted to provide revised or supplemental Infringement Contentions. Citations given in the chart below are merely representative of the respective elements to provide notice to Plaintiffs' of BofA's invalidity theories and are not meant to be exhaustive.

1. Claim 1

Element	'529 Claim Recitation	Exemplary Citations in Reference
[1Pre]	A system comprising:	Rhoads discloses a system Rhoads at Abstract • "A cell phone is equipped with a 2D optical sensor, enabling a variety of applications. For example, such a phone may also be provided with a digital watermark decoder, permitting decoding of steganographic data on imaged objects. Movement of a phone may be inferred by sensing movement of an imaged pattern across the optical sensor's field of view, allowing use of the phone as a gestural input device through which a user can signal instructions to a computer-based process. A variety of other arrangements by which electronic devices can interact with the physical world are also detailed, e.g., involving sensing and responding to digital watermarks, bar codes, RFIDs, etc." Rhoads at 5:54-56. • "Basically, the technology detailed in this disclosure may be regarded as enhanced systems by which users can interact with computer-based devices." Rhoads at 9:35–41 • "The Steganographic decoding process may entail three Steps. In the first, the object is located. In the Second, the objects orientation is discerned. In the third, the Bedoop data is extracted from the image data corresponding to the Bedoop object." Rhoads at Fig. 2 (and accompanying description)

Element	'529 Claim Recitation	Exemplary Citations in Reference
		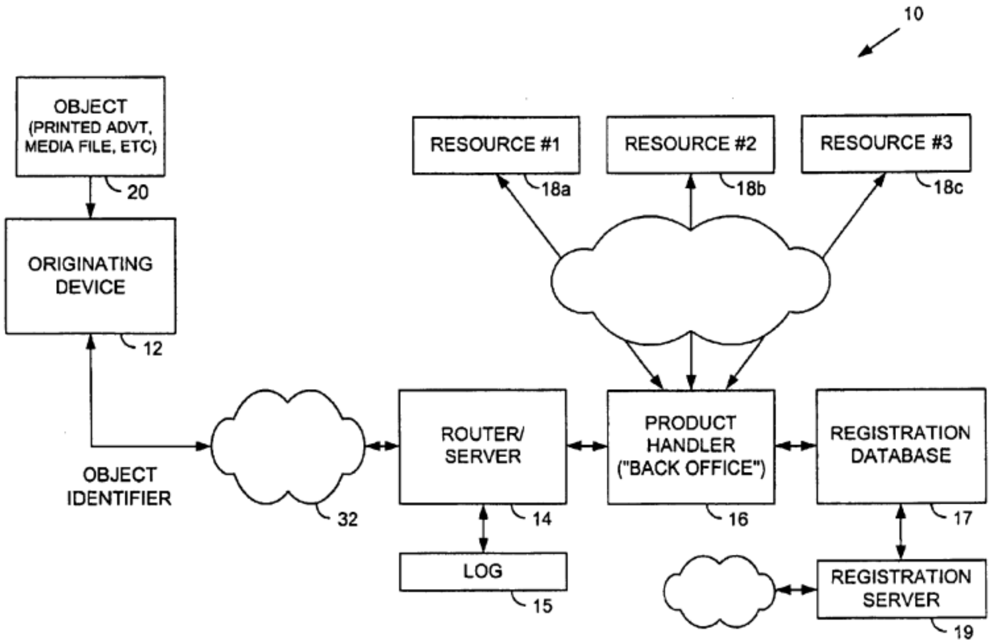 FIG. 2
[1A]	a camera that captures an image;	Rhoads discloses a camera that captures an image. Rhoads at Fig. 2 (and accompanying description)

Element	'529 Claim Recitation	Exemplary Citations in Reference
		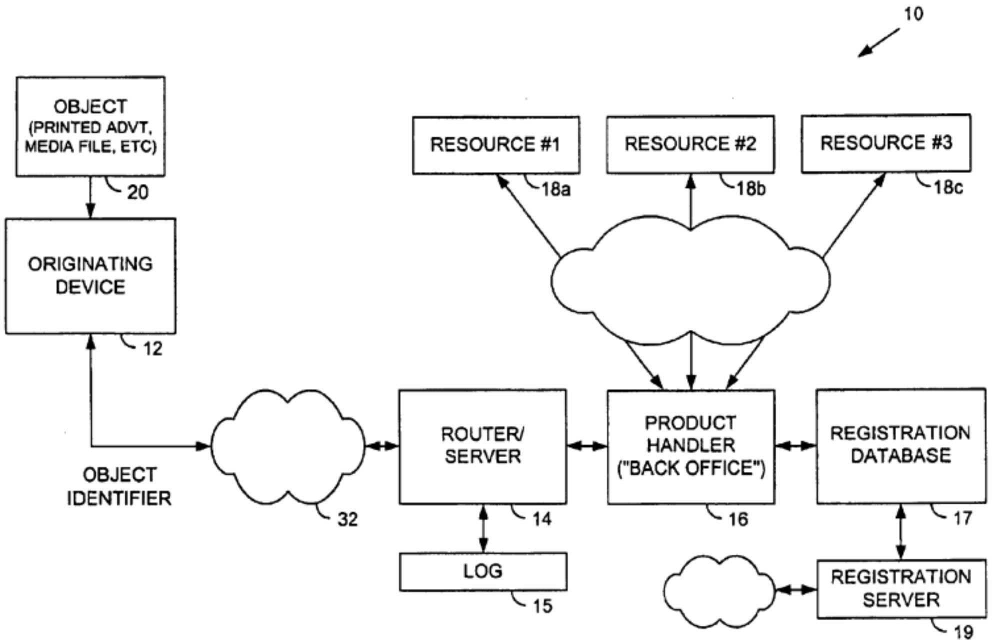 FIG. 2 Rhoads at 3:41–51 • "According to another aspect, the invention includes a promotional method comprising: (a) presenting an object within the field of view of an optical sensor device, the object being selected from the list consisting of a retail product, packaging for a retail product, or printed advertising; (b) acquiring optical data corresponding to the object; (c) decoding plural-bit digital data from the optical data; (d) submitting at least some of said decoded data to a remote computer; and (e) determining at the remote computer whether a prize should be awarded in response to submission of said decoded data." Rhoads at 7:24-42 • "Device 12 interacts with an object 20. The object can be electronic or not.

Element	'529 Claim Recitation	Exemplary Citations in Reference
		Electronic objects can include computer files, representations of audio, video, or still imagery (e.g., files or in streaming form), etc. Non-electronic objects can include physical objects such as newspapers, magazine pages, posters, product packaging, event tickets, credit cards, paper currency, etc. Non-electronic objects can also include sounds produced by loudspeakers. When used with nonelectronic objects, device 12 (FIG. 2) typically includes some form of sensor or transducer 22 to produce electronic signals or data corresponding to the object. Examples include CCD- or CMOS-based optical sensors (either as part of still- or video cameras, flatbed scanners, mice, or otherwise), microphones, barcode scanners, RF ID sensors, mag stripe readers, etc. In such cases, the sensor 22 may be coupled to associated interface electronics 24, which in turn may be coupled to device driver software 26, which in turn may be coupled to one or more application programs 28." Rhoads at Fig. 11(and accompanying description)

Element	'529 Claim Recitation	Exemplary Citations in Reference
		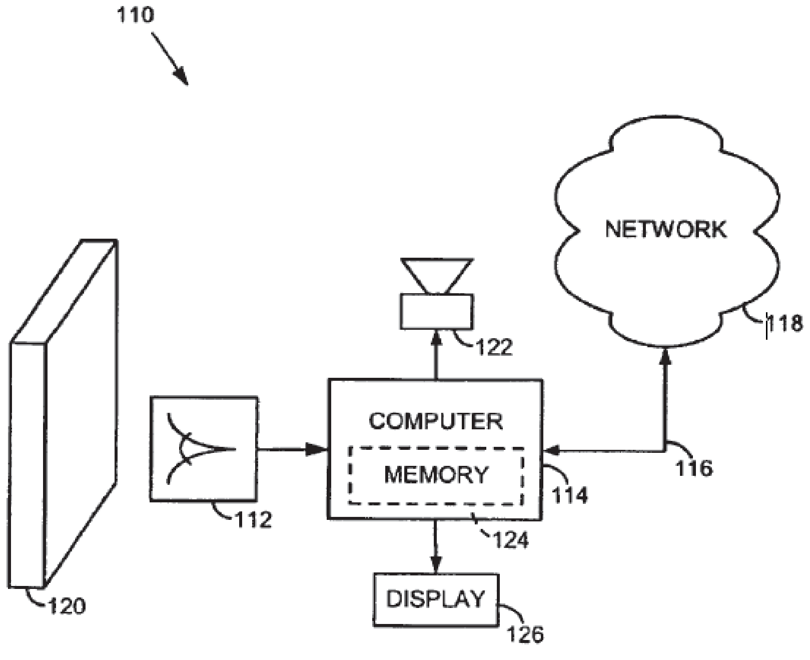 FIG. 11 Rhoads at 7:13-29 - "Referring to FIG. 11, a basic embodiment 110 of the present technology includes an optical sensor 112, a computer 114, and a network connection 116 to the internet 118. The illustrated optical sensor 112 is a digital camera having a resolution of 320 by 200 pixels (color or black and white) that stares out, grabbing frames of image data five times per second and storing same in one or more frame buffers. These frames of image data are analyzed

Element	**'529 Claim Recitation**	**Exemplary Citations in Reference**
		by a computer 114 for the presence of Bedoop data. (Essentially, Bedoop data is any form of plural-bit data encoding recognized by the system 110—data which, in many embodiments, initiates some action.) Once detected, the system responds in accordance with the detected Bedoop data (e.g., by initiating some local action, or by communication with a remote computer, such as over the internet, via an online service such as AOL, or using point-to-point dial-up communications, as with a bulletin board system)." Rhoads 8:66 – 9:12 • "In still other embodiments, the camera (or other sensor) can be equipped with one or more auxiliary fixed-focus lenses that can be selectively used, depending on the particular application. Some such embodiments have a first fixed focused lens that always overlies the sensor, with which one or more auxiliary lenses can be optically cascaded (e.g., by hinge or slide arrangements). Such arrangements are desirable, e.g., when a camera is not a dedicated Bedoop sensor but also performs other imaging tasks. When the camera is to be used for Bedoop, the auxiliary lens is positioned (e.g., flipped into) place, changing the focal length of the first lens (which may by unsuitably long for Bedoop purposes, such as infinity) to an appropriate Bedoop imaging range (such as one foot)." Rhoads at 17:46–54 • "The building can be provided with an image sensor (such as a video camera or the like), an associated Bedoop. detection system, and the proximity detector. When a user wearing the badge approaches, the proximity detector signals the camera to capture image data. The Bedoop detection system identifies the badge photograph (e.g., by clues as are described in the prior applications, or without such aids), captures optical data, and decodes same to extract the steganographically-embedded data hidden therein."

Element	'529 Claim Recitation	Exemplary Citations in Reference
		Rhoads at 46:47-53 - "The provision of image Sensors in cell phones enables other functionality. One is the capture of still or video imagery. Such image data can be compressed (typically by lossy processes such as MPEG, JPEG, or the like, implemented with dedicated hardware CODECs) and transmitted with the audio data. The Screens on Such phones can likewise be used for display of incoming image or Video data." To the extent it is found that Rhoads does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Rhoads does not anticipate this claim, Rhoads renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Rhoads with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1B]	a network-enabled device that conducts a data processing operation on at least a portion of the image to produce data, and sends the data to a service;	Rhoads discloses a network-enabled device that conducts a data processing operation on at least a portion of the image to produce data, and sends the data to a service. Rhoads at Fig. 2 (including accompanying description)

Element	'529 Claim Recitation	Exemplary Citations in Reference
		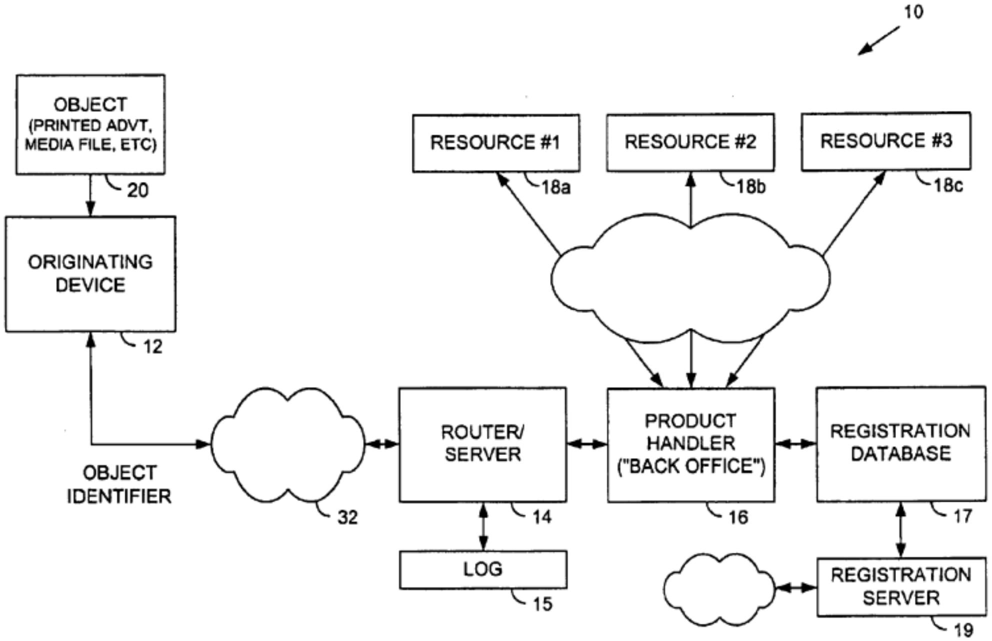 FIG. 2 Rhoads at 9:35–10:10 - "The Steganographic decoding process may entail three Steps. In the first, the object is located. In the Second, the objects orientation is discerned. In the third, the Bedoop data is extracted from the image data corresponding to the Bedoop object. The first step, object location, can be assisted by various clues. One is the placement of the object; typically the center of the image field will be a point on the object. The surrounding data can then be analyzed to try and discern the object's boundaries. Another location technique is slight movement. Although the user will

Element	'529 Claim Recitation	Exemplary Citations in Reference
		typically try to hold the object still, there will usually be some jitter of the Bedoop object within the image frame (e.g., a few pixels back and forth). Background visual clutter, in contrast, will typically be stationary. Such movement may thus be sensed and used to identify the Bedoop object from within the image data. Still another object-location clue is object shape. Many Bedoop objects are rectangular in shape (or trapezoidal as viewed by the camera). Straight edge boundaries can thus be used to define an area of likely Bedoop data. Color is a further object identification clue that may be useful in some contexts. Yet another object location clue is spatial frequency. In imaging systems with well defined focal zones, undesired visual clutter may be at focal distances that results in blurring. The Bedoop object, in contrast, will be in focus and may be characterized by fine detail. Analyzing the image data for the high frequencies associated with fine detail can be used to distinguish the intended object from others. Characteristic markings on the object (as discussed below in connection with determining object orientation) can also be sensed and used in locating the object. Once the Bedoop object has been located within the image data, masking can be applied (if desired) to eliminate image data not corresponding to the intended object." Rhoads at 3:52–60 • "According to another aspect, the invention includes a method of interacting with a magazine using a computer, the computer including an internet web browser, the method including: (a) providing a peripheral device having a

Element	'529 Claim Recitation	Exemplary Citations in Reference
		sensor; (b) positioning the peripheral device adjacent a first advertisement in the magazine to direct the web browser to a first internet address; and (c) positioning the peripheral device adjacent a second advertisement in the magazine to direct the web browser to a second internet address." Rhoads at 49:15–35 • "Referring to FIG. 2, a system 10 according to the exemplary embodiment includes an originating device 12, a router/server 14, a product handler 16, a registration database 17, and one or more remote resources 18. The originating device 12 can take many different forms, e.g., a cell phone, a personal digital assistant (e.g., a Palm Pilot), a personal computer, a barcode scanning system, etc. For expository convenience, the embodiment is described with reference to a personal computer for device 12. Device 12 interacts with an object 20. The object can be electronic or not. Electronic objects can include computer files, representations of audio, video, or still imagery (e.g., files or in streaming form), etc. Non-electronic objects can include physical objects such as newspapers, magazine pages, posters, product packaging, event tickets, credit cards, paper currency, etc. Non-electronic objects can also include sounds produced by loudspeakers. When used with nonelectronic objects, device 12 (FIG. 2) typically includes some form of sensor or transducer 22 to produce electronic signals or data corresponding to the object…." Rhoads at 10:22–30 • "Some embodiments filter the image data at some point in the process to aid in ultimate Bedoop data extraction. One use of such filtering is to mitigate

Element	'529 Claim Recitation	Exemplary Citations in Reference
		image data artifacts due to the particular optical sensor. For example, CCD arrays have regularly-spaced sensors that sample the optical image at uniformly spaced discrete points. This discrete sampling effects a transformation of the image data, leading to certain image artifacts. An appropriately configured filter can reduce some of these artifacts." Rhoads at 19:61 – 20:15 • "The coffee cup is an example of a non-planar object. Another is soft drink cans. Special issues can arise when encoding and decoding markings on such objects. For example, when sensing such an object with a camera or the like, part of the image will be out of focus due to differing distances from the camera to different parts of the can surface. While parts of an image sensed from a non-planar object, such as a can, may be out of focus, they still convey useful image data. The out of focus areas are just blurred—as if filtered by a low pass filter. But to make use of this information, a further complication must first be addressed: warping. When viewed from a camera, the planar artwork with which the can is wrapped becomes warped. Portions of the can nearest the camera appear at a nominal full scale, while areas successively further around the can curvature (as viewed from the camera) appear progressively more and more spatially compressed. Regardless of the watermarking technology being employed, the physical warp of the can's surface is likewise manifested as a warping of the encoded watermark data." Rhoads at 38:1 – 19 • "Post-It® Notes Pads of Post-It® notes, or other pads of paper, can be marked by the

Element	'529 Claim Recitation	Exemplary Citations in Reference
		manufacturer (either by texturing, watermarked tinting, ink-jet spattering, etc.) to convey steganographic data (e.g., Bedoop data). When such a note is presented to a Bedoop system, the system may launch an application that stores a snapshot of the note. More particularly, the application may mask the note-portion of the image data from the other image data, virtually re-map it to a rectangular format of standardized pixel dimensions, JPEG-compress the resulting image, and store it in a particular computer subdirectory with a name indicating the date of image acquisition, together with the color and/or size of the note. (These latter two data may be indicated by data included in the Bedoop payload.) If the color of the note is indicated by digital data (e.g, in the file name), then the image itself may be stored in grey-scale. When later recalled for display, the white image background can be flooded with color in accordance with the digital color data." Rhoads at 42:1–12 • "The optical data collected by the sensor can be processed within the peripheral's processor to extract the steganographically encoded binary Bedoop data therefrom. Or this processing burden can be undertaken by the associated computer system, with the peripheral simply processing and formatting the raw sensor data into sequential frames of image data to be output to that system. While scanning peripherals of the type described above are typically wired to an associated host system, wireless links (e.g., radio, infrared, ultrasonic, etc.) can of course be used, freeing the user from the constraint imposed by the cable." Rhoads at 45:19–55 • "Bedoop technology can be integrated into portable telecommunication

Element	'529 Claim Recitation	Exemplary Citations in Reference
		terminals, such as cell phones (manufactured, e.g., by Motorola, Nokia, Qualcomm, and others). Such a phone can be equipped with a 1D or 2D image sensor, the output of which is applied to Bedoop decoding circuitry within the phone. This decoding circuitry can be the phone's main CPU, or can be a processing circuit dedicated to Bedoop functionality. (In this as in other embodiments, the decoding can be effected by dedicated hardware, by decoding software executing on a general purpose CPU, etc.) Cell phones are already equipped with numerous features that make them well suited for Bedoop operation. One is that cell phones typically include an LCD or similar screen for display of textual or graphic information, and additionally include buttons or other controls for selecting among menu options presented on the screen (e.g., by moving a cursor). Moreover, cell phones naturally include both audio input and output devices (i.e., microphone and speaker). Still further, the protocol by which cell phones transmit data includes data identifying the phone, so that such data need not be separately encoded. And finally, cell phones obviously provide ready links to remote computer systems. Collectively, these capabilities rival those of the most fully-equipped desktop computer system. Thus, essentially all of the applications detailed elsewhere in this specification can be implemented using cell phone Bedoop systems As with the other Bedoop systems, when Bedoop data is sensed, the phone can respond to the data locally, or it can forward Same over the cellular network to a remote System (or computer network) for handling." Rhoads at 46:10–12. • "Bedoop decoding function within the phone, in other embodiments the image data can be transmitted from the phone and decoded at a remote location."

Element	'529 Claim Recitation	Exemplary Citations in Reference
		Rhoads at Fig. 3 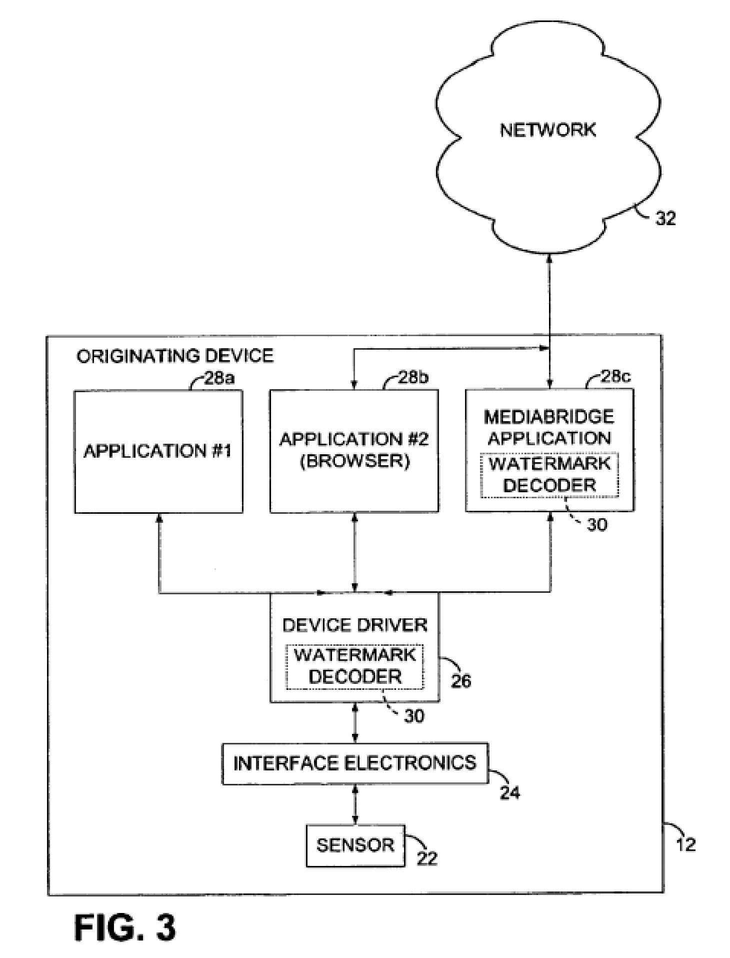 FIG. 3 Rhoads at 46:10–12.

Element	'529 Claim Recitation	Exemplary Citations in Reference
		• "Bedoop decoding function within the phone, in other embodiments the image data can be transmitted from the phone and decoded at a remote location." Rhoads at 46:47 – 53 "The provision of image sensors in cell phones enables other functionality. One is the capture of still or video imagery. Such image data can be compressed (typically by lossy processes such as MPEG, JPEG, or the like, implemented with dedicated hardware CODECs) and transmitted with the audio data. The screens on such phones can likewise be used for display of incoming image or video data." Rhoads at 49:48-65 • "The detailed embodiment contemplates that the object 20 is a magazine advertisement encoded with a steganographic watermark conveying a plural bit object identifier. The watermark is hidden in the advertisement's image in a manner indiscernable to human observers, but detectable by computer analysis. That analysis is performed by a watermark detector 30. Watermark detector 30 can be implemented in various different locations in the system of FIG. 1. Typically, the detector is implemented in the originating device 12, e.g., in the driver software 26, or in application software 28 c that serves to link to external resources based on detected watermarks. But it may be implemented elsewhere, e.g., in hardware in the interface electronics 24, in an operating system associated with the device, or outside device 12 altogether. Some systems may have plural watermark detectors, implemented at different locations throughout the system." Rhoads at Fig. 12 (and accompanying description)

Element	'529 Claim Recitation	Exemplary Citations in Reference
		FIG. 12 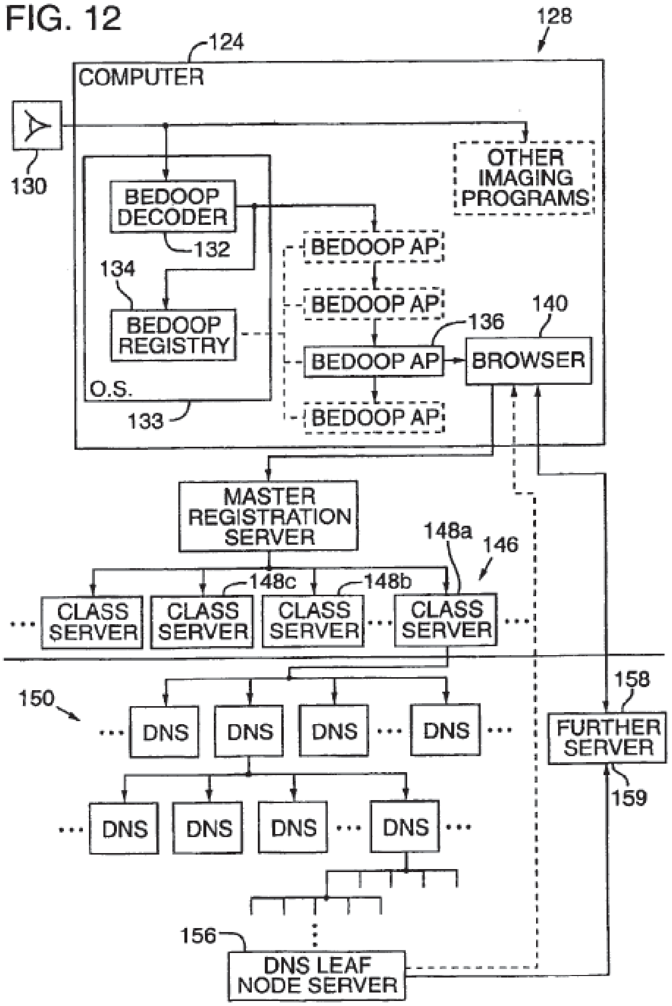 Rhoads at 12:23 – 33 • "Sometimes the computer system may encounter a Bedoop object for which it does not have a registered application program. In such case, a default

Element	'529 Claim Recitation	Exemplary Citations in Reference
		Bedoop application can be invoked. This default application can, e.g., establish an internet link to a remote server computer (or a network of such computers), and can transmit the Bedoop data (or a part of the Bedoop data) to that remote computer. The remote server can undertake the response itself, it can instruct the originating computer how to respond appropriately, or it can undertake some combination of these two responses. (Such arrangements are further considered below.)" Rhoads at 16:66 – 17:10 • "Inside the building the candidate may encounter a security guard. Seeing an unfamiliar person, the guard may visually compare the photo on the badge with the candidate's face. Additionally, the guard can present the badge to a portable Bedoop device, or to one of many Bedoop systems scattered through the building (e g., at every telephone). The Bedoop system extracts the Bedoop data from the card (i.e., from the DMV photograph), interrogates the DMV's DNS server with this Bedoop data, and receives in reply the name of the person depicted in the photograph. (If the Bedoop system is a telephone, the name may be displayed on a small LCD display commonly provided on telephones.)" To the extent it is found that Rhoads does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Rhoads does not anticipate this claim, Rhoads renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Rhoads with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

Element	'529 Claim Recitation	Exemplary Citations in Reference
[1C]	the service programmed to receive the data;	Rhoads discloses the service programmed to receive the data. Rhoads at Fig. 2 (and accompanying description) 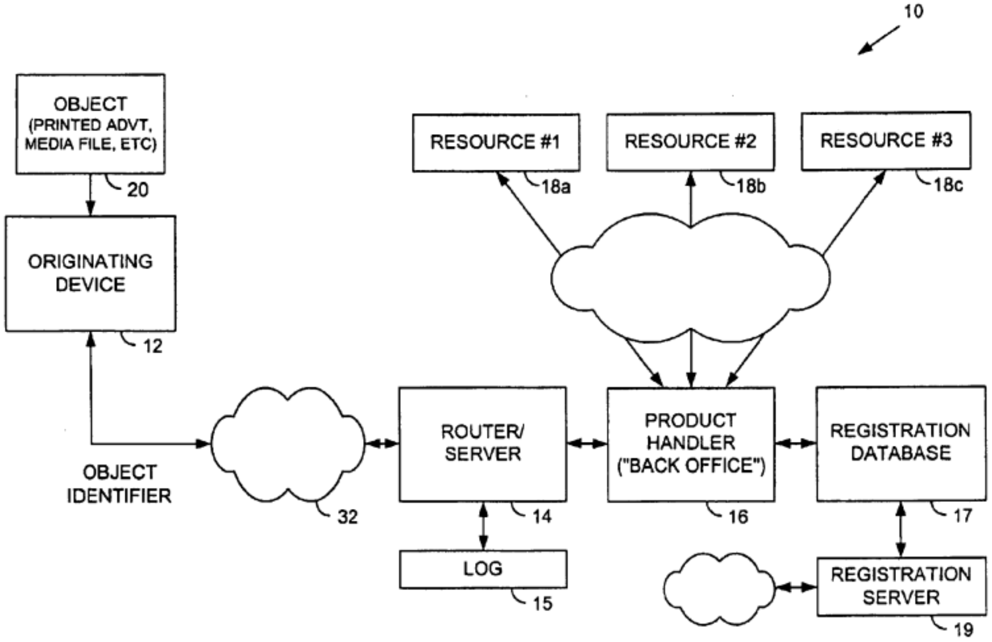 FIG. 2 Rhoads at 12:23 – 33 • "Sometimes the computer system may encounter a Bedoop object for which it does not have a registered application program. In such case, a default Bedoop application can be invoked. This default application can, e.g., establish an internet link to a remote server computer (or a network of such computers), and can transmit the Bedoop data (or a part of the Bedoop data) to that remote computer. The remote server can undertake the response itself, it can instruct the originating computer how to respond appropriately,

Element	'529 Claim Recitation	Exemplary Citations in Reference
		or it can undertake some combination of these two responses. (Such arrangements are further considered below.)" Rhoads at 16:66 – 17:10 • "Inside the building the candidate may encounter a security guard. Seeing an unfamiliar person, the guard may visually compare the photo on the badge with the candidate's face. Additionally, the guard can present the badge to a portable Bedoop device, or to one of many Bedoop systems scattered through the building (e g., at every telephone). The Bedoop system extracts the Bedoop data from the card (i.e., from the DMV photograph), interrogates the DMV's DNS server with this Bedoop data, and receives in reply the name of the person depicted in the photograph. (If the Bedoop system is a telephone, the name may be displayed on a small LCD display commonly provided on telephones.)" To the extent it is found that Rhoads does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Rhoads does not anticipate this claim, Rhoads renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Rhoads with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results." Rhoads at 49:19-47 • "Device 12 interacts with an object 20. The object can be electronic or not. Electronic objects can include computer files, representations of audio,

Element	'529 Claim Recitation	Exemplary Citations in Reference
		video, or still imagery (e.g., files or in streaming form), etc. Non-electronic objects can include physical objects such as newspapers, magazine pages, posters, product packaging, event tickets, credit cards, paper currency, etc. Non-electronic objects can also include sounds produced by loudspeakers. When used with nonelectronic objects, device 12 (FIG. 2) typically includes some form of sensor or transducer 22 to produce electronic signals or data corresponding to the object. Examples include CCD- or CMOS-based optical sensors (either as part of still- or video cameras, flatbed scanners, mice, or otherwise), microphones, barcode scanners, RF ID sensors, mag stripe readers, etc. In such cases, the sensor 22 may be coupled to associated interface electronics 24, which in turn may be coupled to device driver software 26, which in turn may be coupled to one or more application programs 28. Device driver software 26 serves as a software interface, communicating at a relatively high level with the application programs 28 (e.g., through API instructions whose content and format are standardized to facilitate application programming), and at a relatively low level with the interface electronics 24." Rhoads at 49:48-65 • "The detailed embodiment contemplates that the object 20 is a magazine advertisement encoded with a steganographic watermark conveying a plural bit object identifier. The watermark is hidden in the advertisement's image in a manner indiscernable to human observers, but detectable by computer analysis. That analysis is performed by a watermark detector 30. Watermark detector 30 can be implemented in various different locations in the system of FIG. 1. Typically, the detector is implemented in the originating device 12, e.g., in the driver software 26, or in application software 28 c that serves to link to external resources based on detected watermarks. But it may be implemented elsewhere, e.g., in hardware in the

Element	'529 Claim Recitation	Exemplary Citations in Reference
		interface electronics 24, in an operating system associated with the device, or outside device 12 altogether. Some systems may have plural watermark detectors, implemented at different locations throughout the system." Rhoads at 50:21–51:3 • "As more particularly detailed below, the handler 16 provides a response in accordance with the particular watermark payload. The response may be provided directly by the product handler to the device 12, or the handler may respond by communicating with a remote resource 18 (which may be, e.g., a data repository or service provider). In the former case, the handler 16 may identify a URL corresponding to the watermark (using the database 17), and return the URL to the application 28 c. Application 28 c can then pass the URL to a web browser 28 b in the device 12, and initiate a link to the internet site identified by the URL. Or the handler may have some locally stored data (e.g., audio or video, or software updates) and send it to the device 12 in response to the watermark. In the latter case, the handler 16 does not respond directly to the device 12. Instead, the handler responds by communicating with a remote resource 18. The communication can be as simple as logging receipt of the watermark message in a remote repository. Or it can be to authenticate device 12 (or a user thereof) to a remote resource in anticipation of a further transaction (e.g., the communication can form part of an on-line licensing or digital rights management transaction). Or the communication can request the remote resource to provide data or a service back to device 12 or to another destination (e.g., to initiate an FTP file transfer, or to request that a song selection identified by the watermark be downloaded to a user's personal music library, or to update software installed on device 12). In still other cases, hybrids of the two foregoing cases can be employed, e.g., handler 16 can send some data back to device 12, while also

Element	**'529 Claim Recitation**	**Exemplary Citations in Reference**
		communicating with a remote resource 18. In some cases, the response returned to the device 12 by handler 6 (or a remote resource 18) can serve to trigger some further action by the device 12. For example, the response returned to device 12 can include a WindowsMedia audio file, together with a request that the device 12 launch the WindowsMedia player installed on the device. (The launching of a browser pointed to a URL is another example of such triggering.) The illustrated product handler 16 comprises essentially the same hardware elements as the router 14, e.g., CPU, memory, etc. Although FIG. 2 shows just one product handler, several product handlers can be included in the system—either co-located of geographically distributed. Different handlers can be dedicated to different functions (e.g., serving URLs, serving music, etc.) or to different watermark sources (e.g., one responds to watermarks found in audio, another responds to watermarks found in print advertising, etc.)." Rhoads at Fig. 12 (and accompanying description)

Element	'529 Claim Recitation	Exemplary Citations in Reference
		FIG. 12 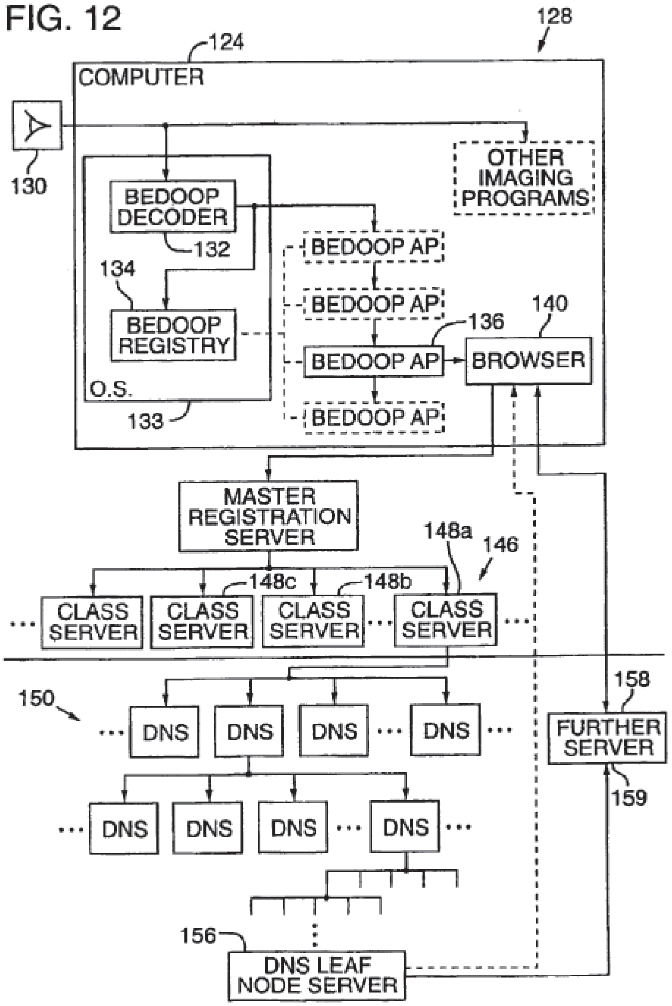 \ To the extent it is found that Rhoads does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Rhoads does not anticipate

Element	'529 Claim Recitation	Exemplary Citations in Reference
		this claim, Rhoads renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Rhoads with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1D]	identify an object within the image:	Rhoads discloses identify an object within the image. Rhoads at 4:34 – 51 • "According to another aspect, the invention includes a method comprising: (a) sensing an object identifier from a first object; (b) sending said first object identifier from a first device to a second device; (c) in response, at said second device, identifying address information corresponding to said first object identifier and sending same to the first device; (d) initiating a link from the first device in accordance with said address information; (e) at said second device, identifying additional objects related to said first object; identifying additional address information corresponding to said additional objects; and sending said additional address information to the first device; and (f) storing said additional address information in a memory at the first device; wherein, if an object included among said identified additional objects is sensed by the first device, the corresponding address information can be retrieved from said memory in the first device without the intervening delays of communicating with the second device." Rhoads at 9:35–41 • "The Steganographic decoding process may entail three Steps. In the first, the object is located. In the Second, the objects orientation is discerned. In

Element	'529 Claim Recitation	Exemplary Citations in Reference
		the third, the Bedoop data is extracted from the image data corresponding to the Bedoop object." Rhoads at Fig. 2 (and accompanying description) 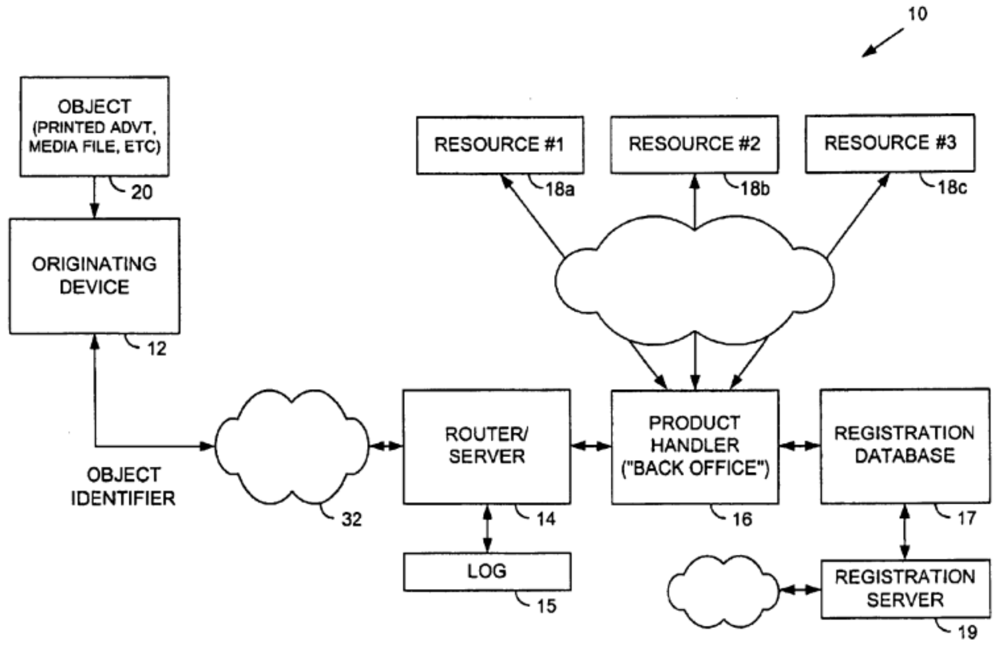 FIG. 2 Rhoads at 7:13-29 • "Referring to FIG. 11, a basic embodiment 110 of the present technology includes an optical sensor 112, a computer 114, and a network connection 116 to the internet 118. The illustrated optical sensor 112 is a digital camera having a resolution of 320 by 200 pixels (color or black and white) that stares out, grabbing frames of image data five times per second and storing same in one or more frame buffers. These frames of image data are analyzed

Element	'529 Claim Recitation	Exemplary Citations in Reference
		by a computer 114 for the presence of Bedoop data. (Essentially, Bedoop data is any form of plural-bit data encoding recognized by the system 110—data which, in many embodiments, initiates some action.) Once detected, the system responds in accordance with the detected Bedoop data (e.g., by initiating some local action, or by communication with a remote computer, such as over the internet, via an online service such as AOL, or using point-to-point dial-up communications, as with a bulletin board system)." Rhoads 7:30 – 46 • "Consider the milk carton example. The artwork on a milk carton can be adapted to convey Bedoop data. In the preferred embodiment, the Bedoop data is steganographically encoded (e.g., digitally watermarked) on the carton Numerous digital watermarking techniques are known—all of which convey data in a hidden form (i.e., on human inspection, it is not apparent that digitally encoded data is present). Exemplary techniques operate by slightly changing the luminance, or contours of selected points on artwork or text printed on the carton, or by splattering tiny droplets of ink on the carton in a seemingly random pattern. Each of these techniques has the effect of changing the local luminance at areas across the carton—luminance changes that can be detected by the computer 114 and decoded to extract the encoded digital data. In the case of a milk carton, the data may serve to identify the object as, e.g., a half gallon carton of Alpenrose brand skim milk." Rhoads at 9:35–10:10 • "The Steganographic decoding process may entail three Steps. In the first, the object is located. In the Second, the objects orientation is discerned. In the third, the Bedoop data is extracted from the image data corresponding to the Bedoop object.

Element	'529 Claim Recitation	Exemplary Citations in Reference
		The first step, object location, can be assisted by various clues. One is the placement of the object; typically the center of the image field will be a point on the object. The surrounding data can then be analyzed to try and discern the object's boundaries. Another location technique is slight movement. Although the user will typically try to hold the object still, there will usually be some jitter of the Bedoop object within the image frame (e.g., a few pixels back and forth). Background visual clutter, in contrast, will typically be stationary. Such movement may thus be sensed and used to identify the Bedoop object from within the image data. Still another object-location clue is object shape. Many Bedoop objects are rectangular in shape (or trapezoidal as viewed by the camera). Straight edge boundaries can thus be used to define an area of likely Bedoop data. Color is a further object identification clue that may be useful in some contexts. Yet another object location clue is spatial frequency. In imaging systems with well defined focal zones, undesired visual clutter may be at focal distances that results in blurring. The Bedoop object, in contrast, will be in focus and may be characterized by fine detail. Analyzing the image data for the high frequencies associated with fine detail can be used to distinguish the intended object from others. Characteristic markings on the object (as discussed below in connection with determining object orientation) can also be sensed and used in locating the object. Once the Bedoop object has been located within the image data, masking can be applied (if desired) to eliminate image data not corresponding to the intended object."

Element	'529 Claim Recitation	Exemplary Citations in Reference
		Rhoads at 22:9–27 • "To aid in identification, the Bedoop objects (e.g., badges) can be colored a distinctive color, permitting the system to more easily identify candidate objects from other items within the optical capture devices' field of view. Or the object can be provided with a retro-reflective coating, and the elevator can be equipped with one or more illumination sources of known spectral or temporal quality (e.g., constant infrared, or constant illumination with a single- or multi-line spectrum, or a pulsed light source of known periodicity; LEDs or semiconductor lasers, each with an associated diffuser, can be used for each the foregoing and can be paired with the image capture devices). Other such tell-tale clues can likewise be used to aid in object location. In all such cases, the optical capture device can sense the tell-tale clue(s) using a wide field of view sensor. The device can then be physically or electronically steered, and/or zoomed, to acquire a higher resolution image of the digitally-encoded object suitable for decoding." Rhoads at 49:3-14: • "Detection is the complementary operation to embedding, i.e., discerning a digital identifier from an object. Response refers to the action taken based on the discerned identifier. The middle two steps—embedding and detection—can employ any of myriad well-known technologies, including 1D and 2D barcodes, magnetic ink character recognition (MICR), optical character recognition (OCR), optical mark recognition (OMR), radio frequency identification (RF/ID), data glyphs, organic transistors, magnetic stripe, metadata, file header information, UV/IR identifiers, and other machine-readable indicia and techniques for associating plural-bit digital data with an electronic or physical object. The detailed embodiment employs watermarking

Element	'529 Claim Recitation	Exemplary Citations in Reference
		technology, although this is illustrative only." Rhoads at 50:21-34 - "As more particularly detailed below, the handler 16 provides a response in accordance with the particular watermark payload. The response may be provided directly by the product handler to the device 12, or the handler may respond by communicating with a remote resource 18 (which may be, e.g., a data repository or service provider). In the former case, the handler 16 may identify a URL corresponding to the watermark (using the database 17), and return the URL to the application 28 c. Application 28 c can then pass the URL to a web browser 28 b in the device 12, and initiate a link to the internet site identified by the URL. Or the handler may have some locally stored data (e.g., audio or video, or software updates) and send it to the device 12 in response to the watermark." To the extent it is found that Rhoads does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Rhoads does not anticipate this claim, Rhoads renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Rhoads with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1E]	distinguish an object present in the image from others using a database that stores data characteristics of	Rhoads discloses distinguish an object present in the image from others using a database that stores data characteristics of target objects.

Element	'529 Claim Recitation	Exemplary Citations in Reference
	target objects;	Rhoads at Fig. 2 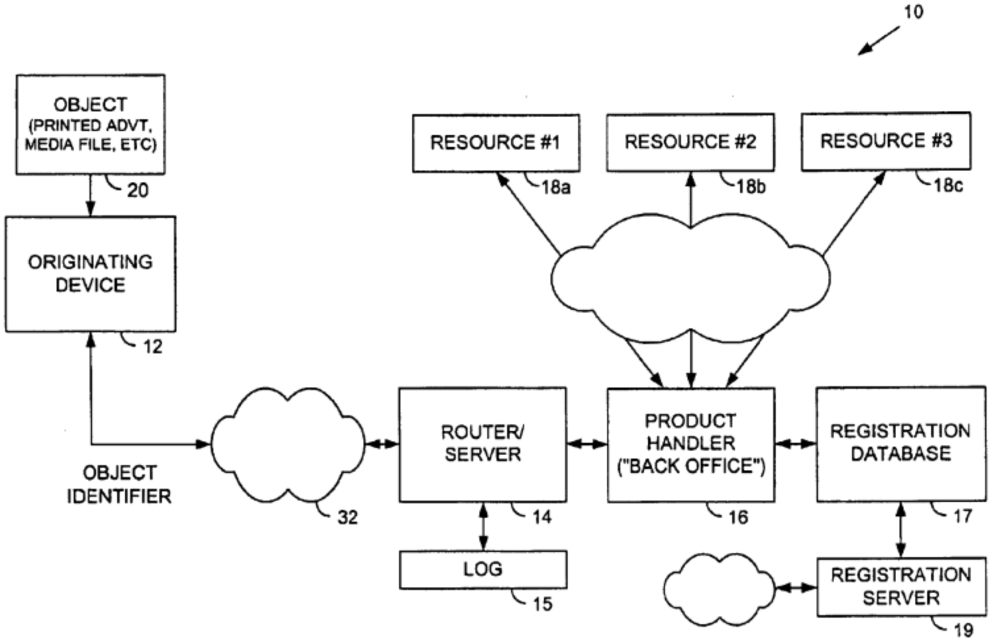 FIG. 2 Rhoads at 9:35–10:21 • "The Steganographic decoding process may entail three Steps. In the first, the object is located. In the Second, the objects orientation is discerned. In the third, the Bedoop data is extracted from the image data corresponding to the Bedoop object. The first step, object location, can be assisted by various clues. One is the placement of the object; typically the center of the image field will be a point on the object. The surrounding data can then be analyzed to try and discern

Element	'529 Claim Recitation	Exemplary Citations in Reference
		the object's boundaries. Another location technique is slight movement. Although the user will typically try to hold the object still, there will usually be some jitter of the Bedoop object within the image frame (e.g., a few pixels back and forth). Background visual clutter, in contrast, will typically be stationary. Such movement may thus be sensed and used to identify the Bedoop object from within the image data. Still another object-location clue is object shape. Many Bedoop objects are rectangular in shape (or trapezoidal as viewed by the camera). Straight edge boundaries can thus be used to define an area of likely Bedoop data. Color is a further object identification clue that may be useful in some contexts. Yet another object location clue is spatial frequency. In imaging systems with well defined focal zones, undesired visual clutter may be at focal distances that results in blurring. The Bedoop object, in contrast, will be in focus and may be characterized by fine detail. Analyzing the image data for the high frequencies associated with fine detail can be used to distinguish the intended object from others. Characteristic markings on the object (as discussed below in connection with determining object orientation) can also be sensed and used in locating the object. Once the Bedoop object has been located within the image data, masking can be applied (if desired) to eliminate image data not corresponding to the intended object. The next step in the decoding process, determining orientation of the Bedoop data, can likewise be discerned by reference to visual clues. For example,

Element	'529 Claim Recitation	Exemplary Citations in Reference
		some objects include subliminal graticule data, or other calibration data, steganographically encoded with the Bedoop data to aid in determining orientation. Others can employ overt markings, either placed for that sole purpose (e.g. reference lines or fiducials), or serving another purpose as well (e.g. lines of text), to discern orientation. Edge-detection algorithms can also be employed to deduce the orientation of the object by reference to its edges." Rhoads at 16:66 – 17:10 • "Inside the building the candidate may encounter a security guard. Seeing an unfamiliar person, the guard may visually compare the photo on the badge with the candidate's face. Additionally, the guard can present the badge to a portable Bedoop device, or to one of many Bedoop systems scattered through the building (e g., at every telephone). The Bedoop system extracts the Bedoop data from the card (i.e., from the DMV photograph), interrogates the DMV's DNS server with this Bedoop data, and receives in reply the name of the person depicted in the photograph. (If the Bedoop system is a telephone, the name may be displayed on a small LCD display commonly provided on telephones.)" Rhoads at 22:9–27 • "To aid in identification, the Bedoop objects (e.g., badges) can be colored a distinctive color, permitting the system to more easily identify candidate objects from other items within the optical capture devices' field of view. Or the object can be provided with a retro-reflective coating, and the elevator can be equipped with one or more illumination sources of known spectral or temporal quality (e.g., constant infra red, or constant illumination with a single- or multi-line spectrum, or a pulsed light source of known periodicity; LEDs or semiconductor lasers, each with an associated diffuser, can be used for each the foregoing and can be paired with the image capture devices).

Element	'529 Claim Recitation	Exemplary Citations in Reference
		Other such tell-tale clues can likewise be used to aid in object location. In all such cases, the optical capture device can sense the tell-tale clue(s) using a wide field of view sensor. The device can then be physically or electronically steered, and/or zoomed, to acquire a higher resolution image of the digitally-encoded object suitable for decoding." Rhoads at 23:33–54 • "When a recipient of a business card holds it in front of a Bedoop sensor, the operating system on the local system launches a local Bedoop application. That local Bedoop application, in turn, establishes an external internet connection to a remote business card server. The address of that server may already be known to the local Bedoop application (e.g., having been stored from previous use), or the local Bedoop system can traverse the above-described public network of DNS servers to reach the business card server. A database on the business card name server maintains a large collection of business card data, one database record per UID. When that server receives Bedoop data from a local Bedoop system, it parses out the UID and accesses the corresponding database record. This record typically includes more information than is commonly printed on conventional business cards. Sample fields from the record may include, for example, name, title, office phone, office fax, home phone, home fax, cellular phone, email address, company name, corporate web page address, personal web page address, secretary's name, spouse's name, and birthday. This record is transmitted back to the originating Bedoop system." Rhoads at 48:32–53 • "As noted, the following disclosure focuses on one particular application—a system for linking print media to electronic content It should be reiterated, however, that the technology is not so limited, and may more generally be

Element	'529 Claim Recitation	Exemplary Citations in Reference
		viewed as a system for linking any object (physical or electronic) to a corresponding networked or local resource. In accordance with an exemplary application of the below-detailed technology, digital watermarking is employed to convey a plural bit identifier within print media, such as magazine advertisements or articles, direct mail coupons or catalogs, bank- or credit-cards, and business cards. This identifier is read by software in a user's computing device and forwarded to a remote database. The remote database identifies a URL corresponding to the identifier, and provides the URL back to the user's computer—permitting a browser on the user's computer to display the URL-identified web page. That web page can provide additional information or services—more timely and/or more extensive than that provided by the print material. By such arrangement, more efficient internet navigation and access is provided to consumers, and more effective means for linking readers to e-commerce points of sale is provided to advertisers." Rhoads at 51:18-37 • "When shown a watermarked image, the application 28 c analyzes the image and extracts the embedded watermark payload (more particularly detailed below) from the image. The application sends some or all of this information in a message format to the router 14. The router 14 decodes the received message, looking for vendor and product information. Based on this information, it passes the message to a corresponding product handler 16. The product handler receives the message and attempts to match the detected watermark serial number to a registered watermark serial number earlier stored in the database 17. If a match is found, the product handler performs the desired action. As noted, typical actions include returning a URL for web redirection, serving up an HTML page for initial user navigation, initiating

Element	'529 Claim Recitation	Exemplary Citations in Reference
		software downloads, etc. If a match is not found, the product handler returns an error code and message to the application 28 c. If a match is found, but the corresponding action is unavailable, incomplete, inactive or invalid, the product handler returns an error code and message to the calling application." Rhoads at 84:11 – 17 - "B55. A method comprising: presenting an object to an optical sensor, the optical sensor producing image data, the object being steganographically encoded with plural bits of digital data; decoding said digital data from the image data; and in response to said digital data, displaying an image of a like object." To the extent it is found that Rhoads does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Rhoads does not anticipate this claim, Rhoads renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Rhoads with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1F]	associate the object with information; and	Rhoads discloses associate the object with information. Rhoads at 22:28–33

Element	'529 Claim Recitation	Exemplary Citations in Reference
		• "Magazine (and newspaper) pages can be steganographically encoded with Bedoop data to provide another "paper as portal" experience. As with the earlier-described office document case, the encoded data yields an address to a computer location (e.g., a web page) having the same, or related, content." Rhoads at 48:32–53 • "As noted, the following disclosure focuses on one particular application—a system for linking print media to electronic content It should be reiterated, however, that the technology is not so limited, and may more generally be viewed as a system for linking any object (physical or electronic) to a corresponding networked or local resource. In accordance with an exemplary application of the below-detailed technology, digital watermarking is employed to convey a plural bit identifier within print media, such as magazine advertisements or articles, direct mail coupons or catalogs, bank- or credit-cards, and business cards. This identifier is read by software in a user's computing device and forwarded to a remote database. The remote database identifies a URL corresponding to the identifier, and provides the URL back to the user's computer—permitting a browser on the user's computer to display the URL-identified web page. That web page can provide additional information or services—more timely and/or more extensive than that provided by the print material. By such arrangement, more efficient internet navigation and access is provided to consumers, and more effective means for linking readers to e-commerce points of sale is provided to advertisers." Rhoads at 51:18-37 • "When shown a watermarked image, the application 28 c analyzes the image and extracts the embedded watermark payload (more particularly detailed

Element	'529 Claim Recitation	Exemplary Citations in Reference
		below) from the image. The application sends some or all of this information in a message format to the router 14. The router 14 decodes the received message, looking for vendor and product information. Based on this information, it passes the message to a corresponding product handler 16. The product handler receives the message and attempts to match the detected watermark serial number to a registered watermark serial number earlier stored in the database 17. If a match is found, the product handler performs the desired action. As noted, typical actions include returning a URL for web redirection, serving up an HTML page for initial user navigation, initiating software downloads, etc. If a match is not found, the product handler returns an error code and message to the application 28 c. If a match is found, but the corresponding action is unavailable, incomplete, inactive or invalid, the product handler returns an error code and message to the calling application." Rhoads at 52:3-29 • "[user] Sees the data/action associated with the object (e.g., views web page)" To the extent it is found that Rhoads does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Rhoads does not anticipate this claim, Rhoads renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Rhoads with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains

Element	'529 Claim Recitation	Exemplary Citations in Reference
		familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1G]	return the information to the network-enabled device; and	Rhoads discloses return the information to the network-enabled device. Rhoads at 22:28–33 • "Magazine (and newspaper) pages can be steganographically encoded with Bedoop data to provide another "paper as portal" experience. As with the earlier-described office document case, the encoded data yields an address to a computer location (e.g., a web page) having the same, or related, content." Rhoads at 48:32–53 • "As noted, the following disclosure focuses on one particular application—a system for linking print media to electronic content It should be reiterated, however, that the technology is not so limited, and may more generally be viewed as a system for linking any object (physical or electronic) to a corresponding networked or local resource. In accordance with an exemplary application of the below-detailed technology, digital watermarking is employed to convey a plural bit identifier within print media, such as magazine advertisements or articles, direct mail coupons or catalogs, bank- or credit-cards, and business cards. This identifier is read by software in a user's computing device and forwarded to a remote database. The remote database identifies a URL corresponding to the identifier, and provides the URL back to the user's computer—permitting a browser on the user's computer to display the URL-identified web page. That web page can provide additional information or services—more timely and/or more extensive than that provided by the print material. By such arrangement, more efficient internet navigation and access is provided to consumers, and more effective means for linking readers to e-commerce points of sale is provided to advertisers."

Element	'529 Claim Recitation	Exemplary Citations in Reference
		Rhoads at 50:27-60 • "In the former case, the handler 16 may identify a URL corresponding to the watermark (using the database 17), and return the URL to the application 28 c. Application 28 c can then pass the URL to a web browser 28 b in the device 12, and initiate a link to the internet site identified by the URL. Or the handler may have some locally stored data (e.g., audio or video, or software updates) and send it to the device 12 in response to the watermark. In the latter case, the handler 16 does not respond directly to the device 12. Instead, the handler responds by communicating with a remote resource 18. The communication can be as simple as logging receipt of the watermark message in a remote repository. Or it can be to authenticate device 12 (or a user thereof) to a remote resource in anticipation of a further transaction (e.g., the communication can form part of an on-line licensing or digital rights management transaction). Or the communication can request the remote resource to provide data or a service back to device 12 or to another destination (e.g., to initiate an FTP file transfer, or to request that a song selection identified by the watermark be downloaded to a user's personal music library, or to update software installed on device 12). In still other cases, hybrids of the two foregoing cases can be employed, e.g., handler 16 can send some data back to device 12, while also communicating with a remote resource 18. In some cases, the response returned to the device 12 by handler 6 (or a remote resource 18) can serve to trigger some further action by the device 12. For example, the response returned to device 12 can include a WindowsMedia audio file, together with a request that the device 12 launch the WindowsMedia player installed on the device. (The launching

Element	**'529 Claim Recitation**	**Exemplary Citations in Reference**
		of a browser pointed to a URL is another example of such triggering.)" Rhoads at 51:18-37 - "When shown a watermarked image, the application 28 c analyzes the image and extracts the embedded watermark payload (more particularly detailed below) from the image. The application sends some or all of this information in a message format to the router 14. The router 14 decodes the received message, looking for vendor and product information. Based on this information, it passes the message to a corresponding product handler 16. The product handler receives the message and attempts to match the detected watermark serial number to a registered watermark serial number earlier stored in the database 17. If a match is found, the product handler performs the desired action. As noted, typical actions include returning a URL for web redirection, serving up an HTML page for initial user navigation, initiating software downloads, etc. If a match is not found, the product handler returns an error code and message to the application 28 c. If a match is found, but the corresponding action is unavailable, incomplete, inactive or invalid, the product handler returns an error code and message to the calling application." Rhoads at 52:3-29 "[user] Sees the data/action associated with the object (e.g., views web page)" To the extent it is found that Rhoads does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Rhoads does not anticipate this claim, Rhoads renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with

Element	'529 Claim Recitation	Exemplary Citations in Reference
		other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Rhoads with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1H]	the network-enabled device further programmed to present the information related to the object to a user.	Rhoads discloses the network-enabled device further programmed to present the information related to the object to a user. Rhoads at 22:60 – 23:10 • "On presenting a magazine to the optical scanner device of a Bedoop-compliant computer, the computer senses the Bedoop data, decodes same, and launches a web browser to an internet address corresponding to the Bedoop data. If the magazine page is an advertisement, the internet address can provide information complementary to the advertisement. For example, if the magazine page is an advertisement for a grocery item, the Bedoop data can identify a web page on which recipes that make use of the advertised item are presented. If the magazine page includes a photo of a tropical beach, the Bedoop data can lead to a travel web page (e.g., hosted by Expedia or other travel enterprise) that presents fare and lodging information useful to a reader who wants to vacation at the illustrated beach. (The fare information can be customized to the reader's home airport by reference to user profile data stored on the user's computer and relayed to the web site to permit customization of the displayed page.)" Rhoads at 48:32–53 • "As noted, the following disclosure focuses on one particular application—a system for linking print media to electronic content It should be reiterated,

Element	'529 Claim Recitation	Exemplary Citations in Reference
		however, that the technology is not so limited, and may more generally be viewed as a system for linking any object (physical or electronic) to a corresponding networked or local resource. In accordance with an exemplary application of the below-detailed technology, digital watermarking is employed to convey a plural bit identifier within print media, such as magazine advertisements or articles, direct mail coupons or catalogs, bank- or credit-cards, and business cards. This identifier is read by software in a user's computing device and forwarded to a remote database. The remote database identifies a URL corresponding to the identifier, and provides the URL back to the user's computer—permitting a browser on the user's computer to display the URL-identified web page. That web page can provide additional information or services—more timely and/or more extensive than that provided by the print material. By such arrangement, more efficient internet navigation and access is provided to consumers, and more effective means for linking readers to e-commerce points of sale is provided to advertisers." Rhoads at 50:5-10 • "The illustrated application 28 c is a software program that serves to communicate the watermark data from the device 12 to the router/server 14 through one or more communications links 32 (e.g., the internet). Application 28 c also receives information from the communication links 32 and presents same to the user (or otherwise uses same)." Rhoads at 50:27-60 • "In the former case, the handler 16 may identify a URL corresponding to the watermark (using the database 17), and return the URL to the application 28 c. Application 28 c can then pass the URL to a web browser 28 b in the device 12, and initiate a link to the internet site identified by the URL. Or the handler may have some locally stored data (e.g., audio or

Element	'529 Claim Recitation	Exemplary Citations in Reference
		video, or software updates) and send it to the device 12 in response to the watermark. In the latter case, the handler 16 does not respond directly to the device 12. Instead, the handler responds by communicating with a remote resource 18. The communication can be as simple as logging receipt of the watermark message in a remote repository. Or it can be to authenticate device 12 (or a user thereof) to a remote resource in anticipation of a further transaction (e.g., the communication can form part of an on-line licensing or digital rights management transaction). Or the communication can request the remote resource to provide data or a service back to device 12 or to another destination (e.g., to initiate an FTP file transfer, or to request that a song selection identified by the watermark be downloaded to a user's personal music library, or to update software installed on device 12). In still other cases, hybrids of the two foregoing cases can be employed, e.g., handler 16 can send some data back to device 12, while also communicating with a remote resource 18. In some cases, the response returned to the device 12 by handler 6 (or a remote resource 18) can serve to trigger some further action by the device 12. For example, the response returned to device 12 can include a WindowsMedia audio file, together with a request that the device 12 launch the WindowsMedia player installed on the device. (The launching of a browser pointed to a URL is another example of such triggering.)" Rhoads at 52:3-29 • "[user] Sees the data/action associated with the object (e.g., views web page)"

Element	'529 Claim Recitation	Exemplary Citations in Reference
		To the extent it is found that Rhoads does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Rhoads does not anticipate this claim, Rhoads renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Rhoads with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

2. Claim 4

Claim	'529 Claim Recitation	Exemplary Citations in Reference
4	The system of claim 1, wherein the service is further programmed to identify the object by applying multiple algorithms to at least a portion of the image or of the data, and identify the object as a function of a confidence levels associated with results of the multiple algorithms.	Rhoads discloses the service is further programmed to identify the object by applying multiple algorithms to at least a portion of the image or of the data, and identify the object as a function of a confidence levels associated with results of the multiple algorithms. Rhoads at 50:21-48 • "In the former case, the handler 16 may identify a URL corresponding to the watermark (using the database 17), and return the URL to the application 28 c. Application 28 c can then pass the URL to a web browser 28 b in the device 12, and initiate a link to the internet site identified by the URL. Or the handler may have some locally stored data (e.g., audio or video, or software updates) and send it to the device 12 in response to the watermark.

		In the latter case, the handler 16 does not respond directly to the device 12. Instead, the handler responds by communicating with a remote resource 18. The communication can be as simple as logging receipt of the watermark message in a remote repository. Or it can be to authenticate device 12 (or a user thereof) to a remote resource in anticipation of a further transaction (e.g., the communication can form part of an on-line licensing or digital rights management transaction). Or the communication can request the remote resource to provide data or a service back to device 12 or to another destination (e.g., to initiate an FTP file transfer, or to request that a song selection identified by the watermark be downloaded to a user's personal music library, or to update software installed on device 12). In still other cases, hybrids of the two foregoing cases can be employed, e.g., handler 16 can send some data back to device 12, while also communicating with a remote resource 18." To the extent it is found that Rhoads does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Rhoads does not anticipate this claim, Rhoads renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Rhoads with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

3. Claim 20

Claim	'529 Claim Recitation	Exemplary Citations in Reference
20	The system of claim 1, wherein the data processing operation is selected from the group consisting of contrast enhancement, noise removal, and de-blurring.	Rhoads discloses the data processing operation is selected from the group consisting of contrast enhancement, noise removal, and de-blurring. Rhoads at 10:22-30 • "Some embodiments filter the image data at some point in the process to aid in ultimate Bedoop data extraction. One use of such filtering is to mitigate image data artifacts due to the particular optical sensor. For example, CCD arrays have regularly-spaced sensors that sample the optical image at uniformly spaced discrete points. This discrete sampling effects a transformation of the image data, leading to certain image artifacts. An appropriately configured filter can reduce some of these artifacts." To the extent it is found that Rhoads does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Rhoads does not anticipate this claim, Rhoads renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Rhoads with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

Exhibit G-11

Bank of America Corporation ("BofA" or "Defendant") provides this chart subject to all reservations, objections, statements, and disclaimers set forth herein and in BofA's Final Invalidity Contentions Cover Pleading, as well as any amendment, supplement, or modification thereof, which are incorporated herein by reference in their entirety.

On information and belief, the application that matured into Ehrhart was filed on July 13, 2001, claiming the priority date of July 13, 2001. Ehrhart published at least as of April 20, 2004, and is available as prior art at least under 35 U.S.C. § 102(e).

As shown in the chart below, Ehrhart anticipates asserted claims 4, 20 of the '529 Patent. To the extent Ehrhart is found not to expressly disclose certain limitations in the asserted claims, such limitations are inherent. To the extent Ehrhart is found not to anticipate any asserted claims or claim elements of the '529 Patent, the reference nevertheless renders those claims or claim elements obvious under 35 U.S.C. § 103, either alone or in combination with other art identified herein or in the cover pleading of BofA's Invalidity Contentions, for the reasons identified in the cover pleading. These Invalidity Contentions are not an admission by BofA that the accused products, including any current or past versions of the accused products, are covered by, or infringe the asserted claims, but are based instead on the recognition that if the claims are interpreted to be broad enough to encompass the accused products, the claims must also be construed to have that same scope when considering whether they are invalid.

The following chart is based on the Court's claim construction rulings (Dkt. Nos. 145, 153, 236), and to the extent not contradictory, Plaintiffs' apparent interpretations of the claims in Plaintiffs' "Preliminary" Final Infringement Contentions. BofA notes that in many instances, Plaintiffs' "Preliminary" Final Infringement Contentions ignore and/or contradict the Court's claim construction rulings. BofA does not accept the assumptions concerning the scope and meaning implicit in Plaintiffs' Preliminary Final Infringement Contentions, to the extent those assumptions are discernible, and reserves the right to challenge Plaintiffs' proposed (or implied) constructions. Furthermore, Plaintiffs' "Preliminary" Final Contentions fail to comply with the level of specificity required by SPR 2.1.2 and 2.1.3, and include improper reservation of rights to supplement with "additional ways in which BoA Accused Products infringe." Any alleged deficient discovery is of Plaintiffs own making. Most if not all of the deficiencies in Plaintiffs' Preliminary Final Contentions have been documented by Defendants through multiple letters since at least May 18, 2021 and up to Bank of America's most recent letter of August 19, 2023. Plaintiffs' shifting infringement contentions, deficiencies in its infringement contentions, and changing claim scope have prejudiced BofA from being able to reasonably prepare its defenses, including its invalidity contentions. BofA also reserves the right to revise and supplement these charts if and when Plaintiffs are permitted to provide revised or supplemental Infringement Contentions. Citations given in the chart below are merely representative of the respective elements to provide notice to Plaintiffs of BofA's invalidity theories and are not meant to be exhaustive.

1. Claim 1

Element	'529 Claim Recitation	Exemplary Citations in Reference
[1Pre]	A system comprising:	Ehrhart discloses a system. "The present invention relates to an optical reader that includes a color imaging assembly that generates color imaging data. An image analysis circuit determines if the acquired image should be characterized as a color photograph or as including a graphical symbol. A processing circuit processes the imaging data based on the image analysis circuit's determination of whether the image is a graphical symbol or a color photograph. The present invention allows a user to acquire and process both color images and graphical symbols, such as bar codes, text, OCR symbols or signatures. The optical reader of the present invention is also configured to associate an acquired image with at least one other acquired image." Ehrhart at Abstract.
[1A]	a camera that captures an image;	Ehrhart discloses a camera that captures an image. *Id.* at [1Pre]. "The optical reader of the present invention automatically, or through manual Selection, determines whether a captured image is a color photographic image or, a color image that includes a graphical Symbol. Subsequently, the optical reader of the present invention processes the acquired imaging data in accordance with that determination. The optical reader of the present invention is operative to acquire and associate a plurality of acquired images. One aspect of the present invention is an optical reader.

Element	'529 Claim Recitation	Exemplary Citations in Reference
		The optical reader includes a color imaging assembly for acquiring an image of an object, the color imaging assembly generating imaging data corresponding to the image. An image analysis circuit is coupled to the color imaging assembly. The image analysis circuit being configured to determine if the color imaging data includes at least one graphical Symbol. The image is classified as a graphical Symbol, or the image is classified as a color photograph if the color imaging data does not include at least one graphical Symbol. A processing circuit is coupled to the image analysis circuit. The processing circuit is operative to process the imaging data based on the determination. In another aspect, the present invention includes an optical reader for capturing an image of an object. The optical reader includes a color imaging assembly for converting the image of the object into color digital data corresponding to the image." Ehrhart at 1:57-2:16. "As embodied herein, and depicted in FIGS. 1A-1D, perspective views of the optical reader in accordance with various embodiments of the present invention are disclosed. FIG. 1A shows the underside of hand held wireless optical reader 10. FIG. 1B shows the top of the optical reader depicted in FIG. 1A. Optical reader 10 includes housing 100, antenna 102, window 104 and trigger 12. Window 104 accommodates illumination assembly 20 and imaging assembly 30. As shown in FIG. 1B, the top side of reader 10 includes function keys 14, alphanumeric key pad 16, and display 60. In one embodiment, function keys 14 include an enter key and up and down cursor keys. FIG. 1C is also a hand held wireless optical reader 10. Reader 10 includes function keys 14, alphanumeric key pad 16, writing stylus 18, display 60, and signature block 62. Stylus 18 is employed by a user to write his signature in signature block 62. FIG. 1D shows yet another embodiment of optical reader 10 of the present invention. In this embodiment, reader 10 includes a gun-shaped housing 100. Display 60 and keypad 16 are disposed on a top portion of gun-shaped housing 100, whereas trigger 12 is disposed on the underside of the top portion of housing 100.

Element	**'529 Claim Recitation**	**Exemplary Citations in Reference**
		Housing 100 also includes window 104 that accommodates illumination assembly 20 and imaging assembly 30. Wire 106 is disposed at the butt-end of housing 100. Wire 106 provides optical reader 10 with a hard wired communication link for external devices such as a host processor or other data collection devices." Ehrhart at 5:51-6:10. To the extent it is found that Ehrhart does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Ehrhart does not anticipate this claim, Ehrhart renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Ehrhart with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1B]	a network-enabled device that conducts a data processing operation on at least a portion of the image to produce data, and sends the data to a service;	Ehrhart discloses a network-enabled device that conducts a data processing operation on at least a portion of the image to produce data, and sends the data to a service. *Id.* at [1Pre, 1A]. "In another aspect, the present invention includes an optical reader for capturing an image of an object. The optical reader includes a color imaging assembly for capturing the image as color imaging data. A classification circuit is coupled to the color imaging assembly, the classification circuit being configured to process at least a portion of the color imaging data to thereby Select one of a plurality of classifications, whereby the image is classified as a color photographic image, or as

Element	'529 Claim Recitation	Exemplary Citations in Reference
		an image that includes at least one graphical Symbol. An automatic mode Selector is coupled to the classification circuit, the automatic mode Selector being configured to Select an optical reader mode in accordance with the Selected classification. A processor is coupled to the classification circuit, the processor being programmed to process the color imaging data in accordance with the optical reader mode Selected by the automatic mode Selector." Ehrhart at 2:26-42. "In accordance with the invention, the present invention for an optical reader includes a color imaging assembly for acquiring color imaging data. An image analysis circuit determines if the acquired image includes at least one graphical symbol. A processing circuit processes the imaging data based on the determination of whether the image includes at least one graphical symbol. The present invention allows a user to read graphical symbols, such as bar codes, text, OCR characters or signatures using a color imager. The color optical reader of the present invention is configured to automatically determine whether a color image includes a graphical symbol, or is merely a color photographic image. The optical reader of the present invention also is operative to associate one acquired image with at least one subsequently acquired image." Ehrhart at 5:36-50. "If processor 40 does not locate a bar code symbol it moves onto step 420 and looks for OCR-A or OCR-B characters. If it finds these characters it performs optical character recognition in Step 422. If it does not, processor evaluates the image for the presence of text. If text is located, the image is cropped, and the text is compressed and Stored in Steps 428 and 430. If the image does not include text, processor 40 evaluates the image for the presence of a signature. If one is present, the image is cropped, and the data is compressed and stored in steps 428 and 430. In another embodiment, optical reader 10 is networked, and processor 40 communicates with remote network resources to provide Signature Verification Services. If processor 40 cannot detect a bar code Symbol, OCR symbols, text, or a signature, the user is asked in step 432 if he desires to store the image. If he does, the color photographic image is Stored in memory in Step 434." Ehrhart at 10:5-

Element	**'529 Claim Recitation**	**Exemplary Citations in Reference**
		20. To the extent it is found that Ehrhart does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Ehrhart does not anticipate this claim, Ehrhart renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Ehrhart with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1C]	the service programmed to receive the data;	Ehrhart discloses the service programmed to receive the data. *Id.* at [1Pre, 1A, 1B]. "LAN 1820 includes server 1822, computer 1824, database 1826, and a plurality of optical readers 10. Database 1826 is used to store associated images along with other data fields. For example, it would be rather useful to store additional information with the associated images shown in FIG. 13. One may want to associate the delivery means, route, driver, and other related information for subsequent analysis. Network 1810 allows reader 10, PAN 1850, and wireless system 1400 a way to store such data in database 1826. System analysts can access this information via personal computer 1802 connected to network 1810. In one embodiment, LAN 1820 includes an Internet website. In this embodiment, users are authenticated before gaining access to database 1826." Ehrhart at 16:1-12.

Element	'529 Claim Recitation	Exemplary Citations in Reference
		To the extent it is found that Ehrhart does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Ehrhart does not anticipate this claim, Ehrhart renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Ehrhart with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1D]	identify an object within the image:	Ehrhart discloses identify an object within the image. *Id.* at [1Pre, 1A, 1B, 1C]. "In accordance with the invention, the present invention for an optical reader includes a color imaging assembly for acquiring color imaging data. An image analysis circuit determines if the acquired image includes at least one graphical symbol. A processing circuit processes the imaging data based on the determination of whether the image includes at least one graphical symbol. The present invention allows a user to read graphical symbols, such as bar codes, text, OCR characters or signatures using a color imager. The color optical reader of the present invention is configured to automatically determine whether a color image includes a graphical symbol, or is merely a color photographic image. The optical reader of the present invention also is operative to associate one acquired image with at least one subsequently acquired image." Ehrhart at 5:36-50. "In the embodiment depicted in FIG. 2, ASIC 44 is implemented using a programmable logic array (PLA) device. In a similar embodiment, ASIC 44 is

Element	'529 Claim Recitation	Exemplary Citations in Reference
		implemented using a field programmable gate array (FPGA) device. ASIC 44 is tasked with controlling the image acquisition process, and the storage of image data. As part of the image acquisition process, ASIC 44 performs various timing and control functions including control of light source 24, control of color imager 34, and control of external interface 56. It will be apparent to those of ordinary skill in the pertinent art that modifications and variations can be made to processor 40 of the present invention depending on the cost, availability, and performance of off-the-shelf microprocessors, and the type of color imager used. In one embodiment, microprocessor 42 and ASIC 44 are replaced by a single microprocessor 40. In one embodiment, microprocessor 40 is implemented using a single RISC processor. In yet another embodiment, microprocessor 40 is implemented using a RISC and DSP hybrid processor." Ehrhart at 7:32-50. "In step 1004, processor 40 identifies the page features. Processor 40 analyzes the page and divides it into blank and non-blank portions. The non-blank portions are analyzed to distinguish text regions from non-text regions. After determining the layout of the page, processor 40 uses black-to-white transitions to determine degrees of skew. In step 1008, horizontal white spaces are identified to separate lines of text. In step 1010, vertical white spaces are identified within each line of text to thereby separate individual words and characters from each other. In Step 1014, a character recognition algorithm is used in an attempt to recognize each individual character." Ehrhart at 12:55-66.

		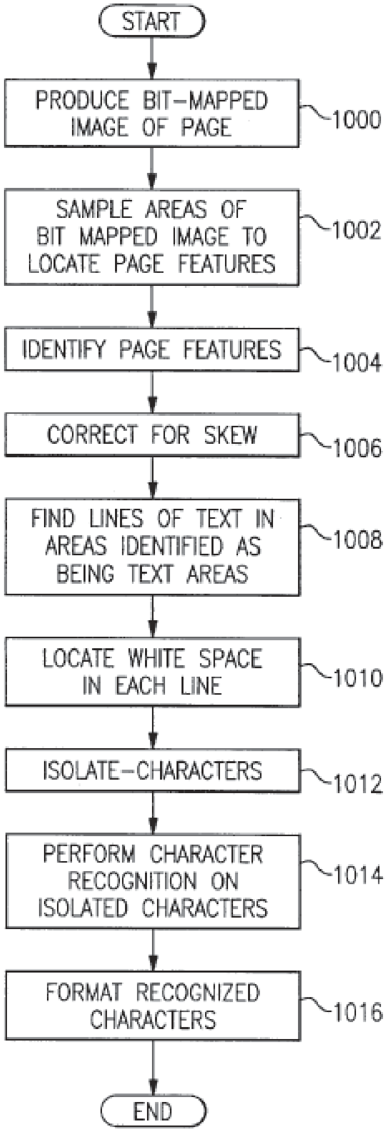START PRODUCE BIT-MAPPED IMAGE OF PAGE 1000 SAMPLE AREAS OF BIT MAPPED IMAGE TO LOCATE PAGE FEATURES 1002 IDENTIFY PAGE FEATURES 1004 CORRECT FOR SKEW 1006 FIND LINES OF TEXT IN AREAS IDENTIFIED AS BEING TEXT AREAS 1008 LOCATE WHITE SPACE IN EACH LINE 1010 ISOLATE-CHARACTERS 1012 PERFORM CHARACTER RECOGNITION ON ISOLATED CHARACTERS 1014 FORMAT RECOGNIZED CHARACTERS 1016 END FIG.10

Element	'529 Claim Recitation	Exemplary Citations in Reference
		Ehrhart at Fig. 10.

		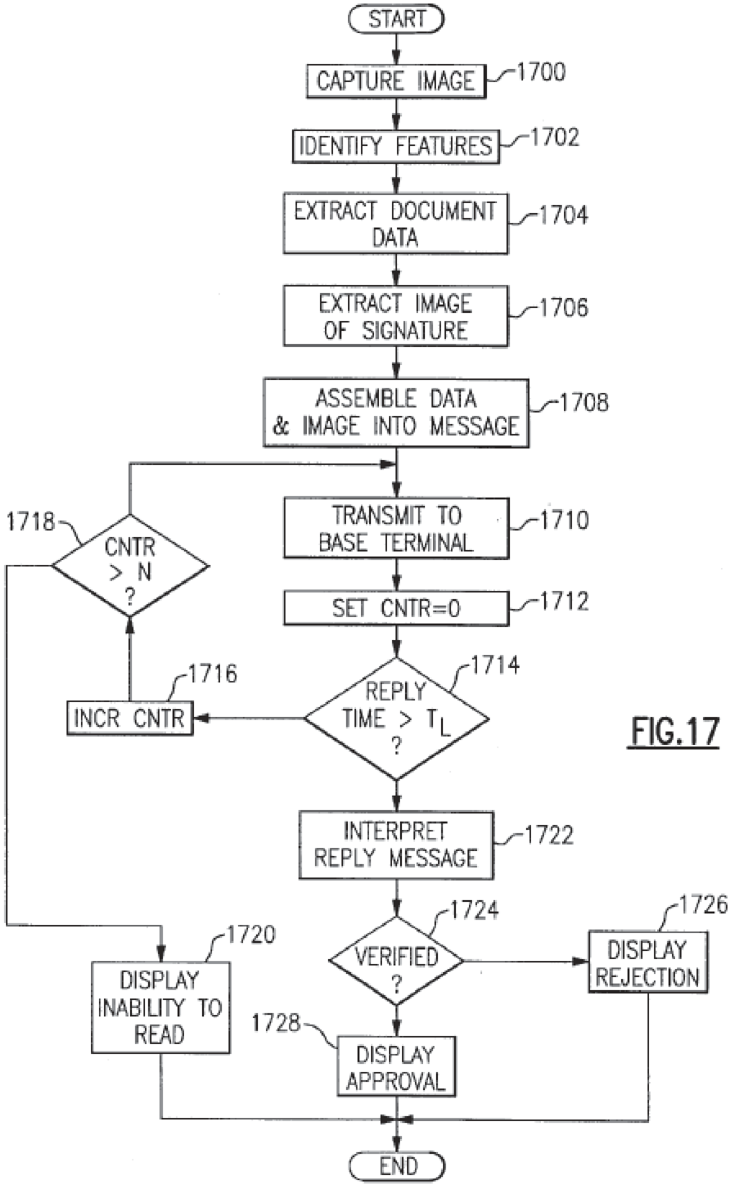START CAPTURE IMAGE 1700 IDENTIFY FEATURES 1702 EXTRACT DOCUMENT DATA 1704 EXTRACT IMAGE OF SIGNATURE 1706 ASSEMBLE DATA & IMAGE INTO MESSAGE 1708 TRANSMIT TO BASE TERMINAL 1710 SET CNTR=0 1712 REPLY TIME > T_L ? 1714 INCR CNTR 1716 CNTR > N ? 1718 INTERPRET REPLY MESSAGE 1722 VERIFIED ? 1724 DISPLAY REJECTION 1726 DISPLAY INABILITY TO READ 1720 DISPLAY APPROVAL 1728 END FIG.17

Element	'529 Claim Recitation	Exemplary Citations in Reference
		Ehrhart at Fig. 17. To the extent it is found that Ehrhart does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Ehrhart does not anticipate this claim, Ehrhart renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Ehrhart with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1E]	distinguish an object present in the image from others using a database that stores data characteristics of target objects;	Ehrhart discloses distinguish an object present in the image from others using a database that stores data characteristics of target objects. *Id.* at [1Pre, 1A, 1B, 1C, 1D]. "In the Automatic imaging mode, processor 40 is programmed to analyze the color imaging data to determine if an acquired image includes a graphical symbol or is merely a color photographic image. If it makes the determination that the color image includes a graphical symbol, it further analyzes the acquired image and classifies it as a bar code, OCR symbol, text, or a signature. Based on the classification, optical reader 10 jumps to the appropriate routine in EROM 48. The semi-automatic mode is similar. Thus, in the automatic or semi-automatic modes, the bar code scanning mode, the OCR/text mode, the signature capture mode, the color photography mode, and the association mode are controlled by the application program, not by the user." Ehrhart at 8:36-49.

Element	'529 Claim Recitation	Exemplary Citations in Reference
		"If processor 40 does not locate a bar code symbol it moves onto step 420 and looks for OCR-A or OCR-B characters. If it finds these characters it performs optical character recognition in step 422. If it does not, processor evaluates the image for the presence of text. If text is located, the image is cropped, and the text is compressed and stored in steps 428 and 430. If the image does not include text, processor 40 evaluates the image for the presence of a signature. If one is present, the image is cropped, and the data is compressed and stored in steps 428 and 430. In another embodiment, optical reader 10 is networked, and processor 40 communicates with remote network resources to provide signature verification services. If processor 40 cannot detect a bar code symbol, OCR symbols, text, or a signature, the user is asked in step 432 if he desires to store the image. If he does, the color photographic image is stored in memory in step 434." Ehrhart at 10:5-20. "In step 1004, processor 40 identifies the page features. Processor 40 analyzes the page and divides it into blank and non-blank portions. The non-blank portions are analyzed to distinguish text regions from non-text regions. After determining the layout of the page, processor 40 uses black-to-white transitions to determine degrees of skew. In step 1008, horizontal white spaces are identified to separate lines of text. In step 1010, vertical white spaces are identified within each line of text to thereby separate individual words and characters from each other. In Step 1014, a character recognition algorithm is used in an attempt to recognize each individual character." Ehrhart at 12:55-66. "As embodied herein and depicted in FIG. 11, a flow chart showing a method for performing OCR in accordance with yet another embodiment of the present invention is disclosed. In step 1100, reader 10 produces a bit-mapped image of the page. Subsequently, processor 40 finds lines of text in the image, locates the white spaces in each line, and isolates the characters. In step 1108, processor 40 performs character recognition, either OCR-A or OCR-B, as desired. The decoded characters are stored in memory." Ehrhart at 13:1-9.

		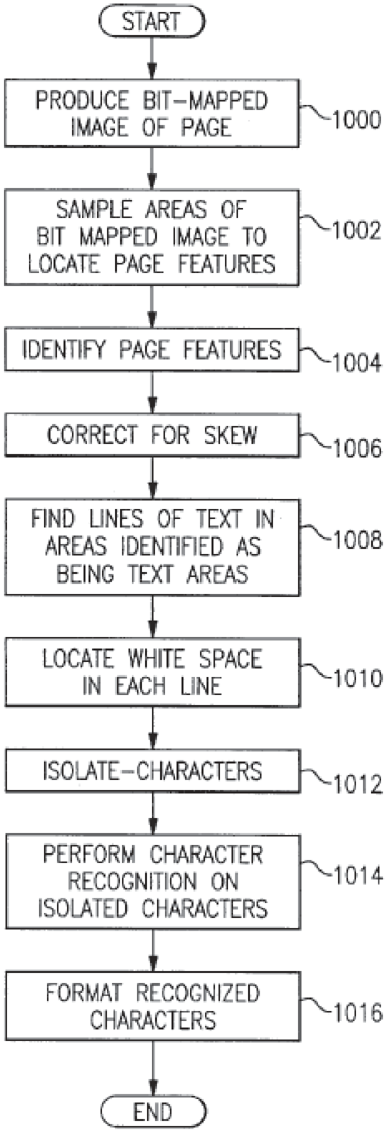START PRODUCE BIT-MAPPED IMAGE OF PAGE 1000 SAMPLE AREAS OF BIT MAPPED IMAGE TO LOCATE PAGE FEATURES 1002 IDENTIFY PAGE FEATURES 1004 CORRECT FOR SKEW 1006 FIND LINES OF TEXT IN AREAS IDENTIFIED AS BEING TEXT AREAS 1008 LOCATE WHITE SPACE IN EACH LINE 1010 ISOLATE-CHARACTERS 1012 PERFORM CHARACTER RECOGNITION ON ISOLATED CHARACTERS 1014 FORMAT RECOGNIZED CHARACTERS 1016 END FIG.10

Element	'529 Claim Recitation	Exemplary Citations in Reference
		Ehrhart at Fig. 10.

		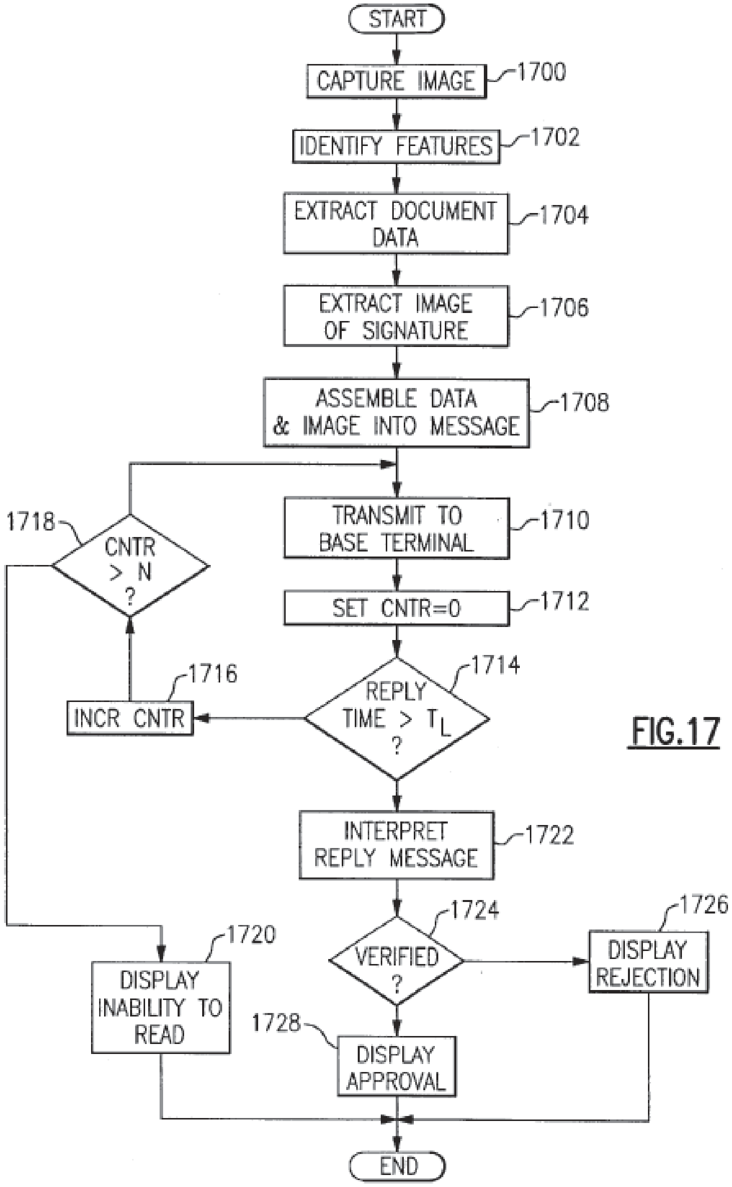START CAPTURE IMAGE 1700 IDENTIFY FEATURES 1702 EXTRACT DOCUMENT DATA 1704 EXTRACT IMAGE OF SIGNATURE 1706 ASSEMBLE DATA & IMAGE INTO MESSAGE 1708 1718 CNTR > N ? TRANSMIT TO BASE TERMINAL 1710 SET CNTR=0 1712 1716 INCR CNTR 1714 REPLY TIME > T_L ? INTERPRET REPLY MESSAGE 1722 1724 1726 1720 VERIFIED ? DISPLAY REJECTION DISPLAY INABILITY TO READ 1728 DISPLAY APPROVAL END FIG.17

Element	'529 Claim Recitation	Exemplary Citations in Reference
		Ehrhart at Fig. 17. To the extent it is found that Ehrhart does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Ehrhart does not anticipate this claim, Ehrhart renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Ehrhart with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1F]	associate the object with information; and	Ehrhart discloses associate the object with information. *Id.* at [1Pre, 1A, 1B, 1C, 1D, 1E]. "As embodied herein and depicted in FIG. 13, an example of image association in accordance with the present invention is disclosed. One or ordinary skill in the art will recognize that associated images 1300 can be disposed on paper, displayed electronically on display 60, or displayed electronically sing other electronic means, such as a computer monitor. In this example, the first image captured is color photograph 1302 which shows a damaged parcel. The second image captured is bar code 1304 affixed to the side of the damaged parcel. Processor 40 decodes bar code 1304 and associates decoded bar code data 1306 with color photograph 1302. In this example, the user elected to associate a third image, signature 1308. Thus, personnel viewing record 1300 may reasonably conclude that a damaged parcel was delivered

Element	'529 Claim Recitation	Exemplary Citations in Reference
		to Company XYZ, and that the person signing for the parcel delivery was someone named John W. Smith." Ehrhart at 13:27-43. "In the Automatic imaging mode, processor 40 is programmed to analyze the color imaging data to determine if an acquired image includes a graphical symbol or is merely a color photographic image. If it makes the determination that the color image includes a graphical symbol, it further analyzes the acquired image and classifies it as a bar code, OCR symbol, text, or a signature. Based on the classification, optical reader 10 jumps to the appropriate routine in EROM 48. The semi-automatic mode is similar. Thus, in the automatic or semi-automatic modes, the bar code scanning mode, the OCR/text mode, the signature capture mode, the color photography mode, and the association mode are controlled by the application program, not by the user." Ehrhart at 8:36-49. "In step 1004, processor 40 identifies the page features. Processor 40 analyzes the page and divides it into blank and non-blank portions. The non-blank portions are analyzed to distinguish text regions from non-text regions. After determining the layout of the page, processor 40 uses black-to-white transitions to determine degrees of skew. In step 1008, horizontal white spaces are identified to separate lines of text. In step 1010, vertical white spaces are identified within each line of text to thereby separate individual words and characters from each other. In Step 1014, a character recognition algorithm is used in an attempt to recognize each individual character." Ehrhart at 12:55-66. "As embodied herein and depicted in FIG. 11, a flow chart showing a method for performing OCR in accordance with yet another embodiment of the present invention is disclosed. In step 1100, reader 10 produces a bit-mapped image of the page. Subsequently, processor 40 finds lines of text in the image, locates the white spaces in each line, and isolates the characters. In step 1108, processor 40 performs character recognition, either OCR-A or OCR-B, as desired. The decoded characters are stored in memory." Ehrhart at 13:1-9.

Element	'529 Claim Recitation	Exemplary Citations in Reference

		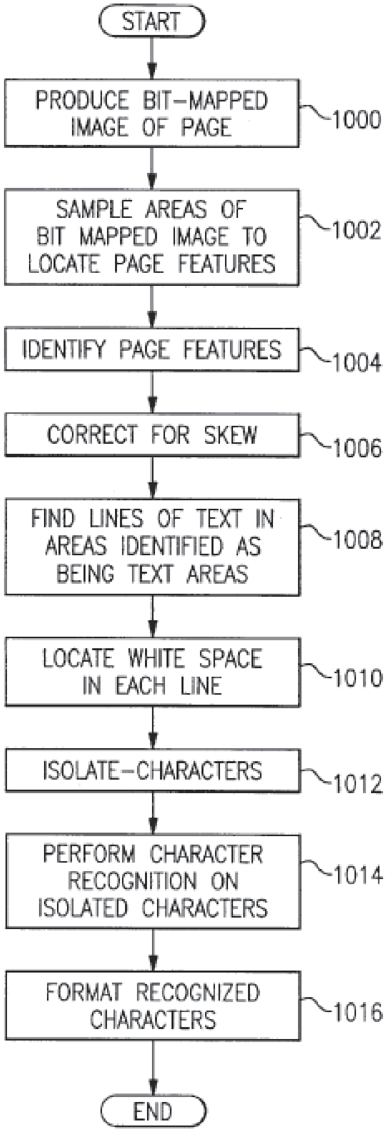START PRODUCE BIT-MAPPED IMAGE OF PAGE 1000 SAMPLE AREAS OF BIT MAPPED IMAGE TO LOCATE PAGE FEATURES 1002 IDENTIFY PAGE FEATURES 1004 CORRECT FOR SKEW 1006 FIND LINES OF TEXT IN AREAS IDENTIFIED AS BEING TEXT AREAS 1008 LOCATE WHITE SPACE IN EACH LINE 1010 ISOLATE-CHARACTERS 1012 PERFORM CHARACTER RECOGNITION ON ISOLATED CHARACTERS 1014 FORMAT RECOGNIZED CHARACTERS 1016 END FIG.10

Element	'529 Claim Recitation	Exemplary Citations in Reference
		Ehrhart at Fig. 10.

		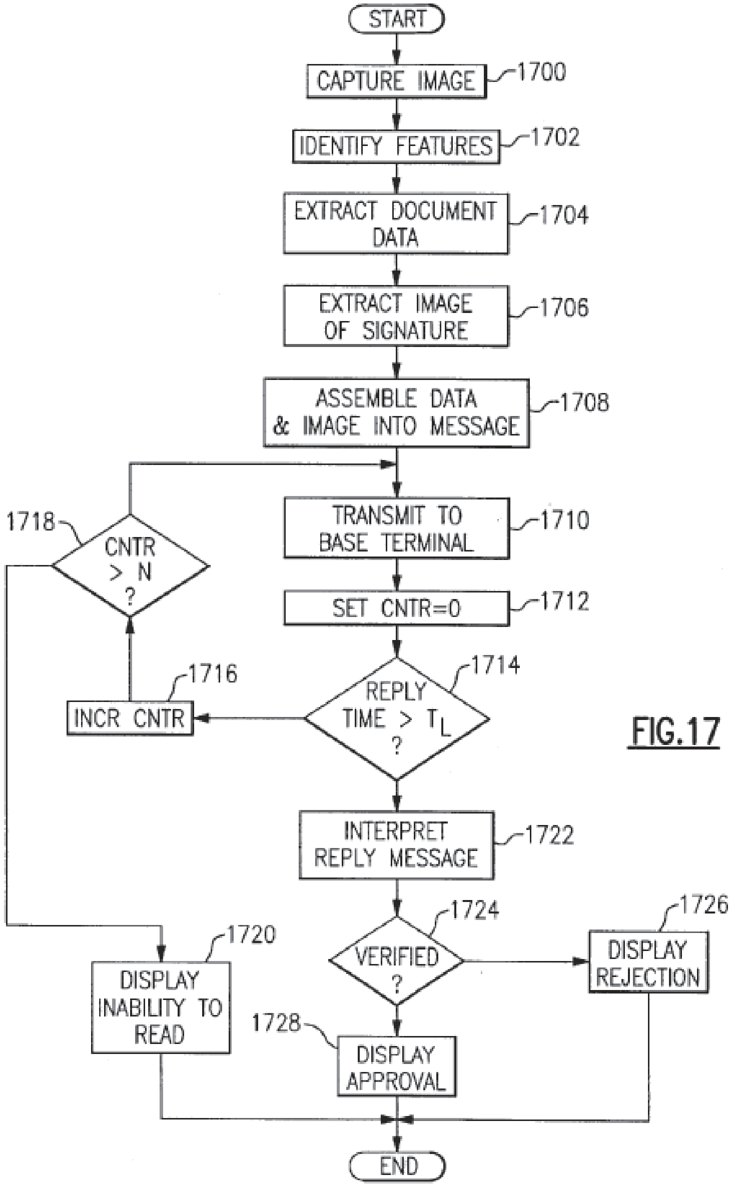START CAPTURE IMAGE 1700 IDENTIFY FEATURES 1702 EXTRACT DOCUMENT DATA 1704 EXTRACT IMAGE OF SIGNATURE 1706 ASSEMBLE DATA & IMAGE INTO MESSAGE 1708 TRANSMIT TO BASE TERMINAL 1710 SET CNTR=0 1712 REPLY TIME > T_L ? 1714 INCR CNTR 1716 CNTR > N ? 1718 INTERPRET REPLY MESSAGE 1722 VERIFIED ? 1724 DISPLAY REJECTION 1726 DISPLAY INABILITY TO READ 1720 DISPLAY APPROVAL 1728 END FIG.17

Element	'529 Claim Recitation	Exemplary Citations in Reference
		Ehrhart at Fig. 17. To the extent it is found that Ehrhart does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Ehrhart does not anticipate this claim, Ehrhart renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Ehrhart with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1G]	return the information to the network-enabled device; and	Ehrhart discloses return the information to the network-enabled device. *Id.* at [1Pre, 1A, 1B, 1C, 1D, 1E, 1F]. "As embodied herein and depicted in FIG. 13, an example of image association in accordance with the present invention is disclosed. One or ordinary skill in the art will recognize that associated images 1300 can be disposed on paper, displayed electronically on display 60, or displayed electronically sing other electronic means, such as a computer monitor. In this example, the first image captured is color photograph 1302 which shows a damaged parcel. The second image captured is bar code 1304 affixed to the side of the damaged parcel. Processor 40 decodes bar code 1304 and associates decoded bar code data 1306 with color photograph 1302. In this example, the user elected to associate a third image, signature 1308. Thus, personnel viewing record 1300 may reasonably conclude that a damaged parcel was delivered to Company XYZ, and that the person signing for the parcel delivery was someone named John W. Smith." Ehrhart at 13:27-43.

Element	'529 Claim Recitation	Exemplary Citations in Reference
		To the extent it is found that Ehrhart does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Ehrhart does not anticipate this claim, Ehrhart renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Ehrhart with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1H]	the network-enabled device further programmed to present the information related to the object to a user.	Ehrhart discloses the network-enabled device further programmed to present the information related to the object to a user. *Id.* at [1Pre, 1A, 1B, 1C, 1D, 1E, 1F, 1G]. "In some embodiments, processor 40 may display the acquired image on display 60 during this step." Ehrhart at 9:35-38. "As embodied herein, and depicted in FIGS. 1A-1D, perspective views of the optical reader in accordance with various embodiments of the present invention are disclosed. FIG. 1A shows the underside of hand held wireless optical reader 10. FIG. 1B shows the top of the optical reader depicted in FIG. 1A. Optical reader 10 includes housing 100, antenna 102, window 104 and trigger 12. Window 104 accommodates illumination assembly 20 and imaging assembly 30. As shown in FIG. 1B, the top side of reader 10 includes function keys 14, alphanumeric key pad 16, and display 60. In one embodiment, function

Element	'529 Claim Recitation	Exemplary Citations in Reference
		keys 14 include an enter key and up and down cursor keys. FIG. 1C is also a hand held wireless optical reader 10. Reader 10 includes function keys 14, alphanumeric key pad 16, writing stylus 18, display 60, and signature block 62. Stylus 18 is employed by a user to write his signature in signature block 62. FIG. 1D shows yet another embodiment of optical reader 10 of the present invention. In this embodiment, reader 10 includes a gun-shaped housing 100. Display 60 and keypad 16 are disposed on a top portion of gun-shaped housing 100, whereas trigger 12 is disposed on the underside of the top portion of housing 100. Housing 100 also includes window 104 that accommodates illumination assembly 20 and imaging assembly 30. Wire 106 is disposed at the butt-end of housing 100. Wire 106 provides optical reader 10 with a hard wired communication link for external devices such as a host processor or other data collection devices." Ehrhart at 5:51-6:10. To the extent it is found that Ehrhart does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Ehrhart does not anticipate this claim, Ehrhart renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Ehrhart with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

2. Claim 4

Claim	'529 Claim Recitation	Exemplary Citations in Reference

4	The system of claim 1, wherein the service is further programmed to identify the object by applying multiple algorithms to at least a portion of the image or of the data, and identify the object as a function of a confidence levels associated with results of the multiple algorithms.	Ehrhart discloses the system of claim 1, wherein the service is further programmed to identify the object by applying multiple algorithms to at least a portion of the image or of the data, and identify the object as a function of a confidence levels associated with results of the multiple algorithms. "The user may also click on OCR/Text icon 660. Clicking icon 660 provides the user with a check validation mode, a text scanning mode, or a bi-tonal image capture mode. The check validation mode is performed in conjunction with network services." Ehrhart at 8:58-62. "In steps 1712 and 1714, processor 40 initializes a counter and begins waiting for a reply from the host computer. In steps 1714-1718, if the reply is not received within time limit TL, the counter CNTR is incremented and the message is re-transmitted. After several attempts, if CNTR>N (N being an integer), processor 40 outputs a fault message. If the reply message is received within time limit TL, processor interprets the reply in step 1722. If the extracted data and the signature match information stored in the database accessible by the host computer, an approval message is displayed. If the extracted data and the signature do not match information stored in the database accessible by the host computer, a disapproval message is displayed. The dynamic signature verification embodiment is similar to the static embodiment described immediately above. In the dynamic version, the user provides his signature using stylus 18 and signature block 62, as shown in FIG. 1C. Signature block 62 provides processor 40 with the dynamic parameters recorded during signature. The dynamic parameters are transmitted to a host processor, as described above." Ehrhart at 15:31-50. To the extent it is found that Ehrhart does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Ehrhart does not anticipate this claim, Ehrhart renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's

		Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Ehrhart with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

3. Claim 20

Claim	'529 Claim Recitation	Exemplary Citations in Reference
20	The system of claim 1, wherein the data processing operation is selected from the group consisting of contrast enhancement, noise removal, and de-blurring.	Ehrhart discloses the system of claim 1, wherein the data processing operation is selected from the group consisting of contrast enhancement, noise removal, and de-blurring. "The user selects the color photography mode by clicking on icon 664. This mode allows the user to select an automatic imaging mode wherein optical reader 10 makes the imaging adjustments(e.g., exposure, etc.) or a manual mode that allows the user to adjust imager settings as he pleases." Ehrhart at 7:10-18. "The user Selects the color photography mode by clicking on icon 664. This mode allows the user to select an automatic imaging mode wherein optical reader 10 makes the imaging adjustments(e.g., exposure, etc.) or a manual mode that allows the user to adjust imager Settings as he pleases." Ehrhart at 9:11-15. To the extent it is found that Ehrhart does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Ehrhart does not anticipate this claim, Ehrhart renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's

		Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Ehrhart with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

Nantworks, LLC v. Bank of America Corporation
Central District of California, No. 2:20-cv-7872
Invalidity Chart for U.S. Patent 7,881,529 based on U.S. Patent No. 6,097,834 to Krouse ("Krouse")

Exhibit G-18

Bank of America Corporation ("BofA" or "Defendant") provides this chart subject to all reservations, objections, statements, and disclaimers set forth herein and in BofA's Final Invalidity Contentions Cover Pleading, as well as any amendment, supplement, or modification thereof, which are incorporated herein by reference in their entirety.

On information and belief, the application that matured into Krouse was filed on June 13, 1997, claiming the priority date of June 13, 1997. Krouse published at least as of August 1, 2000, and is available as prior art at least under 35 U.S.C. § 102(a).

As shown in the chart below, Krouse anticipates asserted claims 4, 20 of the '529 Patent. To the extent Krouse is found not to expressly disclose certain limitations in the asserted claims, such limitations are inherent. To the extent Krouse is found not to anticipate any asserted claims or claim elements of the '529 Patent, the reference nevertheless renders those claims or claim elements obvious under 35 U.S.C. § 103, either alone or in combination with other art identified herein or in the cover pleading of BofA's Invalidity Contentions, for the reasons identified in the cover pleading. These Invalidity Contentions are not an admission by BofA that the accused products, including any current or past versions of the accused products, are covered by, or infringe the asserted claims, but are based instead on the recognition that if the claims are interpreted to be broad enough to encompass the accused products, the claims must also be construed to have that same scope when considering whether they are invalid.

The following chart is based on the Court's claim construction rulings (Dkt. Nos. 145, 153, 236), and to the extent not contradictory, Plaintiffs' apparent interpretations of the claims in Plaintiffs' "Preliminary" Final Infringement Contentions. BofA notes that in many instances, Plaintiffs' "Preliminary" Final Infringement Contentions ignore and/or contradict the Court's claim construction rulings. BofA does not accept the assumptions concerning the scope and meaning implicit in Plaintiffs' Preliminary Final Infringement Contentions, to the extent those assumptions are discernible, and reserves the right to challenge Plaintiffs' proposed (or implied) constructions. Furthermore, Plaintiffs' "Preliminary" Final Contentions fail to comply with the level of specificity required by SPR 2.1.2 and 2.1.3, and include improper reservation of rights to supplement with "additional ways in which BoA Accused Products infringe." Any alleged deficient discovery is of Plaintiffs own making. Most if not all of the deficiencies in Plaintiffs' Preliminary Final Contentions have been documented by Defendants through multiple letters since at least May 18, 2021 and up to Bank of America's most recent letter of August 19, 2023. Plaintiffs' shifting infringement contentions, deficiencies in its infringement contentions, and changing claim scope have prejudiced BofA from being able to reasonably prepare its defenses, including its invalidity contentions. BofA also reserves the right to revise and supplement these charts if and when Plaintiffs are permitted to provide revised or supplemental Infringement Contentions. Citations given in the chart below are merely representative of the respective elements to provide notice to Plaintiffs of BofA's invalidity theories and are not meant to be exhaustive.

1. Claim 1

Element	'529 Claim Recitation	Exemplary Citations in Reference
[1Pre]	A system comprising:	Krouse discloses a system. Krouse at Abstract - "In accordance with the present invention, financial transaction processing systems and methods are provided. One preferred embodiment of a method according to one aspect of the present invention includes generating an optically scanned image of at least a portion of document containing visual data, in a particular format, representing information related to the financial transaction. Recognition characteristics are generated from the scanned image and are compared to respective sets of reference recognition characteristics generated from respective other transaction documents having different respective formats to determine therefrom whether the particular format of the visual data matches one of the respective formats of the other documents. When such a match is found to exist, location is determined of a field in the scanned image to which optical character recognition may be applied to generate therefrom the information, based upon the respective format found to match the particular format of the visual data. Optical character recognition is then utilized to generate said visual data from said location."
[1A]	a camera that captures an image;	Krouse discloses a camera that captures an image. *Id.* at [1Pre]. Krouse at 1:6-12: - The present invention relates generally to financial transaction processing Systems and methods, and more. Specifically, to Such Systems and methods wherein at least one document containing financial transaction-related information is optically Scanned to generate at least one computer-

Element	'529 Claim Recitation	Exemplary Citations in Reference
		readable image from which the information maybe extracted for use in processing the transaction. Krouse at 6:4–14 • "Each of the terminals 18, 20, 22 includes an optical scanner 34 for scanning negotiable instrument documents, (exemplified by check 110 shown in FIG. 6) tendered at point of sale in payment for goods or services provided by the merchant or retailer associated with that point of sale, to generate respective computer-readable scanned images of the documents. Preferably, scanner 34 comprises a conventional monotonal image scanner having a resolution of at least 300 dots-per-inch and being adapted for being controlled by local controller/processor 40 to scan one or both sides of a document 110 placed in the scanner 34." To the extent it is found that Krouse does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Krouse does not anticipate this claim, Krouse renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Krouse with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1B]	a network-enabled device that conducts a data processing operation on at least a portion of the image to produce data, and sends the data to a service;	Krouse discloses a network-enabled device that conducts a data processing operation on at least a portion of the image to produce data, and sends the data to a service. *Id.* at [1Pre-A].

Element	'529 Claim Recitation	Exemplary Citations in Reference
		Krouse at 4:13–57 • "Also, in a second aspect of the present invention, a system and method are provided for use in processing a financial transaction based at least in part upon visual data formed on a document. The visual data represents information related to the transaction and is in a particular format. One preferred embodiment of the system of the second aspect of the present invention includes an optical scanner for generating a scanned image of at least a portion of the document containing the visual representation. An image characterization generator is provided for generating recognition characteristics from the scanned image, and a recognition characteristic comparitor is provided for comparing the recognition characteristics to respective sets of reference recognition characteristics generated from respective other transaction documents having different respective formats and for determining therefrom whether the particular format of the visual data on the scanned document matches one of the respective formats of the other documents. A field location generator is also provided for determining, based upon the one respective format when the particular format is determined to match the one respective format, location of a field in the scanned image to which optical character recognition (OCR) may be applied to generate therefrom the transaction-related information represented by the visual data. OCR is then applied by an image processor to this location to generate the visual data for use in processing the transaction. One preferred embodiment of the method according to the second aspect of the present invention includes generating an optically scanned image of at least a portion of document containing visual data, in a particular format, representing information related to the financial transaction. Recognition characteristics are generated from the scanned image and are compared to respective sets of reference recognition characteristics generated from respective other transaction documents having different respective formats

Element	'529 Claim Recitation	Exemplary Citations in Reference
		to determine therefrom whether the particular format of the visual data matches one of the respective formats of the other documents. When such a match is found to exist, location is determined of a field in the scanned image to which optical character recognition may be applied to generate therefrom the information, based upon the respective format found to match the particular format of the visual data. Optical character recognition is then utilized to generate said visual data from said location." Krouse at 5:60–6:14 • "As shown in FIG. 1, system **10** comprises at least one **18**, and preferably a plurality **18, 20, 22** of point of sale payment terminals. Each terminal **18, 20, 22** preferably is located in its own respective point of sale location **12, 14, 16** (e.g., at a merchant's or retailer's place of business, and/or at the location of third party transaction agent acting on behalf of the merchant or retailer). It should be understood that although three terminals **18, 20, 22** are shown as being comprised in system **10**, system **10** could be modified to include any number of such terminals without departing from this aspect of the present invention. Each of the terminals **18, 20, 22** includes an optical scanner **34** for scanning negotiable instrument documents, (exemplified by check **110** shown in FIG. 6) tendered at point of sale in payment for goods or services provided by the merchant or retailer associated with that point of sale, to generate respective computer-readable scanned images of the documents. Preferably, scanner **34** comprises a conventional monotonal image scanner having a resolution of at least 300 dots-per-inch and being adapted for being controlled by local controller/processor **40** to scan one or both sides of a document **110** placed in the scanner **34**." Krouse at 7:7–31 • "Once the aforesaid data (i.e., payment amount being tendered and payee) has been input via the device **30**, document **110** has been placed in

Element	'529 Claim Recitation	Exemplary Citations in Reference
		the scanner **34**, and presence of the document **110** in the scanner **34** has been signalled to the terminal operator, the terminal operator continues processing the document **110** by commanding the processor **40** via the device **30** to initiate scanning of the document **110**. In response, the processor **40** issues commands to the scanner **34** to scan the document **110** placed therein, whereupon, the scanner **34** scans at least one entire side **112** of the draft **110** so as to generate a computer-readable scanned image of at least that entire side **112** of the draft document **110** (i.e., computer-readable data from which an appropriately provisioned computer may generate a display image of the at least one scanned side **112** of the document **110**). The computer-readable scanned image data is then transmitted by the scanner **34** to the processor **40**; the processor **40** then transmits to the optical character recognizer **32** a copy of the scanned image data generated by the scanner **34** in order to permit the recognizer **32** to process same in the manner described below. A separate copy of the scanned image data is also transmitted to the storage system **42** where it is stored in association with the previously stored payment amount and payee data of the negotiable instrument being processed." Krouse at 13:12 – 14:6 • " When the scanned image is first received by the generator 204, it is processed by an image de-skewing processor 210, such as Scanfix Image Processor available from TMS Sequoia, of Burlingame, Calif., which adjusts the scanned image data by rotating and/or translating the scanned image so as to cause two adjacent edges (e.g., the top and left side edges) of the scanned image (i.e., the edges of the scanned image generated from the top and left side edges of the document 400) to be aligned with the ordinate and abscissa of an orthogonal coordinate system whose center is chosen so as to coincide with the dark pixel of the scanned image closest to those two adjacent edges. Once the processor 210 has accomplished this processing of the scanned image data, the scanned image data is next processed by the noise eliminator 212, which processes the image data to

Element	'529 Claim Recitation	Exemplary Citations in Reference
		remove dark pixels (i.e., non-whitespace pixels) therefrom that were erroneously generated as a result of noise. The noise eliminator 212 accomplishes this by locating isolated groups of contiguous dark pixels in the scanned image data having respective sizes (i.e., respective numbers of pixels) smaller than three dark pixels; these isolated groups of contiguous dark pixels are then removed from the scanned image data and replaced with whitespace by the processor 212. • After being processed by the noise eliminator 212, the scanned image data is processed by the image enhancer processor 214, which processor 214 enhances the dark pixels remaining in the scanned image after processing of the scanned image data by the noise eliminator processor 212. Processor 214 accomplishes this by locating the dark pixels remaining in the scanned image after processing of the scanned image by the processor 212, and by extending by predetermined numbers of pixels (e.g., 5 pixels in each of the positive and negative ordinate directions, and 25 pixels in each of the positive and negative abscissa directions) each of these remaining dark pixels in both of the two mutually orthogonal directions of the coordinate system to which the scanned image was aligned by the de-skewing processor 210. That is, predetermined numbers of whitespace pixels surrounding and extending from the remaining dark pixels in the two mutually orthogonal directions of the coordinate system are replaced with dark pixels. Any resulting extension of the dark pixels that would extend beyond the edge of the scanned image is prevented from doing so by truncating the extension by a predetermined number of pixels sufficient to prevent the dark pixel extensions from coming within 5 pixels of the edges of the scanned image. • After being processed by the image enhancer 214, the result, enhanced scanned image is processed by dark pixel group locator processor 222. Locator processor 222 then generates respective imaginary rectangles around the respective groups of contiguous dark pixels such that the respective rectangles each include two sides are parallel to the ordinate

Element	'529 Claim Recitation	Exemplary Citations in Reference
		and two sides are parallel to the abscissa of the coordinate system to which the scanned image was aligned by the processor 210, and such that the respective rectangles include all of the respective dark pixels of each of the respective groups of contiguous dark pixels. The locator processor 222 then selects those imaginary rectangles, closest to the origin of the coordinate system, whose respective longitudinal dimensions exceed a predetermined threshold (e.g., 1 inch), and determines the respective coordinates of at least one predetermined vertex (e.g., the respective vertex having the minimal abscissal coordinate value and the maximal ordinate coordinate value for the respective rectangle) and dimensions of a predetermined number (e.g., no more than 25) of these selected imaginary rectangles. The respective vertex coordinates and dimensions of these imaginary rectangles are output by the locator processor 222 as recognition characteristics for the scanned document 400, which recognition characteristics are transmitted from the generator 204 to the recognition characteristic comparison processor 220." To the extent it is found that Krouse does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Krouse does not anticipate this claim, Krouse renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Krouse with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1C]	the service programmed to receive the data;	Krouse discloses the service programmed to receive the data. *Id.* at [1Pre-B].

Element	'529 Claim Recitation	Exemplary Citations in Reference
		Krouse at 7:7–31 • "Once the aforesaid data (i.e., payment amount being tendered and payee) has been input via the device **30**, document **110** has been placed in the scanner **34**, and presence of the document **110** in the scanner **34** has been signalled to the terminal operator, the terminal operator continues processing the document **110** by commanding the processor **40** via the device **30** to initiate scanning of the document **110**. In response, the processor **40** issues commands to the scanner **34** to scan the document **110** placed therein, whereupon, the scanner **34** scans at least one entire side **112** of the draft **110** so as to generate a computer-readable scanned image of at least that entire side **112** of the draft document **110** (i.e., computer-readable data from which an appropriately provisioned computer may generate a display image of the at least one scanned side **112** of the document **110**). The computer-readable scanned image data is then transmitted by the scanner **34** to the processor **40**; the processor **40** then transmits to the optical character recognizer **32** a copy of the scanned image data generated by the scanner **34** in order to permit the recognizer **32** to process same in the manner described below. A separate copy of the scanned image data is also transmitted to the storage system **42** where it is stored in association with the previously stored payment amount and payee data of the negotiable instrument being processed." Krouse at 7:38–67 • "One such standardized portion in check **110** is Magnetic Ink Character Recognizable (MICR) line **94**, which is positioned at the bottom of the check **110** and in check **110**, includes three fields **114, 116, 118** of printed digits separated by special control characters. Fields **114, 116** contain the financial institution transaction routing number and account number on which the check **110** is drawn, and field **118** contains the check number of the check **110**. It should be noted, however, that other than the left most nine digits (and associated separator characters) of the

Element	'529 Claim Recitation	Exemplary Citations in Reference
		MICR line **94**, the contents of said line may vary, depending upon the financial institution upon which the negotiable instrument is drawn. System **10** may be adapted to recognize and parse information contained in different types of negotiable instrument MICR lines (not shown). Based upon the standardized location, size, style, etc. of the MICR line **94** in document **110** and the fields **114, 116, 118** contained therein, recognizer **32** utilizes optical character recognition and image processing techniques to extract from the scanned image data the data generated by the scanner **34** from scanning the MICR line portion **94** of the document **110**, and once thus extracted, to generate numeric transaction data indicative of the visual information (i.e., the numbers printed in fields **114, 116, 118**) in MICR line **94**. This numeric data is then transmitted to the processor **40**, which transmits a copy of the numeric data to the storage system **42** where it is stored in association with the scanned image data and the payee and payment amount data previously stored therein." To the extent it is found that Krouse does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Krouse does not anticipate this claim, Krouse renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Krouse with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1D]	identify an object within the image:	Krouse discloses identify an object within the image.

Element	'529 Claim Recitation	Exemplary Citations in Reference
		Id. at [1Pre-C]. • "Accordingly, in a first aspect of the present invention, a System and method are provided for processing an electronic financial transaction based upon visual information related to the transaction contained on a document tendered by a party to the transaction. One preferred embodiment of the System of this aspect of the present invention includes an optical Scanner for generating a computer-readable Scanned image of the document, and an optical character recognizer for generating, from the Scanned image, transaction data indicative of at least a portion of the visual information." Krouse at 7:38–67 • "One such standardized portion in check **110** is Magnetic Ink Character Recognizable (MICR) line **94**, which is positioned at the bottom of the check **110** and in check **110**, includes three fields **114, 116, 118** of printed digits separated by special control characters. Fields **114, 116** contain the financial institution transaction routing number and account number on which the check **110** is drawn, and field **118** contains the check number of the check **110**. It should be noted, however, that other than the left most nine digits (and associated separator characters) of the MICR line **94**, the contents of said line may vary, depending upon the financial institution upon which the negotiable instrument is drawn. System **10** may be adapted to recognize and parse information contained in different types of negotiable instrument MICR lines (not shown). Based upon the standardized location, size, style, etc. of the MICR line **94** in document **110** and the fields **114, 116, 118** contained therein, recognizer **32** utilizes optical character recognition and image processing techniques to extract from the scanned image data the data generated by the scanner **34** from scanning the MICR line portion **94** of the document **110**, and once thus extracted, to generate numeric transaction data indicative of the visual information (i.e., the numbers printed in fields

Element	'529 Claim Recitation	Exemplary Citations in Reference
		114, 116, 118) in MICR line **94**. This numeric data is then transmitted to the processor **40**, which transmits a copy of the numeric data to the storage system **42** where it is stored in association with the scanned image data and the payee and payment amount data previously stored therein." Krouse at 12:37-52: • "Processing of a financial transaction using System 200 begins by placing a bill document (Such as bill document 400 shown in FIG. 8A) to be processed into scanner 208 so as to permit at least one side 402 of the document 400 containing an optical character recognition line (OCR) 404 which contains numerical data 405 visually displaying information related to the transaction to be processed (e.g., amount due 406 and customer account number 408). It should be understood that although the functional components and operation of system 200 will be described in connection with processing of document 400, advantageously, System 200 can be used to Scan other types of billing documents (e.g., those illustrated in FIGS. 8B-8F), in accordance with this aspect of the present invention." To the extent it is found that Krouse does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Krouse does not anticipate this claim, Krouse renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Krouse with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

Element	'529 Claim Recitation	Exemplary Citations in Reference
[1E]	distinguish an object present in the image from others using a database that stores data characteristics of target objects;	Krouse discloses distinguish an object present in the image from others using a database that stores data characteristics of target objects. *Id.* at [1Pre-D]. Krouse at Abstract • "In accordance with the present invention, financial transaction processing systems and methods are provided. One preferred embodiment of a method according to one aspect of the present invention includes generating an optically scanned image of at least a portion of document containing visual data, in a particular format, representing information related to the financial transaction. Recognition characteristics are generated from the scanned image and are compared to respective sets of reference recognition characteristics generated from respective other transaction documents having different respective formats to determine therefrom whether the particular format of the visual data matches one of the respective formats of the other documents. When such a match is found to exist, location is determined of a field in the scanned image to which optical character recognition may be applied to generate therefrom the information, based upon the respective format found to match the particular format of the visual data. Optical character recognition is then utilized to generate said visual data from said location." Krouse at 4:13–57 • "Also, in a second aspect of the present invention, a system and method are provided for use in processing a financial transaction based at least in part upon visual data formed on a document. The visual data represents information related to the transaction and is in a particular format. One preferred embodiment of the system of the second aspect of the present invention includes an optical scanner for generating a scanned image of at least a portion of the document containing the visual

Element	'529 Claim Recitation	Exemplary Citations in Reference
		representation. An image characterization generator is provided for generating recognition characteristics from the scanned image, and a recognition characteristic comparitor is provided for comparing the recognition characteristics to respective sets of reference recognition characteristics generated from respective other transaction documents having different respective formats and for determining therefrom whether the particular format of the visual data on the scanned document matches one of the respective formats of the other documents. A field location generator is also provided for determining, based upon the one respective format when the particular format is determined to match the one respective format, location of a field in the scanned image to which optical character recognition (OCR) may be applied to generate therefrom the transaction-related information represented by the visual data. OCR is then applied by an image processor to this location to generate the visual data for use in processing the transaction. One preferred embodiment of the method according to the second aspect of the present invention includes generating an optically scanned image of at least a portion of document containing visual data, in a particular format, representing information related to the financial transaction. Recognition characteristics are generated from the scanned image and are compared to respective sets of reference recognition characteristics generated from respective other transaction documents having different respective formats to determine therefrom whether the particular format of the visual data matches one of the respective formats of the other documents. When such a match is found to exist, location is determined of a field in the scanned image to which optical character recognition may be applied to generate therefrom the information, based upon the respective format found to match the particular format of the visual data. Optical character recognition is then utilized to generate said visual data from said location." Krouse at 14:64–15:67

Element	'529 Claim Recitation	Exemplary Citations in Reference
		• "Processor **220** compares the recognition characteristics of the particular document **400** being processed with the sets of reference recognition characteristics stored in the archive **228**. More specifically, the vertex coordinates and dimensions of the imaginary rectangles generated by the generator **204** from the scanned image of the document **400** are compared with respective reference recognition characteristics (i.e., vertex coordinates and dimensions of imaginary rectangles) generated by the process used to generate the recognition characteristics from the scanned image of document **400**, to determine whether a match exists between the recognition characteristics from the scanned image of document **400** and any one of the sets of reference recognition characteristics stored in the database archive system **228**. Preferably, this is accomplished by first comparing the recognition characteristics from the scanned image of document **400** to the sets of reference characteristics, and unless the respective dimensions of at least three of the imaginary rectangles from the scanned image of document **400** match to within a predetermined error tolerance (e.g., 2 percent) of corresponding respective vertices and dimensions in a respective set of reference recognition characteristics stored in the system **228**, processor **220** signals that the format of the scanned document 400 does not match any of the formats of the reference documents from which the reference recognition characteristics stored in the database archive **228** were generated. If such a "no match" condition is determined by the processor **220** to exist, the bill document **400** may be sent to the reference characterization generator station to permit the human operator to determine whether to input the reference characteristics and other bill document identification information (i.e., the other types of information stored in association with each set of reference recognition characteristics previously stored in the database archive **228**) manually for storage in the archive system **228** as a new set of reference recognition characteristics to permit other bills having the same format as document **400** to be processable by terminal **202** in the future, or to process the bill payment transaction as a "one-of-a-kind" transaction (i.e., without storing such

Element	'529 Claim Recitation	Exemplary Citations in Reference
		recognition characteristics and information in the database archive **228** for future use by the processor **220**). Otherwise, if the predetermined match error tolerance conditions are found to exist, processor **220** selects those sets of reference recognition characteristics that are found to satisfy such conditions, and calculates the displacement differences that exist between each of the imaginary rectangles of the scanned image of the document **400** and those described in these selected sets of reference recognition characteristics (i.e., the displacements that must be applied to the imaginary rectangles of the scanned image of document **400** to accomplish a "best fit" match of same with the corresponding imaginary rectangles described in these selected sets of reference recognition characteristics). Such "best fit" match conditions are determined to exist for given corresponding rectangles of the recognition characteristics of the scanned image of the document **400** and a given respective set of reference recognition characteristics, when the absolute pixel displacement differences between the respective abscissa and ordinate coordinates of the vertices of the rectangles are within 40 pixels and 60 pixels of each other, respectively, and the respective lengths and widths of the rectangles are within 20 pixels and 2 pixels of each other, respectively. Once these absolute displacement differences have been calculated by the processor **220**, processor **220** "scores" the degree of best fit match condition that exists between the imaginary rectangles of the scanned image of the document **400** and those described in the aforesaid selected sets of reference recognition characteristics, whereby to determine which one of the aforesaid selected sets of reference recognition characteristics best matches the recognition characteristics of the scanned document **400**." To the extent it is found that Krouse does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Krouse does not anticipate this claim, Krouse renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's

Element	'529 Claim Recitation	Exemplary Citations in Reference
		Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Krouse with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1F]	associate the object with information; and	Krouse discloses associate the object with information. *Id.* at [1Pre-E]. Krouse at 3:43–4:12 • "One preferred embodiment of the system of this aspect of the present invention includes an optical scanner for generating a computer-readable scanned image of the document, and an optical character recognizer for generating, from the scanned image, transaction data indicative of at least a portion of the visual information. This embodiment of the system of this aspect of the present invention also includes an authorization record generator for generating, based upon the transaction data, a record of the transaction for being assented to by the party tendering the document prior to completion of the transaction, and for generating from the record once it is assented to by the party, a computer-readable record of assent of the party to the transaction. A collection system is also included in this embodiment of the system of this aspect of the present invention for generating in response to the transaction data and the computer-readable record, an electronic financial transaction request for being used by an electronic transaction system to execute the electronic financial transaction. One embodiment of the method of the first aspect of the present invention includes generating a computer-readable scanned image of the document, and generating, from the scanned image, transaction data

Element	'529 Claim Recitation	Exemplary Citations in Reference
		indicative of at least a portion of the visual information contained on the document. A record is generated, from the transaction data, of the electronic financial transaction being processed, which record is for being assented to by the party tendering the document, prior to completion of said transaction. Once the record has been assented to by the party tendering the document, a computer-readable record of the party's assent to the transaction is generated from the assented-to record. An electronic financial transaction request is then generated in response to the transaction data and the computer-readable assented-to record, which request is for being used by an electronic financial transaction system to execute the electronic financial transaction." Krouse at 7:38–67 • "One such standardized portion in check **110** is Magnetic Ink Character Recognizable (MICR) line **94**, which is positioned at the bottom of the check **110** and in check **110**, includes three fields **114, 116, 118** of printed digits separated by special control characters. Fields **114, 116** contain the financial institution transaction routing number and account number on which the check **110** is drawn, and field **118** contains the check number of the check **110**. It should be noted, however, that other than the left most nine digits (and associated separator characters) of the MICR line **94**, the contents of said line may vary, depending upon the financial institution upon which the negotiable instrument is drawn. System **10** may be adapted to recognize and parse information contained in different types of negotiable instrument MICR lines (not shown). Based upon the standardized location, size, style, etc. of the MICR line **94** in document **110** and the fields **114, 116, 118** contained therein, recognizer **32** utilizes optical character recognition and image processing techniques to extract from the scanned image data the data generated by the scanner **34** from scanning the MICR line portion **94** of the document **110**, and once thus extracted, to generate numeric transaction data

Element	'529 Claim Recitation	Exemplary Citations in Reference
		indicative of the visual information (i.e., the numbers printed in fields **114, 116, 118**) in MICR line **94**. This numeric data is then transmitted to the processor **40**, which transmits a copy of the numeric data to the storage system **42** where it is stored in association with the scanned image data and the payee and payment amount data previously stored therein." Krouse at 14:64–15:67 • "Processor **220** compares the recognition characteristics of the particular document **400** being processed with the sets of reference recognition characteristics stored in the archive **228**. More specifically, the vertex coordinates and dimensions of the imaginary rectangles generated by the generator **204** from the scanned image of the document **400** are compared with respective reference recognition characteristics (i.e., vertex coordinates and dimensions of imaginary rectangles) generated by the process used to generate the recognition characteristics from the scanned image of document **400**, to determine whether a match exists between the recognition characteristics from the scanned image of document **400** and any one of the sets of reference recognition characteristics stored in the database archive system **228**. Preferably, this is accomplished by first comparing the recognition characteristics from the scanned image of document **400** to the sets of reference characteristics, and unless the respective dimensions of at least three of the imaginary rectangles from the scanned image of document **400** match to within a predetermined error tolerance (e.g., 2 percent) of corresponding respective vertices and dimensions in a respective set of reference recognition characteristics stored in the system **228**, processor **220** signals that the format of the scanned document 400 does not match any of the formats of the reference documents from which the reference recognition characteristics stored in the database archive **228** were generated. If such a "no match" condition is determined by the processor **220** to exist, the bill document **400** may be sent to the reference characterization generator station to permit the

Element	'529 Claim Recitation	Exemplary Citations in Reference
		human operator to determine whether to input the reference characteristics and other bill document identification information (i.e., the other types of information stored in association with each set of reference recognition characteristics previously stored in the database archive **228**) manually for storage in the archive system **228** as a new set of reference recognition characteristics to permit other bills having the same format as document **400** to be processable by terminal **202** in the future, or to process the bill payment transaction as a "one-of-a-kind" transaction (i.e., without storing such recognition characteristics and information in the database archive **228** for future use by the processor **220**). Otherwise, if the predetermined match error tolerance conditions are found to exist, processor **220** selects those sets of reference recognition characteristics that are found to satisfy such conditions, and calculates the displacement differences that exist between each of the imaginary rectangles of the scanned image of the document **400** and those described in these selected sets of reference recognition characteristics (i.e., the displacements that must be applied to the imaginary rectangles of the scanned image of document **400** to accomplish a "best fit" match of same with the corresponding imaginary rectangles described in these selected sets of reference recognition characteristics). Such "best fit" match conditions are determined to exist for given corresponding rectangles of the recognition characteristics of the scanned image of the document **400** and a given respective set of reference recognition characteristics, when the absolute pixel displacement differences between the respective abscissa and ordinate coordinates of the vertices of the rectangles are within 40 pixels and 60 pixels of each other, respectively, and the respective lengths and widths of the rectangles are within 20 pixels and 2 pixels of each other, respectively. Once these absolute displacement differences have been calculated by the processor **220**, processor **220** "scores" the degree of best fit match condition that exists between the imaginary rectangles of the scanned image of the document **400** and those described in the

Element	'529 Claim Recitation	Exemplary Citations in Reference
		aforesaid selected sets of reference recognition characteristics, whereby to determine which one of the aforesaid selected sets of reference recognition characteristics best matches the recognition characteristics of the scanned document **400**." To the extent it is found that Krouse does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Krouse does not anticipate this claim, Krouse renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Krouse with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1G]	return the information to the network-enabled device; and	Krouse discloses return the information to the network-enabled device. *Id.* at [1Pre-1F]. Krouse at 8:13–25 • "Generator **38** utilizes the numeric transaction data, payee and payment amount data, transaction date and time data, and terminal identification number to generate, for each point of sale financial transaction processed by the terminal, a unique respective record of the transaction for being assented to by the party tendering the negotiable instrument document being processed by the terminal. For each such transaction, the generator **38** causes the unique respective record generated by the generator **38** to be printed out by printer **36** (which comprises, e.g., a conventional laser printer) as a respective hard copy transaction record. An example of one such hard copy transaction record **70** is shown in FIG. 4."

Element	’529 Claim Recitation	Exemplary Citations in Reference
		Krouse at 7:7–31 • "Once the aforesaid data (i.e., payment amount being tendered and payee) has been input via the device **30**, document **110** has been placed in the scanner **34**, and presence of the document **110** in the scanner **34** has been signalled to the terminal operator, the terminal operator continues processing the document **110** by commanding the processor **40** via the device **30** to initiate scanning of the document **110**. In response, the processor **40** issues commands to the scanner **34** to scan the document **110** placed therein, whereupon, the scanner **34** scans at least one entire side **112** of the draft **110** so as to generate a computer-readable scanned image of at least that entire side **112** of the draft document **110** (i.e., computer-readable data from which an appropriately provisioned computer may generate a display image of the at least one scanned side **112** of the document **110**). The computer-readable scanned image data is then transmitted by the scanner **34** to the processor **40**; the processor **40** then transmits to the optical character recognizer **32** a copy of the scanned image data generated by the scanner **34** in order to permit the recognizer **32** to process same in the manner described below. A separate copy of the scanned image data is also transmitted to the storage system **42** where it is stored in association with the previously stored payment amount and payee data of the negotiable instrument being processed." Krouse at 8:44–64 • "Each transaction record **70** also includes a portion **78** which specifies the date and time of processing of the transaction, a portion **76** specifying the terminal identification number of the terminal processing the transaction, and a transaction data section **82** specifying the numerical transaction data generated from the scanned image. Each transaction record **70** also contains a unique transaction identification number generated by the generator **38** by concatenating a 5-digit form of the terminal identification number to the time and date of the transaction in numerically expressed in

Element	'529 Claim Recitation	Exemplary Citations in Reference
		the form YYYYMMDDHHMM (wherein YYYY represents the year, the left most MM represents the month, DD represents the day of the month, HH represents the hour in said day, and the right most MM represents the minutes in said hour of the transaction) and to a system-wide, unique 5-digit daily transaction sequence number (which is reset system-wide, periodically, e.g., once per day), and a tendering party assent portion **84** for receiving a signature **72** from the tendering party or other authorized representative thereof for signifying assent by the tendering party to the terms and conditions specified in the record **70**." To the extent it is found that Krouse does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Krouse does not anticipate this claim, Krouse renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Krouse with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1H]	the network-enabled device further programmed to present the information related to the object to a user.	Krouse discloses the network-enabled device further programmed to present the information related to the object to a user. *Id.* at [1Pre-G]. Krouse at 3:44–57 • "One preferred embodiment of the system of this aspect of the present invention includes an optical scanner for generating a computer-readable scanned image of the document, and an optical character recognizer for generating, from the scanned image, transaction data indicative of at least

Element	'529 Claim Recitation	Exemplary Citations in Reference
		a portion of the visual information. This embodiment of the system of this aspect of the present invention also includes an authorization record generator for generating, based upon the transaction data, a record of the transaction for being assented to by the party tendering the document prior to completion of the transaction, and for generating from the record once it is assented to by the party, a computer-readable record of assent of the party to the transaction." Krouse at 8:13–25 • "Generator **38** utilizes the numeric transaction data, payee and payment amount data, transaction date and time data, and terminal identification number to generate, for each point of sale financial transaction processed by the terminal, a unique respective record of the transaction for being assented to by the party tendering the negotiable instrument document being processed by the terminal. For each such transaction, the generator **38** causes the unique respective record generated by the generator 38 to be printed out by printer **36** (which comprises, e.g., a conventional laser printer) as a respective hard copy transaction record. An example of one such hard copy transaction record **70** is shown in FIG. 4." Krouse at 7:7–31 • "Once the aforesaid data (i.e., payment amount being tendered and payee) has been input via the device **30**, document **110** has been placed in the scanner **34**, and presence of the document **110** in the scanner **34** has been signalled to the terminal operator, the terminal operator continues processing the document **110** by commanding the processor **40** via the device **30** to initiate scanning of the document **110**. In response, the processor **40** issues commands to the scanner **34** to scan the document **110** placed therein, whereupon, the scanner **34** scans at least one entire side **112** of the draft **110** so as to generate a computer-readable scanned image of at least that entire side **112** of the draft document **110** (i.e., computer-readable data from which an appropriately provisioned

Element	'529 Claim Recitation	Exemplary Citations in Reference
		computer may generate a display image of the at least one scanned side **112** of the document **110**). The computer-readable scanned image data is then transmitted by the scanner **34** to the processor **40**; the processor **40** then transmits to the optical character recognizer **32** a copy of the scanned image data generated by the scanner **34** in order to permit the recognizer **32** to process same in the manner described below. A separate copy of the scanned image data is also transmitted to the storage system **42** where it is stored in association with the previously stored payment amount and payee data of the negotiable instrument being processed." Krouse at 8:44–64 • "Each transaction record **70** also includes a portion **78** which specifies the date and time of processing of the transaction, a portion **76** specifying the terminal identification number of the terminal processing the transaction, and a transaction data section **82** specifying the numerical transaction data generated from the scanned image. Each transaction record **70** also contains a unique transaction identification number generated by the generator **38** by concatenating a 5-digit form of the terminal identification number to the time and date of the transaction in numerically expressed in the form YYYYMMDDHHMM (wherein YYYY represents the year, the left most MM represents the month, DD represents the day of the month, HH represents the hour in said day, and the right most MM represents the minutes in said hour of the transaction) and to a system-wide, unique 5-digit daily transaction sequence number (which is reset system-wide, periodically, e.g., once per day), and a tendering party assent portion **84** for receiving a signature **72** from the tendering party or other authorized representative thereof for signifying assent by the tendering party to the terms and conditions specified in the record **70**." To the extent it is found that Krouse does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Krouse does not anticipate this claim, Krouse renders it obvious, either alone, in combination with

Element	'529 Claim Recitation	Exemplary Citations in Reference
		the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Krouse with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

2. Claim 4

Claim	'529 Claim Recitation	Exemplary Citations in Reference
4	The system of claim 1, wherein the service is further programmed to identify the object by applying multiple algorithms to at least a portion of the image or of the data, and identify the object as a function of a confidence levels associated with results of the multiple algorithms.	Krouse discloses the service is further programmed to identify the object by applying multiple algorithms to at least a portion of the image or of the data, and identify the object as a function of a confidence levels associated with results of the multiple algorithms. *Id.* at [1]. Krouse at Abstract • "In accordance with the present invention, financial transaction processing systems and methods are provided. One preferred embodiment of a method according to one aspect of the present invention includes generating an optically scanned image of at least a portion of document containing visual data, in a particular format, representing information related to the financial transaction. Recognition characteristics are generated from the scanned image and are compared to respective sets of reference recognition characteristics generated from respective other transaction documents having different respective formats to determine therefrom whether the particular format of the visual data matches one of the respective formats of the

		other documents. When such a match is found to exist, location is determined of a field in the scanned image to which optical character recognition may be applied to generate therefrom the information, based upon the respective format found to match the particular format of the visual data. Optical character recognition is then utilized to generate said visual data from said location." Krouse at 13:12 – 14:6 • " When the scanned image is first received by the generator 204, it is processed by an image de-skewing processor 210, such as Scanfix Image Processor available from TMS Sequoia, of Burlingame, Calif., which adjusts the scanned image data by rotating and/or translating the scanned image so as to cause two adjacent edges (e.g., the top and left side edges) of the scanned image (i.e., the edges of the scanned image generated from the top and left side edges of the document 400) to be aligned with the ordinate and abscissa of an orthogonal coordinate system whose center is chosen so as to coincide with the dark pixel of the scanned image closest to those two adjacent edges. Once the processor 210 has accomplished this processing of the scanned image data, the scanned image data is next processed by the noise eliminator 212, which processes the image data to remove dark pixels (i.e., non-whitespace pixels) therefrom that were erroneously generated as a result of noise. The noise eliminator 212 accomplishes this by locating isolated groups of contiguous dark pixels in the scanned image data having respective sizes (i.e., respective numbers of pixels) smaller than three dark pixels; these isolated groups of contiguous dark pixels are then removed from the scanned image data and replaced with whitespace by the processor 212. • After being processed by the noise eliminator 212, the scanned image data is processed by the image enhancer processor 214, which processor 214 enhances the dark pixels remaining in the scanned image after processing of the scanned image data by the noise eliminator processor 212. Processor 214 accomplishes this by locating the dark

		pixels remaining in the scanned image after processing of the scanned image by the processor 212, and by extending by predetermined numbers of pixels (e.g., 5 pixels in each of the positive and negative ordinate directions, and 25 pixels in each of the positive and negative abscissa directions) each of these remaining dark pixels in both of the two mutually orthogonal directions of the coordinate system to which the scanned image was aligned by the de-skewing processor 210. That is, predetermined numbers of whitespace pixels surrounding and extending from the remaining dark pixels in the two mutually orthogonal directions of the coordinate system are replaced with dark pixels. Any resulting extension of the dark pixels that would extend beyond the edge of the scanned image is prevented from doing so by truncating the extension by a predetermined number of pixels sufficient to prevent the dark pixel extensions from coming within 5 pixels of the edges of the scanned image. • After being processed by the image enhancer 214, the result, enhanced scanned image is processed by dark pixel group locator processor 222. Locator processor 222 then generates respective imaginary rectangles around the respective groups of contiguous dark pixels such that the respective rectangles each include two sides are parallel to the ordinate and two sides are parallel to the abscissa of the coordinate system to which the scanned image was aligned by the processor 210, and such that the respective rectangles include all of the respective dark pixels of each of the respective groups of contiguous dark pixels. The locator processor 222 then selects those imaginary rectangles, closest to the origin of the coordinate system, whose respective longitudinal dimensions exceed a predetermined threshold (e.g., 1 inch), and determines the respective coordinates of at least one predetermined vertex (e.g., the respective vertex having the minimal abscissal coordinate value and the maximal ordinate coordinate value for the respective rectangle) and dimensions of a predetermined number (e.g., no more than 25) of these selected imaginary rectangles. The respective vertex coordinates and dimensions of these imaginary rectangles are

		output by the locator processor 222 as recognition characteristics for the scanned document 400, which recognition characteristics are transmitted from the generator 204 to the recognition characteristic comparison processor 220." To the extent it is found that Krouse does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Krouse does not anticipate this claim, Krouse renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Krouse with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

3. Claim 20

Claim	'529 Claim Recitation	Exemplary Citations in Reference
20	The system of claim 1, wherein the data processing operation is selected from the group consisting of contrast enhancement, noise removal, and de-blurring.	Krouse discloses the system of claim 1, wherein the data processing operation is selected from the group consisting of contrast enhancement, noise removal, and de-blurring. *Id.* at [1]. Krouse at 13:12–14:16 • "When the scanned image is first received by the generator **204**, it is processed by an image de-skewing processor **210**, such as Scanfix Image Processor available from TMS Sequoia, of Burlingame, Calif., which adjusts the scanned image data by rotating and/or translating the scanned image so as to cause two adjacent edges (e.g., the top and left side edges)

		of the scanned image (i.e., the edges of the scanned image generated from the top and left side edges of the document **400**) to be aligned with the ordinate and abscissa of an orthogonal coordinate system whose center is chosen so as to coincide with the dark pixel of the scanned image closest to those two adjacent edges. Once the processor **210** has accomplished this processing of the scanned image data, the scanned image data is next processed by the noise eliminator **212**, which processes the image data to remove dark pixels (i.e., non-whitespace pixels) therefrom that were erroneously generated as a result of noise. The noise eliminator **212** accomplishes this by locating isolated groups of contiguous dark pixels in the scanned image data having respective sizes (i.e., respective numbers of pixels) smaller than three dark pixels; these isolated groups of contiguous dark pixels are then removed from the scanned image data and replaced with whitespace by the processor **212**. After being processed by the noise eliminator **212**, the scanned image data is processed by the image enhancer processor **214**, which processor **214** enhances the dark pixels remaining in the scanned image after processing of the scanned image data by the noise eliminator processor **212**. Processor **214** accomplishes this by locating the dark pixels remaining in the scanned image after processing of the scanned image by the processor **212**, and by extending by predetermined numbers of pixels (e.g., 5 pixels in each of the positive and negative ordinate directions, and 25 pixels in each of the positive and negative abscissa directions) each of these remaining dark pixels in both of the two mutually orthogonal directions of the coordinate system to which the scanned image was aligned by the de-skewing processor 210. That is, predetermined numbers of whitespace pixels surrounding and extending from the remaining dark pixels in the two mutually orthogonal directions of the coordinate system are replaced with dark pixels. Any resulting extension of the dark pixels that would extend beyond the edge of the scanned image is prevented from doing so by truncating the extension by a predetermined number of pixels sufficient to

		prevent the dark pixel extensions from coming within 5 pixels of the edges of the scanned image. After being processed by the image enhancer 214, the result, enhanced scanned image is processed by dark pixel group locator processor **222**. Locator processor **222** then generates respective imaginary rectangles around the respective groups of contiguous dark pixels such that the respective rectangles each include two sides are parallel to the ordinate and two sides are parallel to the abscissa of the coordinate system to which the scanned image was aligned by the processor **210**, and such that the respective rectangles include all of the respective dark pixels of each of the respective groups of contiguous dark pixels. The locator processor **222** then selects those imaginary rectangles, closest to the origin of the coordinate system, whose respective longitudinal dimensions exceed a predetermined threshold (e.g., 1 inch), and determines the respective coordinates of at least one predetermined vertex (e.g., the respective vertex having the minimal abscissal coordinate value and the maximal ordinate coordinate value for the respective rectangle) and dimensions of a predetermined number (e.g., no more than 25) of these selected imaginary rectangles. The respective vertex coordinates and dimensions of these imaginary rectangles are output by the locator processor **222** as recognition characteristics for the scanned document **400**, which recognition characteristics are transmitted from the generator **204** to the recognition characteristic comparison processor **220**." To the extent it is found that Krouse does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Krouse does not anticipate this claim, Krouse renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Krouse with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the

		reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

Nantworks, LLC v. Bank of America Corporation
Central District of California, No. 2:20-cv-7872
Invalidity Chart for U.S. Patent 7,881,529 based on the Query by Image and Video Content ("QBIC") System

Exhibit G-23

Bank of America Corporation ("BofA" or "Defendant") provides this chart subject to all reservations, objections, statements, and disclaimers set forth herein and in BofA's Final Invalidity Contentions Cover Pleading, as well as any amendment, supplement, or modification thereof, which are incorporated herein by reference in their entirety. A statement of reasons and good cause for each supplement in this chart is identified in the cover pleading.

QBIC was known or used by others in the United States and publicly used or on sale in the United States by November 5, 2000. *See generally* W. Niblack et al., The QBIC Project: Querying Images By Content Using Color, Texture, and Shape, SPIE Vol. 1908 (1993), pp. 173–187 (IBM0002390–404); M. Flickner et al., Query by Image and Video Content: The QBIC System ("QBIC"), IEEE Sept. 1995, pp. 23–32 (IBM0000893–902); W. Niblack et al, Updates to the QBIC System, SPIE Vol. 3312 (1997), pp. 150–161 (IBM 0002419–430); and Query By Image Content QBIC Demonstration Program, IBM Research, Sept. 1998 (IBM000747). QBIC is available as prior art at least under 35 U.S.C. § 102(a), (b), and (g), having been known or used by others in the United States and publicly used or on sale in the United States by at least 1995 and certainly by November 5, 2000, including, for example, at http://libra.uc.davis.edu through the Art History Department at U.C. Davis; at www..thinker.org/imagebase/index-2 .htm3 through the Fine Arts Museum of San Francisco; through demos at http://wwwqbic.almaden.ibm.com; and through IBM's commercial products, including in IBM's Ultimedia Manager, the DB2 series of products such as the Image Extender in DB2 Universal Database, and the Enterprise Information Portal. *See generally* Deposition of IBM Corporate Representative Myron Flickner, *Nantworks, LLC and Nant Holdings IP, LLC v. Bank of America Corporation and Bank of America, N.A.*, 2:20-cv-07872-GW-PVC (C.D.C.A. 2020) ("Flickner 8/29/23 Depo") at e.g., 32–33, 39–40, 44, 62–63, 66–67,72–74, 103–107, 110–112, 127, 163–174, 183–189, 191–195, 215, 219–220, 241, 307–309, 311–317, Exs. 3–13; *see also* Deposition of IBM Corporate Representative Myron Flickner, *Network-1 Technologies, Inc., v. Google LLC and Youtube LLC*, 1:14-cv-09558-PGG (S.D.N.Y. 2019) ("Flickner Depo"), including, e.g., at 13, 14, 15–18, 29–30, 34, 36, 37, 43–46, 48, 58, 60, 61–62, 67–69, 72, 74–76, 84–86, 88–89, 98–100, 102, 106, 113, 125, 135, 152–154, 160–161, 170–172, 180–181, 194, 198–201, 203–205, 215–216, 220–221, and Exs. 1–16.

As shown in the chart below, QBIC anticipates asserted claims 4 and 20 of the '529 Patent. The disclosures/citations for each element incorporate the disclosures citations of the proceeding limitations. To the extent QBIC is found not to expressly disclose certain limitations in the asserted claims, such limitations are inherent. To the extent QBIC is found not to anticipate any asserted claims or claim elements of the '529 Patent, the reference nevertheless renders those claims or claim elements obvious under 35 U.S.C. § 103, either alone or in combination with other art identified in the cited references and/or in the cover pleading of BofA's Final Invalidity Contentions. These Invalidity Contentions are not an admission by BofA that the accused products, including any current or past versions of the accused products, are covered by, or infringe the asserted claims, but are based instead on the recognition that if the claims are interpreted to be broad enough to encompass the accused products, the claims must also be construed to have that same scope when considering whether they are invalid. The following chart incorporates by reference the analysis from Exhibit G-21 for each limitation.

The following chart is based on the Court's claim construction rulings (Dkt. Nos. 145, 153, 236), and to the extent not contradictory, Plaintiffs' apparent interpretations of the claims in Plaintiffs' "Preliminary" Final Infringement Contentions. BofA notes that in many instances, Plaintiffs' "Preliminary" Final Infringement Contentions ignore and/or contradict the Court's claim construction rulings. BofA does not accept the assumptions concerning the scope and meaning implicit in Plaintiffs' Preliminary Final Infringement Contentions, to the extent those assumptions are discernible, and reserves the right to challenge Plaintiffs' proposed (or implied) constructions. Furthermore, Plaintiffs' "Preliminary" Final Contentions fail to comply with the level of specificity required by SPR 2.1.2 and 2.1.3, and include improper reservation of rights to supplement with "additional ways in which BoA Accused Products infringe." Any alleged deficient discovery is of Plaintiffs own making. Most if not all of the deficiencies in Plaintiffs' Preliminary Final Contentions have been documented by Defendants through multiple letters since at least May 18, 2021 and up to Bank of America's most recent letter of August 19, 2023. Plaintiffs' shifting infringement contentions, deficiencies in its infringement contentions, and changing claim scope have prejudiced BofA from being able to reasonably prepare its defenses, including its invalidity contentions. BofA also reserves the right to revise and supplement these charts if and when Plaintiffs are permitted to provide revised or supplemental Infringement Contentions. Citations given in the chart below are merely representative of the respective elements to provide notice to Plaintiffs of BofA's invalidity theories and are not meant to be exhaustive.

1. **Claim 1 (Non-Asserted)**

Element	'529 Claim Recitation	Exemplary Citations in Reference
[1Pre]	A system comprising:	QBIC discloses a system. The QBIC Project: Querying Images by Content Using Color, Texture, and Shape (IBM 0002390–404) at 2390: 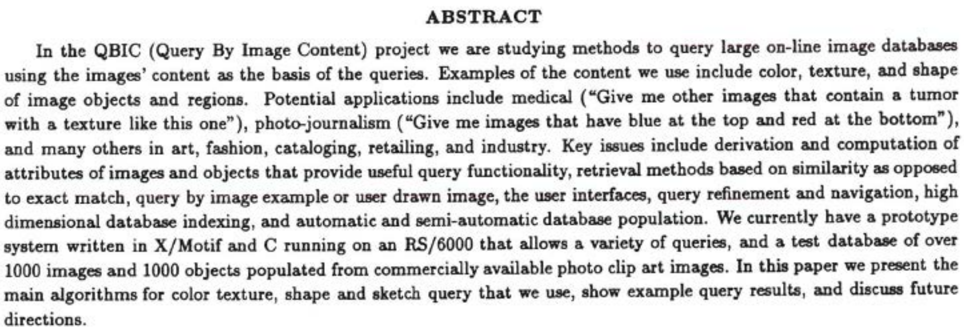 ABSTRACT In the QBIC (Query By Image Content) project we are studying methods to query large on-line image databases using the images' content as the basis of the queries. Examples of the content we use include color, texture, and shape of image objects and regions. Potential applications include medical ("Give me other images that contain a tumor with a texture like this one"), photo-journalism ("Give me images that have blue at the top and red at the bottom"), and many others in art, fashion, cataloging, retailing, and industry. Key issues include derivation and computation of attributes of images and objects that provide useful query functionality, retrieval methods based on similarity as opposed to exact match, query by image example or user drawn image, the user interfaces, query refinement and navigation, high dimensional database indexing, and automatic and semi-automatic database population. We currently have a prototype system written in X/Motif and C running on an RS/6000 that allows a variety of queries, and a test database of over 1000 images and 1000 objects populated from commercially available photo clip art images. In this paper we present the main algorithms for color texture, shape and sketch query that we use, show example query results, and discuss future directions. QBIC Demo (IBM000747):  QBIC(TM) -- IBM's Query By Image Content Goto: How to Use QBIC — QBIC allows you to make queries of image databases large and small based on visual characteristics -- properties such as color percentages, color layout and textures occurring in the images. The queries match image colors, textures and their positions without having to describe them with in words. QBIC's content-based queries can also be combined with keywords to provide a powerful multimedia search and retrieval tool. As image collections grow larger and become more prolific, QBIC is the best tool to more efficiently search and navigate through them. IBM Image Search Demo Stop QBIC Http Server — In this demo, we use QBIC to index about 470 color images based on their color and texture information. You can evaluate the QBIC Search capability by following the link located at left called IBM Image Search Demo. For license and business related questins, please contact Ted Loewenberg by phone 408-927-1202, or by email Flickner Depo at, e.g., 24, 41-42, 62-63 (user could take a picture or video and submit it

Element	'529 Claim Recitation	Exemplary Citations in Reference
		to the QBIC system). 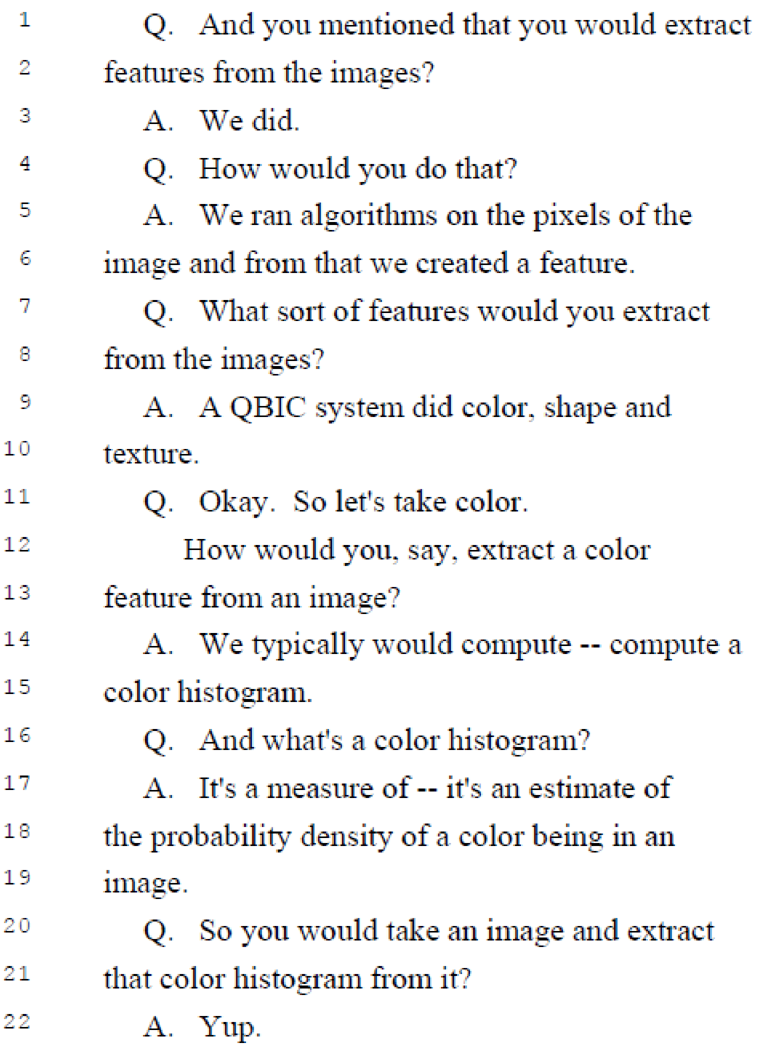 Q. And you mentioned that you would extract features from the images? A. We did. Q. How would you do that? A. We ran algorithms on the pixels of the image and from that we created a feature. Q. What sort of features would you extract from the images? A. A QBIC system did color, shape and texture. Q. Okay. So let's take color. How would you, say, extract a color feature from an image? A. We typically would compute -- compute a color histogram. Q. And what's a color histogram? A. It's a measure of -- it's an estimate of the probability density of a color being in an image. Q. So you would take an image and extract that color histogram from it? A. Yup.

		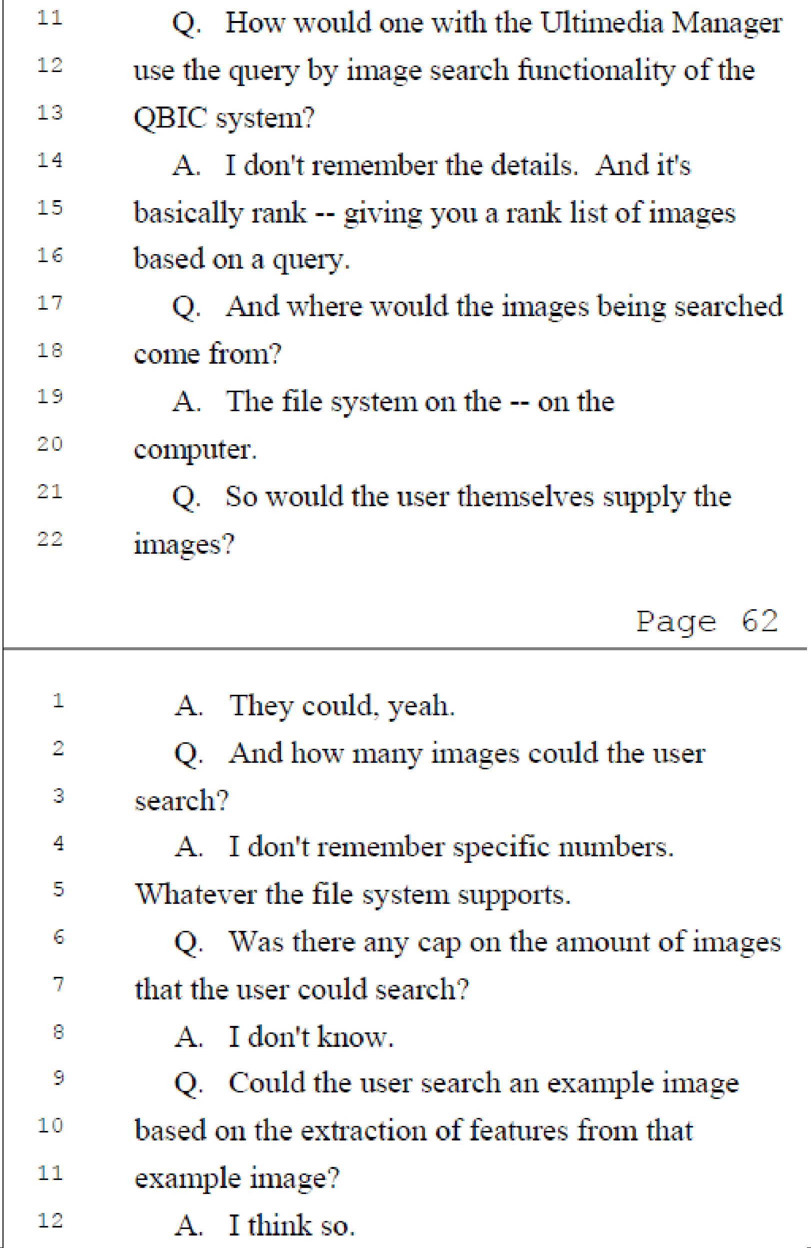 Q. How would one with the Ultimedia Manager use the query by image search functionality of the QBIC system? A. I don't remember the details. And it's basically rank -- giving you a rank list of images based on a query. Q. And where would the images being searched come from? A. The file system on the -- on the computer. Q. So would the user themselves supply the images? Page 62 A. They could, yeah. Q. And how many images could the user search? A. I don't remember specific numbers. Whatever the file system supports. Q. Was there any cap on the amount of images that the user could search? A. I don't know. Q. Could the user search an example image based on the extraction of features from that example image? A. I think so.

Element	'529 Claim Recitation	Exemplary Citations in Reference
		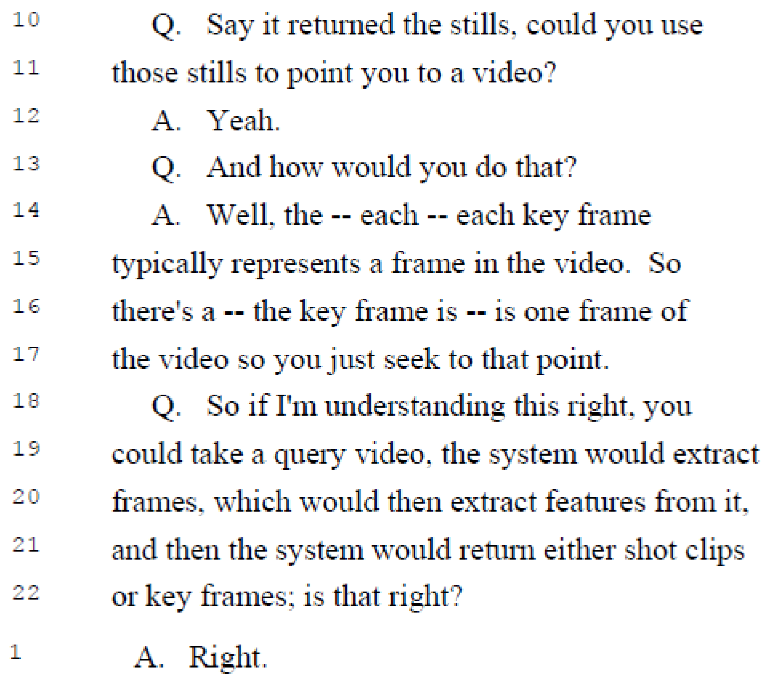 Q. Say it returned the stills, could you use those stills to point you to a video? A. Yeah. Q. And how would you do that? A. Well, the -- each -- each key frame typically represents a frame in the video. So there's a -- the key frame is -- is one frame of the video so you just seek to that point. Q. So if I'm understanding this right, you could take a query video, the system would extract frames, which would then extract features from it, and then the system would return either shot clips or key frames; is that right? A. Right. Flickner 8/29/23 Depo at 42:

Element	'529 Claim Recitation	Exemplary Citations in Reference
		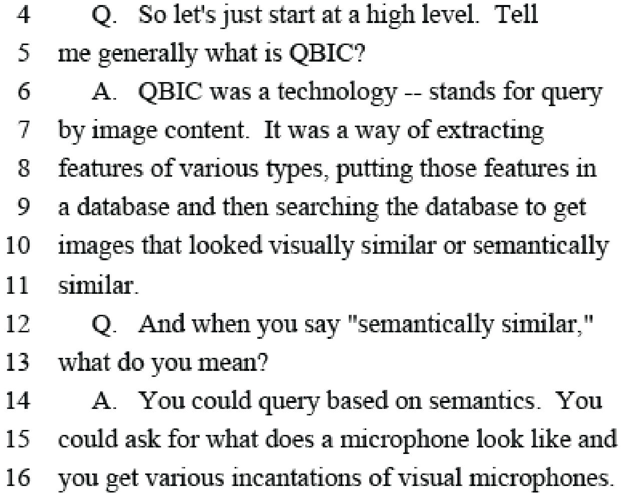 Q. So let's just start at a high level. Tell me generally what is QBIC? A. QBIC was a technology -- stands for query by image content. It was a way of extracting features of various types, putting those features in a database and then searching the database to get images that looked visually similar or semantically similar. Q. And when you say "semantically similar," what do you mean? A. You could query based on semantics. You could ask for what does a microphone look like and you get various incantations of visual microphones. *Id.* at 205–206:

Element	'529 Claim Recitation	Exemplary Citations in Reference
		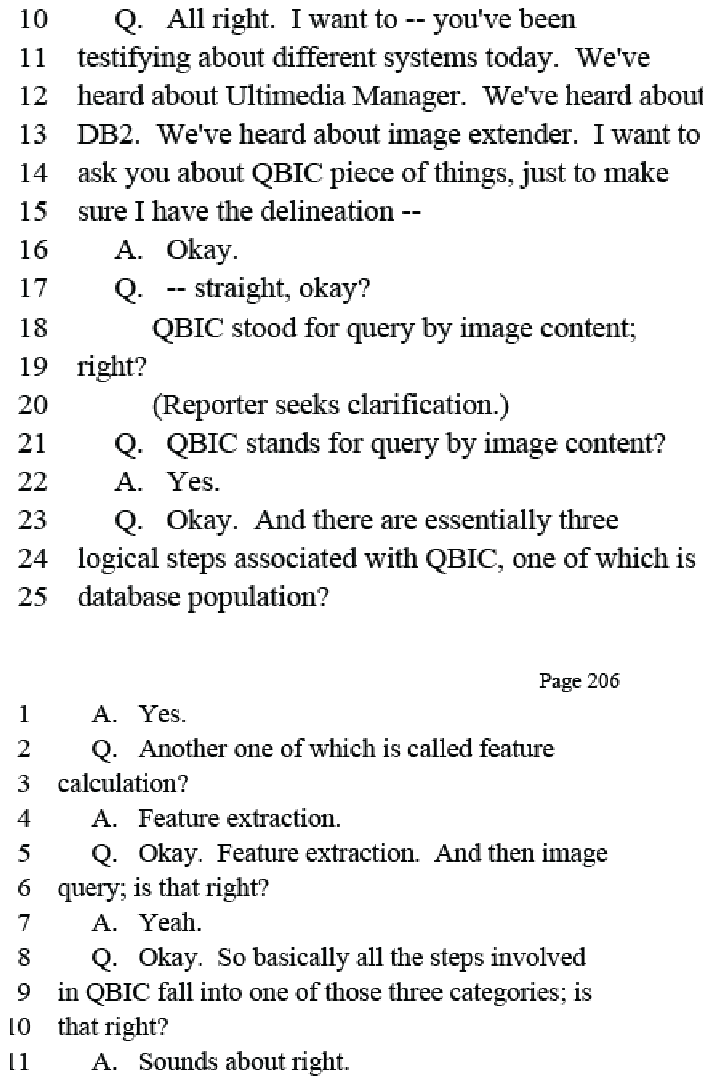 Q. All right. I want to -- you've been testifying about different systems today. We've heard about Ultimedia Manager. We've heard about DB2. We've heard about image extender. I want to ask you about QBIC piece of things, just to make sure I have the delineation -- A. Okay. Q. -- straight, okay? QBIC stood for query by image content; right? (Reporter seeks clarification.) Q. QBIC stands for query by image content? A. Yes. Q. Okay. And there are essentially three logical steps associated with QBIC, one of which is database population? Page 206 A. Yes. Q. Another one of which is called feature calculation? A. Feature extraction. Q. Okay. Feature extraction. And then image query; is that right? A. Yeah. Q. Okay. So basically all the steps involved in QBIC fall into one of those three categories; is that right? A. Sounds about right. Flickner 8/29/23 Depo at Ex. 5:

Element	'529 Claim Recitation	Exemplary Citations in Reference
		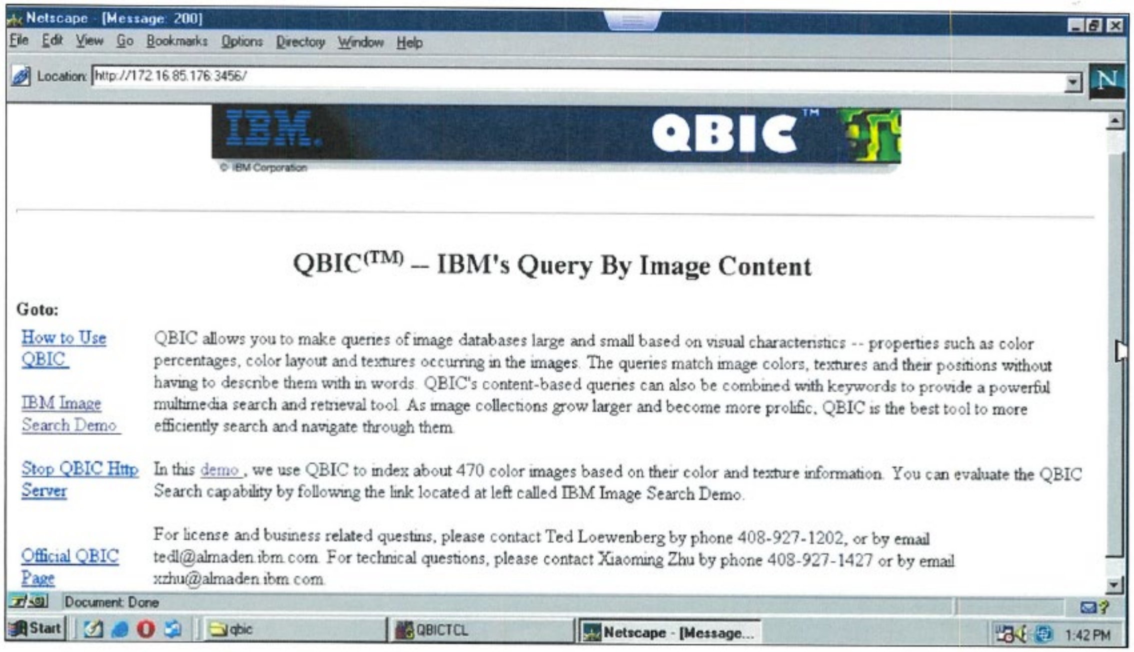 QBIC(TM) -- IBM's Query By Image Content Goto: How to Use QBIC IBM Image Search Demo Stop QBIC Http Server Official QBIC Page QBIC allows you to make queries of image databases large and small based on visual characteristics -- properties such as color percentages, color layout and textures occurring in the images. The queries match image colors, textures and their positions without having to describe them with in words. QBIC's content-based queries can also be combined with keywords to provide a powerful multimedia search and retrieval tool. As image collections grow larger and become more prolific, QBIC is the best tool to more efficiently search and navigate through them. In this demo, we use QBIC to index about 470 color images based on their color and texture information. You can evaluate the QBIC Search capability by following the link located at left called IBM Image Search Demo. For license and business related questins, please contact Ted Loewenberg by phone 408-927-1202, or by email tedl@almaden.ibm.com. For technical questions, please contact Xiaoming Zhu by phone 408-927-1427 or by email xzhu@almaden.ibm.com. The QBIC system was implemented in IBM commercial products, including in IBM's Ultimedia Manager, the DB2 series of products such as the Image Extender in DB2 Universal Database, and the Enterprise Information Portal. All of the features referenced in this chart were incorporated into these commercial implementations before November 6, 2000, as reflected in a combination of scientific papers, IBM's internal documentation, and the QBIC demo software applications. *See, e.g.*, Flickner 8/29/23 Depo Tr. at 33:9-24, 71:21-72:10, 164:21-166:2, 166:11-175:24, 178:5-186:17, 186:21-192:4, 192:14-194:1, 195:16-196:4, 194:1-308:1; IBM 000001-617; IBM 000747; IBM 000777-880; IBM 000893-902; IBM 0002315-320; IBM 0002390-404; IBM 0002418-430.
[1A]	a camera that captures an image;	QBIC uses still images and video (frames) for query and database population:

Element	'529 Claim Recitation	Exemplary Citations in Reference
		Query by Image and Video Content: The QBIC System (IBM 0000893–902): "We have developed the QBIC (Query by Image Content) system to explore content-based retrieval methods. QBIC allows queries on large image and video databases based on • example images, • user-constructed sketches and drawings, • selected color and texture patterns, • camera and object motion, and • other graphical information." *Id*. at pp. 1–3. "To achieve this functionality, QBIC has two main components: database population (the process of creating an image database) and database query. During the population, images and videos are processed to extract features describing their content-colors, textures, shapes, and camera and object motion-and the features are stored in a database. During the query, the user composes a query graphically. Features are generated from the graphical query and then input to a matching engine that finds images or videos from the database with similar features. Figure 2 shows the system architecture." *Id*. at p. 3. "For both population and query, the QBIC data model has • still images or scenes (full images) that contain objects (subsets of an image), and • video shots that consist of sets of contiguous frames and contain motion objects. For still images, the QBIC data model distinguishes between 'scenes' (or images) and 'objects.' A scene is an image or single representative frame of video. An object is a part of a scene-for example, the fox in Figure 3-or a moving entity in a video. For still image database population, features are extracted from images and objects and stored in a database as shown in the top left part of Figure 2. Videos are broken into clips called shots. Representative frames, or r-frames, are generated for each extracted shot. R-frames are treated as still images, and features are extracted and stored in the database. Further processing of shots generates motion objects-for example, a car moving across the screen. Queries are allowed on objects ('Find images with a red, round object'), scenes ('Find images that have approximately 30-percent red and IS-percent blue colors'), shots ('Find all shots panning from left to right'), or any combination ('Find images that have 30 percent red and contain a blue textured object')." *Id*. at p. 3.

Element	'529 Claim Recitation	Exemplary Citations in Reference
		"In still image database population, the images are reduced to a standard-sized icon called a thumbnail and annotated with any available text information. Object identification is an optional but key part of this step. It lets users manually, semiautomatically, or fully automatically identify interesting regions-which we call objects-in the images. . . . For video data, database population has three major components: • shot detection, • representative frame creation for each shot, and • derivation of a layered representation of coherently moving structures/ objects. . . . Once the shot boundaries have been detected, each shot is represented using an r-frame. R-frames are used for several purposes. First, during database population, r-frames are treated as still images in which objects can be identified by using the previously described methods." *Id*. at pp. 3–6. Image, Audio, and Video Extenders Administration and Programming (IBM 000001–617): "Your applications can even search for images by content. Imagine an application that uses visual examples to search for images. With such an application, users could select an example image and have the application find other images that have colors or textures similar to those in the example. With DB2 extenders' Query by Image Content (QBIC) capability, you can create applications that search for images in this visual way." *Id*. at p. 4. "To make an image available for searching by content, you catalog the image. When you catalog an image, the Image Extender analyzes the image, by computing the feature values for the image, and stores the values in a QBIC catalog. When you search for an image by content, your query identifies one or more features for the search (such as average color), a source for each feature (such as an example image), and a target set of cataloged images. The Image Extender computes the feature value of the source and compares it to the cataloged feature values for the target images. It then computes a score that indicates how similar the feature values of the target images

Element	'529 Claim Recitation	Exemplary Citations in Reference
		are to the source. You can have the Image Extender return the images whose features are most similar to the source. The Image Extender will return the handle of each image and the image score. You can also have the Image Extender return only the score of a single image." *Id*. at pp. 20–21. Flickner Depo at 88: 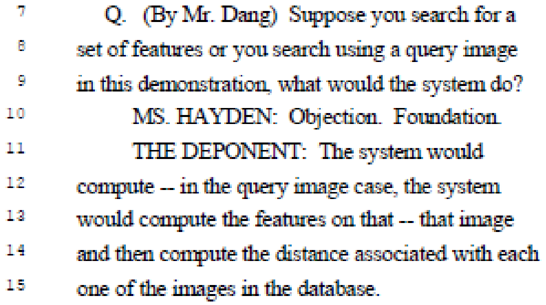 Q. (By Mr. Dang) Suppose you search for a set of features or you search using a query image in this demonstration, what would the system do? MS. HAYDEN: Objection. Foundation. THE DEPONENT: The system would compute -- in the query image case, the system would compute the features on that -- that image and then compute the distance associated with each one of the images in the database. Flickner Depo at 89-90:

Element	'529 Claim Recitation	Exemplary Citations in Reference
		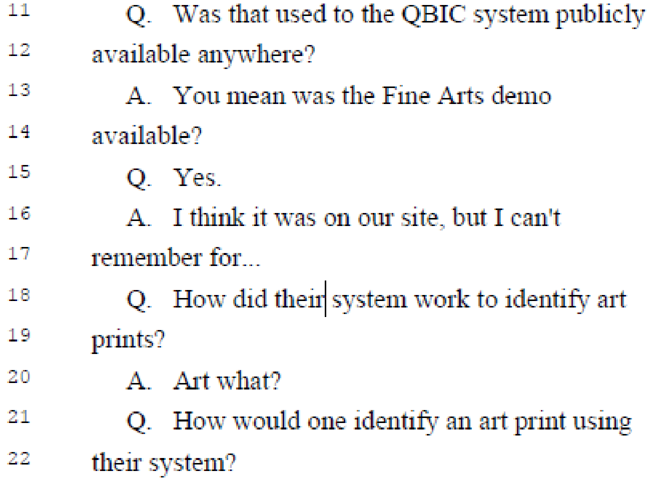Q. Was that used to the QBIC system publicly available anywhere? A. You mean was the Fine Arts demo available? Q. Yes. A. I think it was on our site, but I can't remember for... Q. How did their system work to identify art prints? A. Art what? Q. How would one identify an art print using their system? 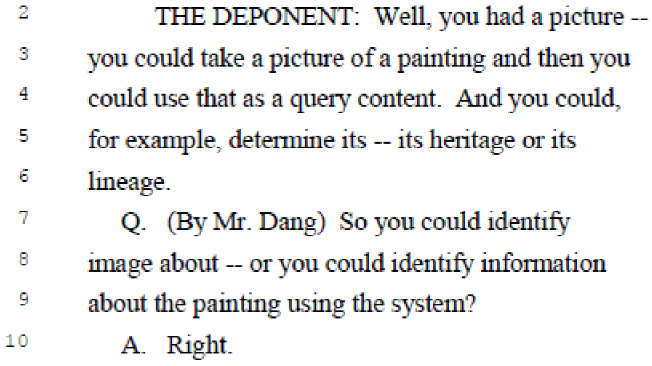THE DEPONENT: Well, you had a picture -- you could take a picture of a painting and then you could use that as a query content. And you could, for example, determine its -- its heritage or its lineage. Q. (By Mr. Dang) So you could identify image about -- or you could identify information about the painting using the system? A. Right. Flickner Depo at 39-40: 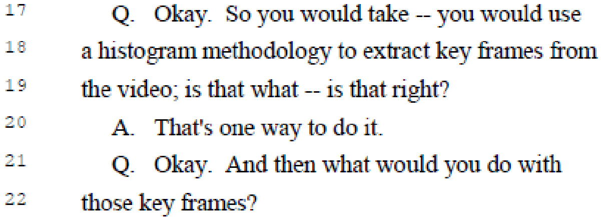Q. Okay. So you would take -- you would use a histogram methodology to extract key frames from the video; is that what -- is that right? A. That's one way to do it. Q. Okay. And then what would you do with those key frames?

Element	'529 Claim Recitation	Exemplary Citations in Reference
		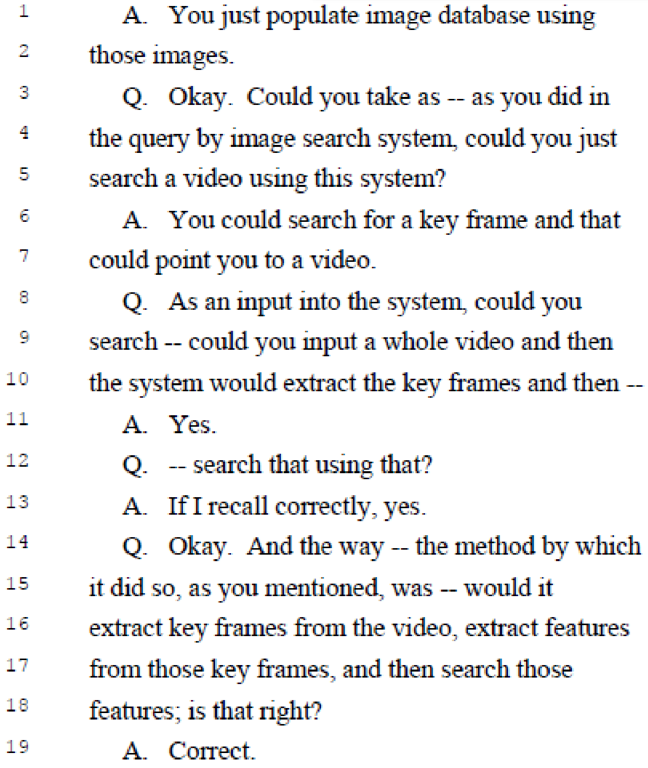 A. You just populate image database using those images. Q. Okay. Could you take as -- as you did in the query by image search system, could you just search a video using this system? A. You could search for a key frame and that could point you to a video. Q. As an input into the system, could you search -- could you input a whole video and then the system would extract the key frames and then -- A. Yes. Q. -- search that using that? A. If I recall correctly, yes. Q. Okay. And the way -- the method by which it did so, as you mentioned, was -- would it extract key frames from the video, extract features from those key frames, and then search those features; is that right? A. Correct. Flickner Depo at 54-55:

Element	'529 Claim Recitation	Exemplary Citations in Reference
		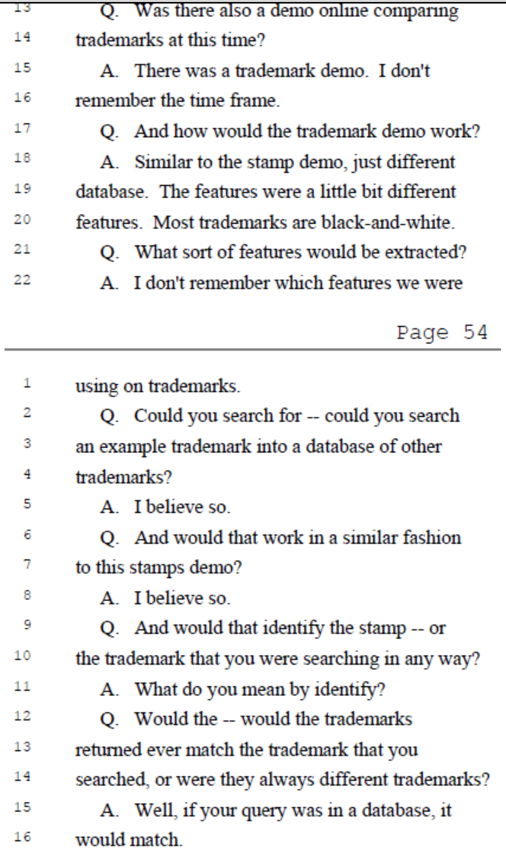 Q. Was there also a demo online comparing trademarks at this time? A. There was a trademark demo. I don't remember the time frame. Q. And how would the trademark demo work? A. Similar to the stamp demo, just different database. The features were a little bit different features. Most trademarks are black-and-white. Q. What sort of features would be extracted? A. I don't remember which features we were Page 54 using on trademarks. Q. Could you search for -- could you search an example trademark into a database of other trademarks? A. I believe so. Q. And would that work in a similar fashion to this stamps demo? A. I believe so. Q. And would that identify the stamp -- or the trademark that you were searching in any way? A. What do you mean by identify? Q. Would the -- would the trademarks returned ever match the trademark that you searched, or were they always different trademarks? A. Well, if your query was in a database, it would match. QBIC could integrate with a camera, such as a portable camera; a user could take an image using a portable camera and submit that image to the QBIC system for query. Flickner Depo at, e.g., 24, 41-42, 62-63 (user could take a picture or video and submit it to the QBIC system):

Element	’529 Claim Recitation	Exemplary Citations in Reference
		Q. And you mentioned that you would extract features from the images? A. We did. Q. How would you do that? A. We ran algorithms on the pixels of the image and from that we created a feature. Q. What sort of features would you extract from the images? A. A QBIC system did color, shape and texture. Q. Okay. So let's take color. How would you, say, extract a color feature from an image? A. We typically would compute -- compute a color histogram. Q. And what's a color histogram? A. It's a measure of -- it's an estimate of the probability density of a color being in an image. Q. So you would take an image and extract that color histogram from it? A. Yup.

		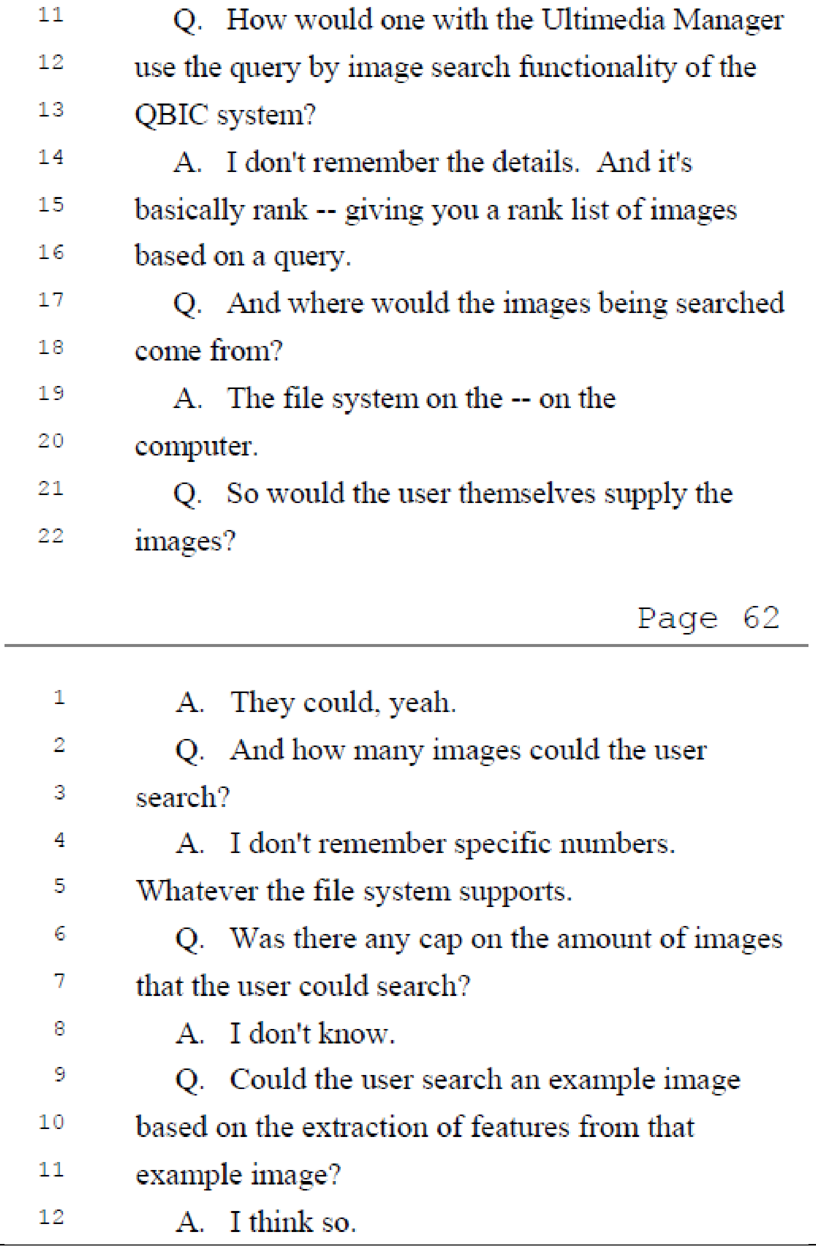 Q. How would one with the Ultimedia Manager use the query by image search functionality of the QBIC system? A. I don't remember the details. And it's basically rank -- giving you a rank list of images based on a query. Q. And where would the images being searched come from? A. The file system on the -- on the computer. Q. So would the user themselves supply the images? Page 62 A. They could, yeah. Q. And how many images could the user search? A. I don't remember specific numbers. Whatever the file system supports. Q. Was there any cap on the amount of images that the user could search? A. I don't know. Q. Could the user search an example image based on the extraction of features from that example image? A. I think so.

Element	'529 Claim Recitation	Exemplary Citations in Reference
		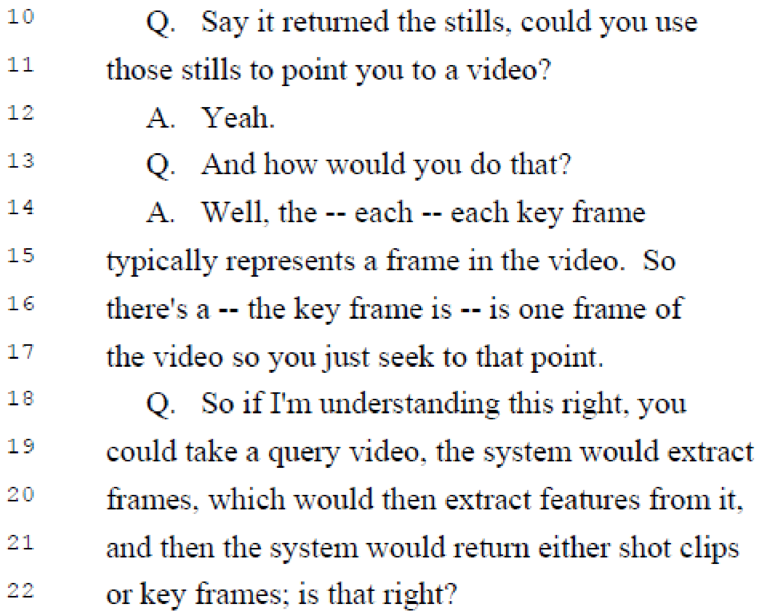 Q. Say it returned the stills, could you use those stills to point you to a video? A. Yeah. Q. And how would you do that? A. Well, the -- each -- each key frame typically represents a frame in the video. So there's a -- the key frame is -- is one frame of the video so you just seek to that point. Q. So if I'm understanding this right, you could take a query video, the system would extract frames, which would then extract features from it, and then the system would return either shot clips or key frames; is that right?  A. Right. Flickner 8/29/23 Depo at 34:

Element	'529 Claim Recitation	Exemplary Citations in Reference
		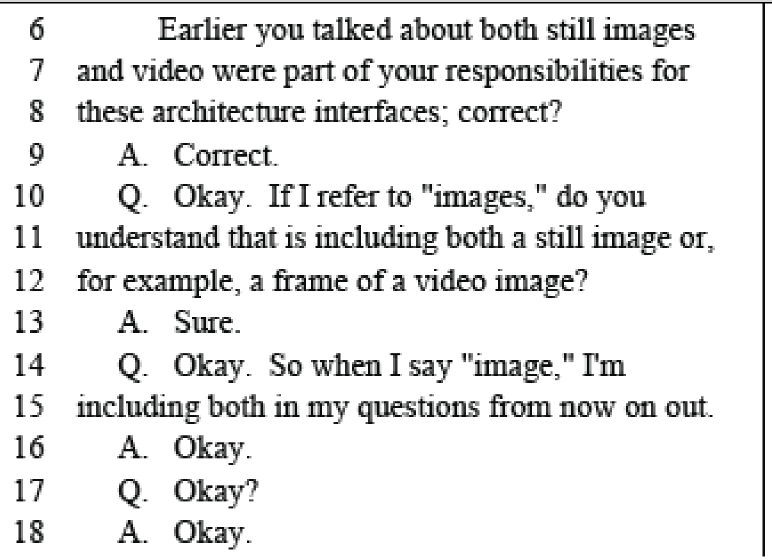 Earlier you talked about both still images and video were part of your responsibilities for these architecture interfaces; correct? A. Correct. Q. Okay. If I refer to "images," do you understand that is including both a still image or, for example, a frame of a video image? A. Sure. Q. Okay. So when I say "image," I'm including both in my questions from now on out. A. Okay. Q. Okay? A. Okay. *Id*. at 53: 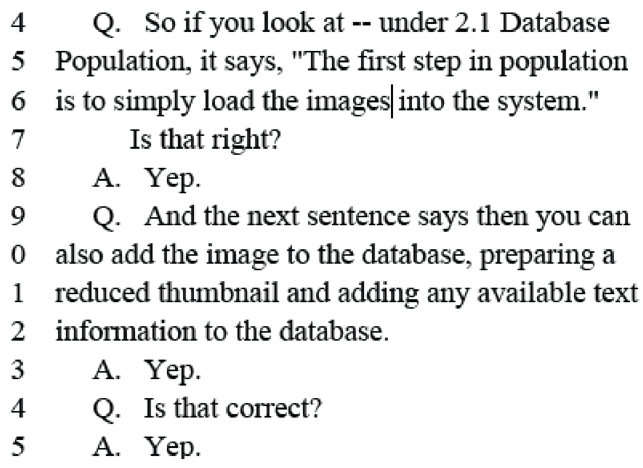 Q. So if you look at -- under 2.1 Database Population, it says, "The first step in population is to simply load the images into the system." Is that right? A. Yep. Q. And the next sentence says then you can also add the image to the database, preparing a reduced thumbnail and adding any available text information to the database. A. Yep. Q. Is that correct? A. Yep.

Element	'529 Claim Recitation	Exemplary Citations in Reference
		Id. at 60–61: 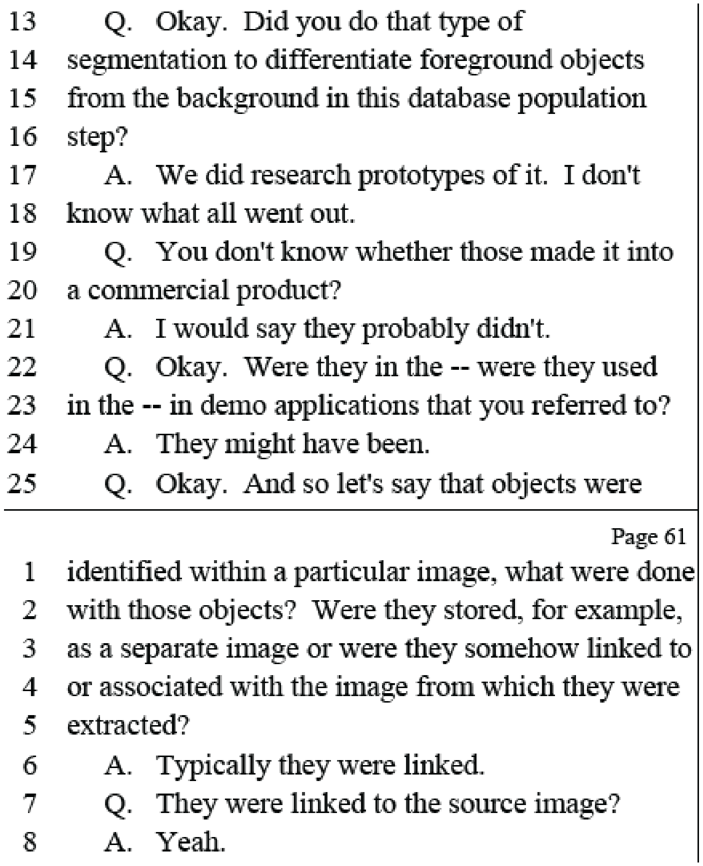 Q. Okay. Did you do that type of segmentation to differentiate foreground objects from the background in this database population step? A. We did research prototypes of it. I don't know what all went out. Q. You don't know whether those made it into a commercial product? A. I would say they probably didn't. Q. Okay. Were they in the -- were they used in the -- in demo applications that you referred to? A. They might have been. Q. Okay. And so let's say that objects were Page 61 identified within a particular image, what were done with those objects? Were they stored, for example, as a separate image or were they somehow linked to or associated with the image from which they were extracted? A. Typically they were linked. Q. They were linked to the source image? A. Yeah. *Id*. at 32–33:

Element	'529 Claim Recitation	Exemplary Citations in Reference
		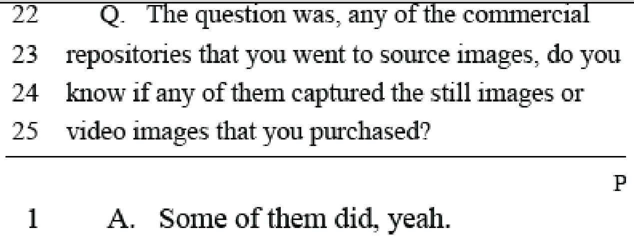 Q. The question was, any of the commercial repositories that you went to source images, do you know if any of them captured the still images or video images that you purchased? P A. Some of them did, yeah. *Id.* at 74–75:

			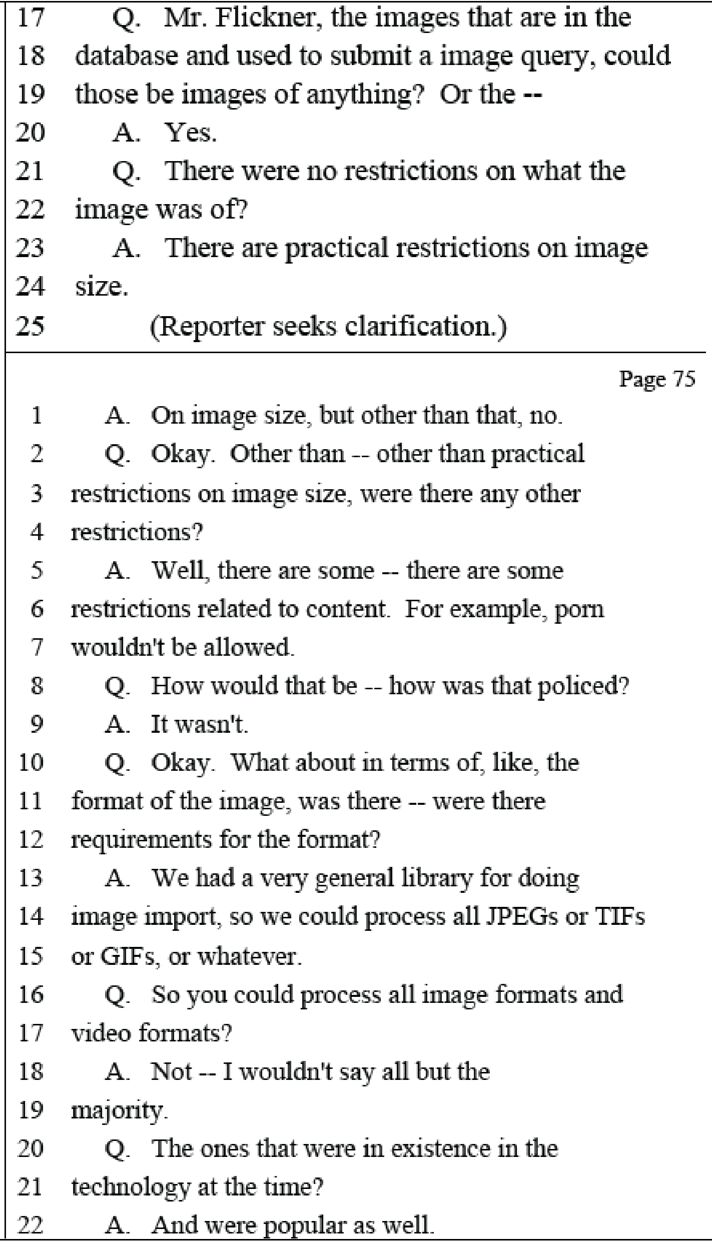 Q. Mr. Flickner, the images that are in the database and used to submit a image query, could those be images of anything? Or the -- A. Yes. Q. There were no restrictions on what the image was of? A. There are practical restrictions on image size. (Reporter seeks clarification.) Page 75 A. On image size, but other than that, no. Q. Okay. Other than -- other than practical restrictions on image size, were there any other restrictions? A. Well, there are some -- there are some restrictions related to content. For example, porn wouldn't be allowed. Q. How would that be -- how was that policed? A. It wasn't. Q. Okay. What about in terms of, like, the format of the image, was there -- were there requirements for the format? A. We had a very general library for doing image import, so we could process all JPEGs or TIFs or GIFs, or whatever. Q. So you could process all image formats and video formats? A. Not -- I wouldn't say all but the majority. Q. The ones that were in existence in the technology at the time? A. And were popular as well.

Element	'529 Claim Recitation	Exemplary Citations in Reference
		Id. at 114: 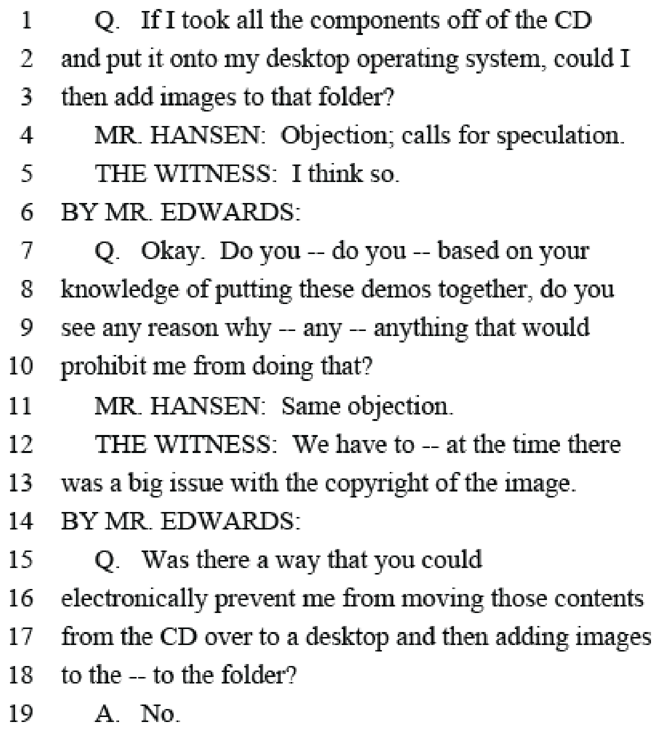 Q. If I took all the components off of the CD and put it onto my desktop operating system, could I then add images to that folder? MR. HANSEN: Objection; calls for speculation. THE WITNESS: I think so. BY MR. EDWARDS: Q. Okay. Do you -- do you -- based on your knowledge of putting these demos together, do you see any reason why -- any -- anything that would prohibit me from doing that? MR. HANSEN: Same objection. THE WITNESS: We have to -- at the time there was a big issue with the copyright of the image. BY MR. EDWARDS: Q. Was there a way that you could electronically prevent me from moving those contents from the CD over to a desktop and then adding images to the -- to the folder? A. No. *Id.* at 202:

Element	'529 Claim Recitation	Exemplary Citations in Reference
		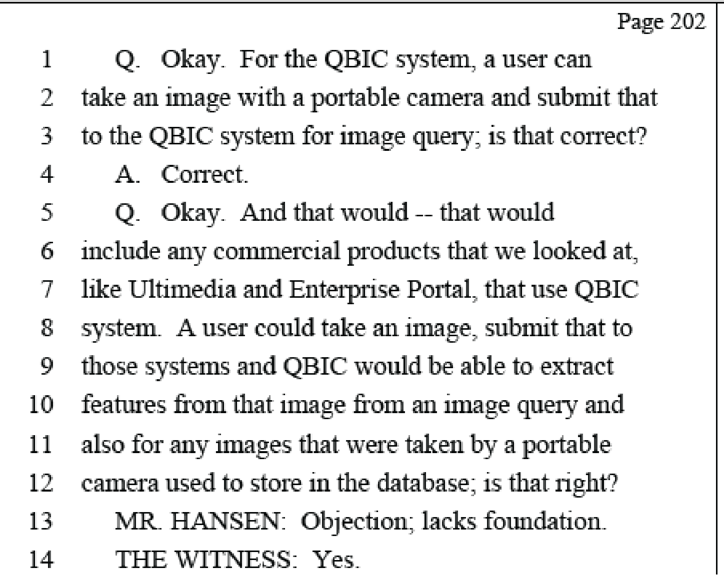 Page 202 Q. Okay. For the QBIC system, a user can take an image with a portable camera and submit that to the QBIC system for image query; is that correct? A. Correct. Q. Okay. And that would -- that would include any commercial products that we looked at, like Ultimedia and Enterprise Portal, that use QBIC system. A user could take an image, submit that to those systems and QBIC would be able to extract features from that image from an image query and also for any images that were taken by a portable camera used to store in the database; is that right? MR. HANSEN: Objection; lacks foundation. THE WITNESS: Yes. *Id.* at 300: 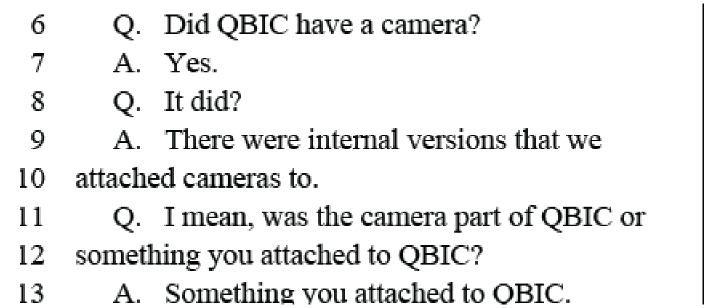 Q. Did QBIC have a camera? A. Yes. Q. It did? A. There were internal versions that we attached cameras to. Q. I mean, was the camera part of QBIC or something you attached to QBIC? A. Something you attached to QBIC. *Id.* at 312:

Element	'529 Claim Recitation	Exemplary Citations in Reference
		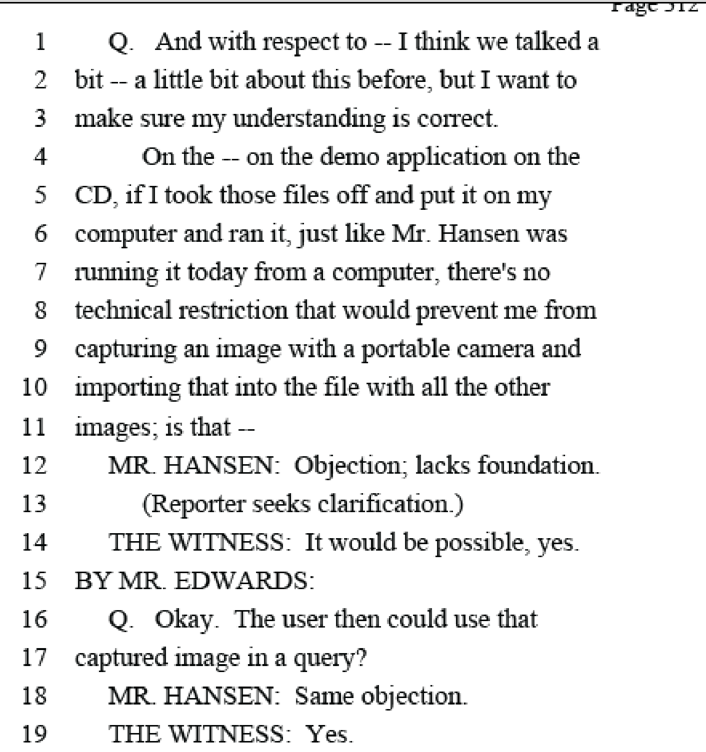Page 312 Q. And with respect to -- I think we talked a bit -- a little bit about this before, but I want to make sure my understanding is correct. On the -- on the demo application on the CD, if I took those files off and put it on my computer and ran it, just like Mr. Hansen was running it today from a computer, there's no technical restriction that would prevent me from capturing an image with a portable camera and importing that into the file with all the other images; is that -- MR. HANSEN: Objection; lacks foundation. (Reporter seeks clarification.) THE WITNESS: It would be possible, yes. BY MR. EDWARDS: Q. Okay. The user then could use that captured image in a query? MR. HANSEN: Same objection. THE WITNESS: Yes.

Element	'529 Claim Recitation	Exemplary Citations in Reference
		QBIC Demo (IBM000747) – showing compatibility with images captured by a mobile phone camera  Usage: I: Get Info C: Color Histogram L: Layout T: Texture S: Special Hybrid (Use a Java-capable Web browser to enable additional functionality.) Query was: Example: =AS_test/check1.jpg Query Type: Color Layout

Element	'529 Claim Recitation	Exemplary Citations in Reference
		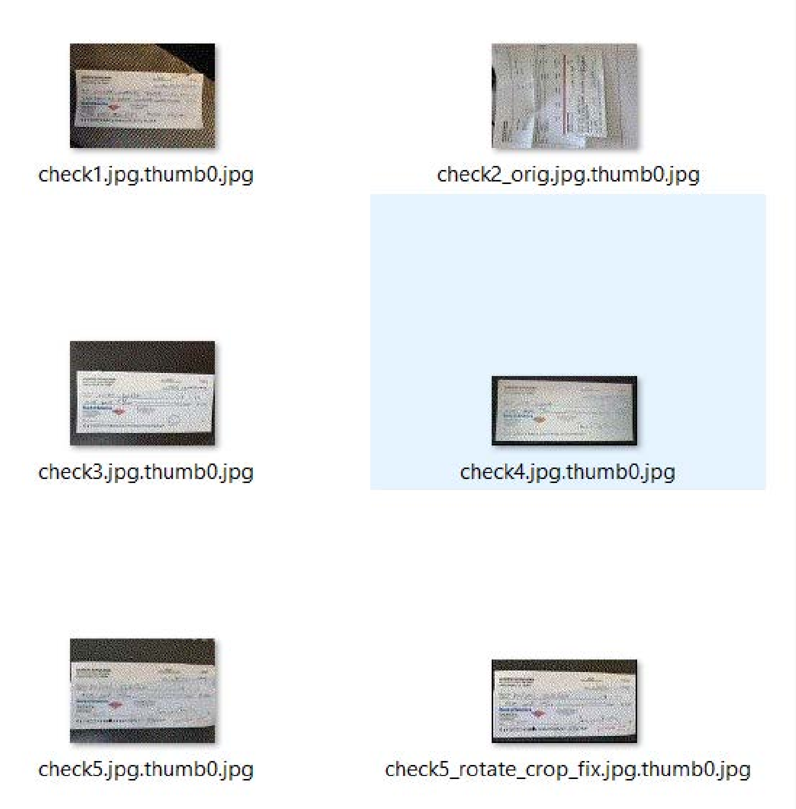

Element	'529 Claim Recitation	Exemplary Citations in Reference
		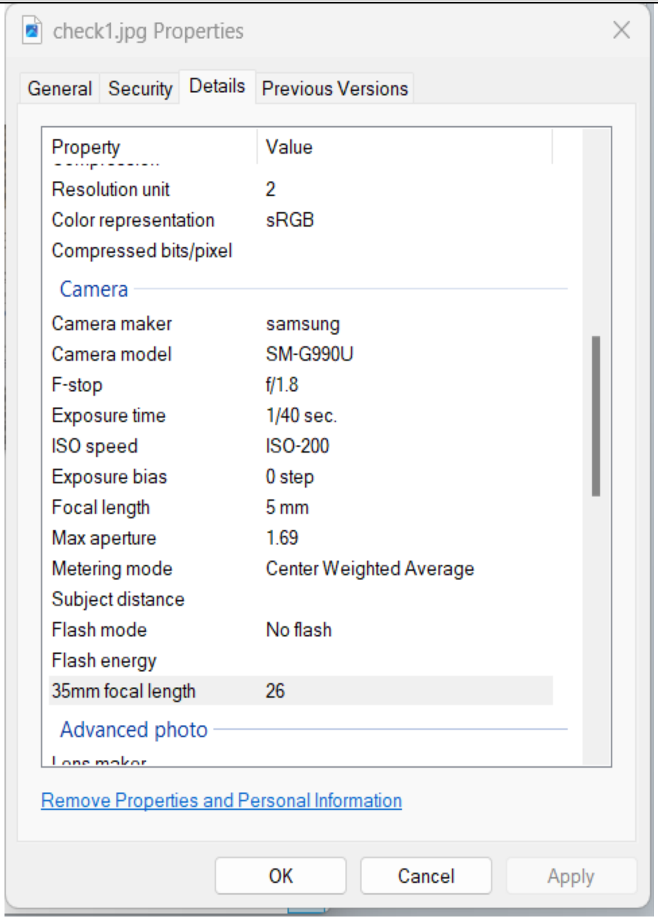 Property: Value Resolution unit: 2 Color representation: sRGB Compressed bits/pixel: Camera Camera maker: samsung Camera model: SM-G990U F-stop: f/1.8 Exposure time: 1/40 sec. ISO speed: ISO-200 Exposure bias: 0 step Focal length: 5 mm Max aperture: 1.69 Metering mode: Center Weighted Average Subject distance: Flash mode: No flash Flash energy: 35mm focal length: 26 Advanced photo

Element	'529 Claim Recitation	Exemplary Citations in Reference
		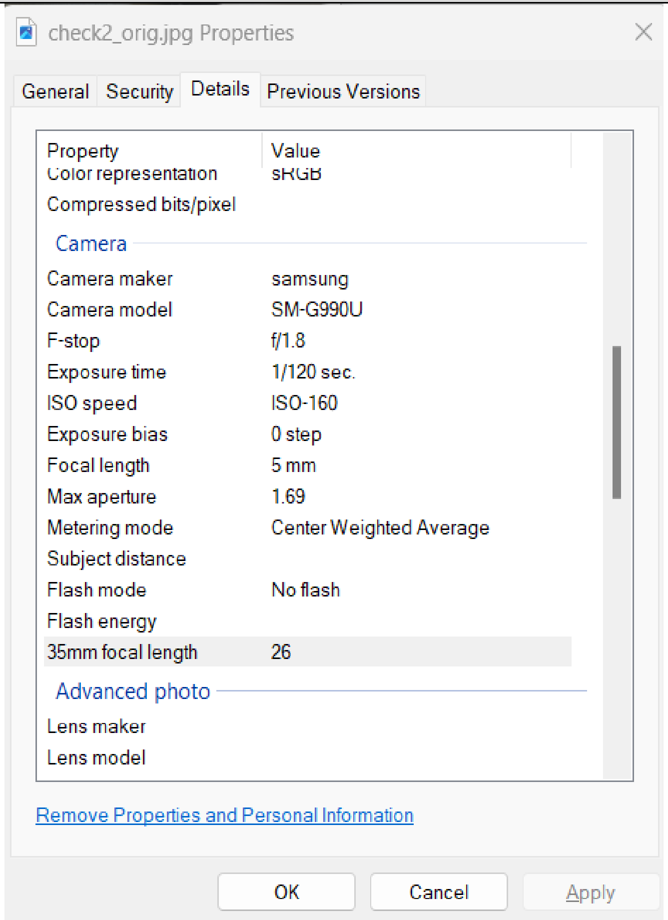

Property	Value
Color representation	sRGB
Compressed bits/pixel	
Camera	
Camera maker	samsung
Camera model	SM-G990U
F-stop	f/1.8
Exposure time	1/120 sec.
ISO speed	ISO-160
Exposure bias	0 step
Focal length	5 mm
Max aperture	1.69
Metering mode	Center Weighted Average
Subject distance	
Flash mode	No flash
Flash energy	
35mm focal length	26
Advanced photo	
Lens maker	
Lens model	

Element	'529 Claim Recitation	Exemplary Citations in Reference
		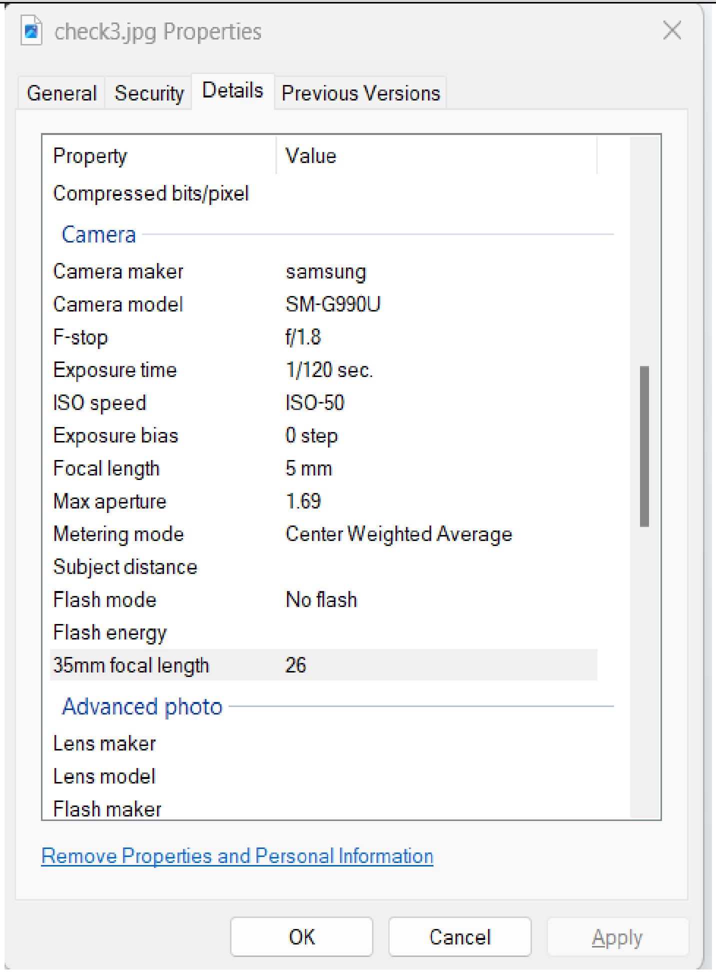

Property	Value
Compressed bits/pixel	
Camera	
Camera maker	samsung
Camera model	SM-G990U
F-stop	f/1.8
Exposure time	1/120 sec.
ISO speed	ISO-50
Exposure bias	0 step
Focal length	5 mm
Max aperture	1.69
Metering mode	Center Weighted Average
Subject distance	
Flash mode	No flash
Flash energy	
35mm focal length	26
Advanced photo	
Lens maker	
Lens model	
Flash maker	

Element	'529 Claim Recitation	Exemplary Citations in Reference
		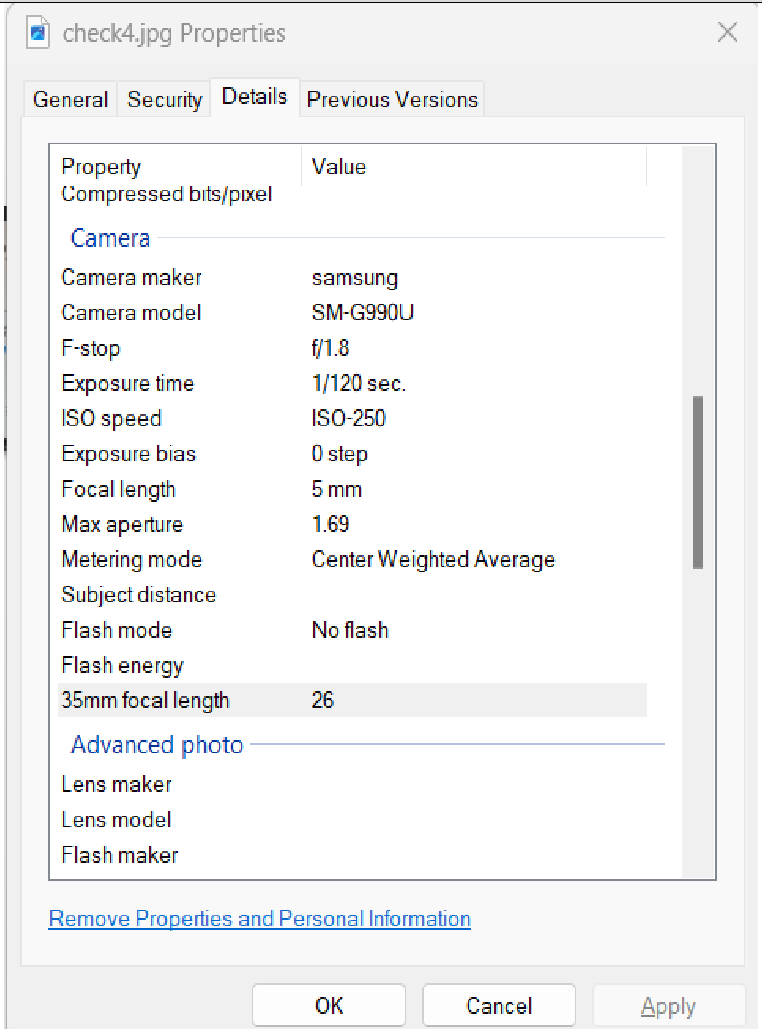

check4.jpg Properties

Property	Value
Compressed bits/pixel	
Camera	
Camera maker	samsung
Camera model	SM-G990U
F-stop	f/1.8
Exposure time	1/120 sec.
ISO speed	ISO-250
Exposure bias	0 step
Focal length	5 mm
Max aperture	1.69
Metering mode	Center Weighted Average
Subject distance	
Flash mode	No flash
Flash energy	
35mm focal length	26
Advanced photo	
Lens maker	
Lens model	
Flash maker	

Element	'529 Claim Recitation	Exemplary Citations in Reference
		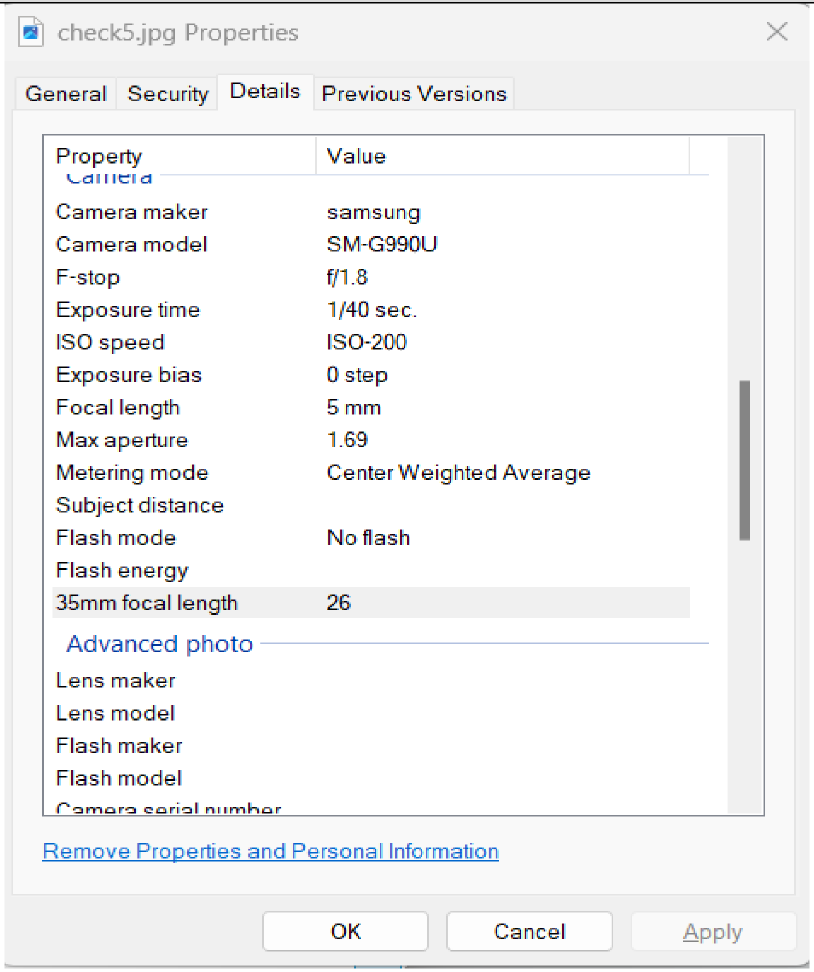 check5.jpg Properties Property Value Camera Camera maker samsung Camera model SM-G990U F-stop f/1.8 Exposure time 1/40 sec. ISO speed ISO-200 Exposure bias 0 step Focal length 5 mm Max aperture 1.69 Metering mode Center Weighted Average Subject distance Flash mode No flash Flash energy 35mm focal length 26 Advanced photo Lens maker Lens model Flash maker Flash model Camera serial number To the extent it is found that QBIC does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that QBIC does not anticipate this

Element	'529 Claim Recitation	Exemplary Citations in Reference
		claim, QBIC renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine QBIC with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1B]	a network-enabled device that conducts a data processing operation on at least a portion of the image to produce data, and sends the data to a service;	QBIC discloses a network-enabled device that conducts a data processing operation on at least a portion of the image to produce data, and sends the data to a service. The QBIC Project: Querying Images by Content Using Color, Texture, and Shape (IBM 0002390–404) IBM0002390 at 2390: 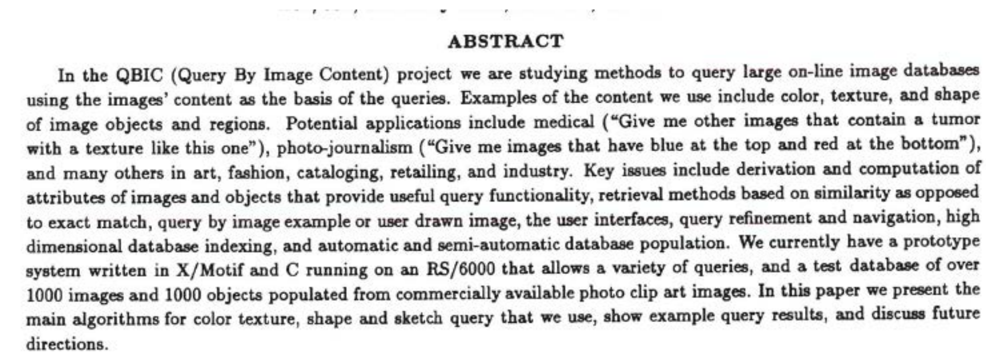**ABSTRACT** In the QBIC (Query By Image Content) project we are studying methods to query large on-line image databases using the images' content as the basis of the queries. Examples of the content we use include color, texture, and shape of image objects and regions. Potential applications include medical ("Give me other images that contain a tumor with a texture like this one"), photo-journalism ("Give me images that have blue at the top and red at the bottom"), and many others in art, fashion, cataloging, retailing, and industry. Key issues include derivation and computation of attributes of images and objects that provide useful query functionality, retrieval methods based on similarity as opposed to exact match, query by image example or user drawn image, the user interfaces, query refinement and navigation, high dimensional database indexing, and automatic and semi-automatic database population. We currently have a prototype system written in X/Motif and C running on an RS/6000 that allows a variety of queries, and a test database of over 1000 images and 1000 objects populated from commercially available photo clip art images. In this paper we present the main algorithms for color texture, shape and sketch query that we use, show example query results, and discuss future directions. IBM0002390 at 2391: 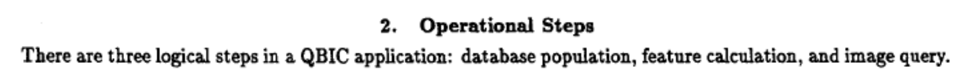**2. Operational Steps** There are three logical steps in a QBIC application: database population, feature calculation, and image query.

Element	'529 Claim Recitation	Exemplary Citations in Reference
		IBM0002390 at 2391-92: 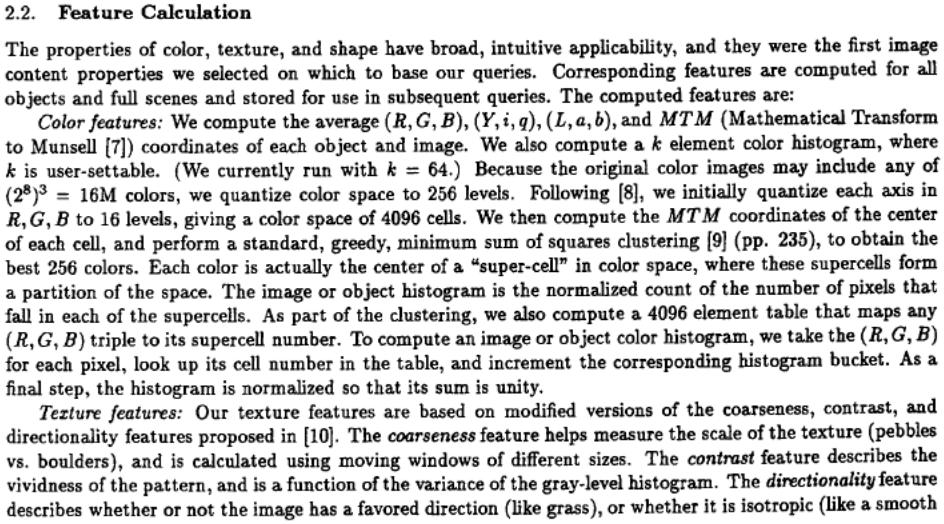 2.2. **Feature Calculation** The properties of color, texture, and shape have broad, intuitive applicability, and they were the first image content properties we selected on which to base our queries. Corresponding features are computed for all objects and full scenes and stored for use in subsequent queries. The computed features are: *Color features:* We compute the average (R, G, B), (Y, i, q), (L, a, b), and MTM (Mathematical Transform to Munsell [7]) coordinates of each object and image. We also compute a k element color histogram, where k is user-settable. (We currently run with $k = 64$.) Because the original color images may include any of $(2^8)^3 = 16$M colors, we quantize color space to 256 levels. Following [8], we initially quantize each axis in R, G, B to 16 levels, giving a color space of 4096 cells. We then compute the MTM coordinates of the center of each cell, and perform a standard, greedy, minimum sum of squares clustering [9] (pp. 235), to obtain the best 256 colors. Each color is actually the center of a "super-cell" in color space, where these supercells form a partition of the space. The image or object histogram is the normalized count of the number of pixels that fall in each of the supercells. As part of the clustering, we also compute a 4096 element table that maps any (R, G, B) triple to its supercell number. To compute an image or object color histogram, we take the (R, G, B) for each pixel, look up its cell number in the table, and increment the corresponding histogram bucket. As a final step, the histogram is normalized so that its sum is unity. *Texture features:* Our texture features are based on modified versions of the coarseness, contrast, and directionality features proposed in [10]. The *coarseness* feature helps measure the scale of the texture (pebbles vs. boulders), and is calculated using moving windows of different sizes. The *contrast* feature describes the vividness of the pattern, and is a function of the variance of the gray-level histogram. The *directionality* feature describes whether or not the image has a favored direction (like grass), or whether it is isotropic (like a smooth

Element	'529 Claim Recitation	Exemplary Citations in Reference
		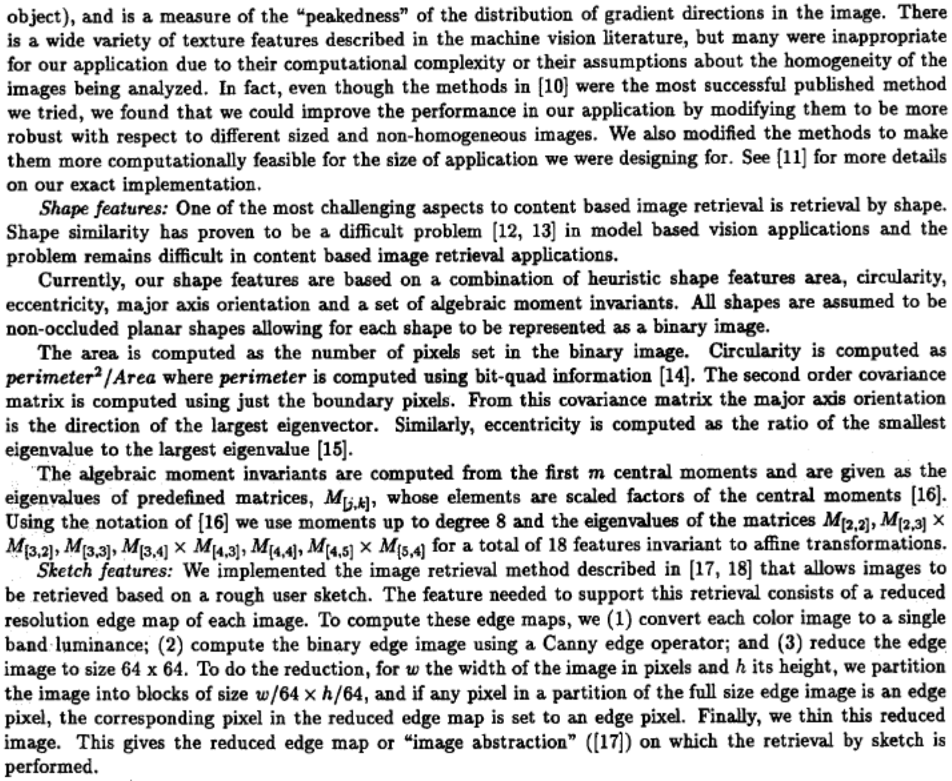 object), and is a measure of the "peakedness" of the distribution of gradient directions in the image. There is a wide variety of texture features described in the machine vision literature, but many were inappropriate for our application due to their computational complexity or their assumptions about the homogeneity of the images being analyzed. In fact, even though the methods in [10] were the most successful published method we tried, we found that we could improve the performance in our application by modifying them to be more robust with respect to different sized and non-homogeneous images. We also modified the methods to make them more computationally feasible for the size of application we were designing for. See [11] for more details on our exact implementation. *Shape features:* One of the most challenging aspects to content based image retrieval is retrieval by shape. Shape similarity has proven to be a difficult problem [12, 13] in model based vision applications and the problem remains difficult in content based image retrieval applications. Currently, our shape features are based on a combination of heuristic shape features area, circularity, eccentricity, major axis orientation and a set of algebraic moment invariants. All shapes are assumed to be non-occluded planar shapes allowing for each shape to be represented as a binary image. The area is computed as the number of pixels set in the binary image. Circularity is computed as $perimeter^2/Area$ where $perimeter$ is computed using bit-quad information [14]. The second order covariance matrix is computed using just the boundary pixels. From this covariance matrix the major axis orientation is the direction of the largest eigenvector. Similarly, eccentricity is computed as the ratio of the smallest eigenvalue to the largest eigenvalue [15]. The algebraic moment invariants are computed from the first m central moments and are given as the eigenvalues of predefined matrices, $M_{[j,k]}$, whose elements are scaled factors of the central moments [16]. Using the notation of [16] we use moments up to degree 8 and the eigenvalues of the matrices $M_{[2,2]}$, $M_{[2,3]} \times M_{[3,2]}$, $M_{[3,3]}$, $M_{[3,4]} \times M_{[4,3]}$, $M_{[4,4]}$, $M_{[4,5]} \times M_{[5,4]}$ for a total of 18 features invariant to affine transformations. *Sketch features:* We implemented the image retrieval method described in [17, 18] that allows images to be retrieved based on a rough user sketch. The feature needed to support this retrieval consists of a reduced resolution edge map of each image. To compute these edge maps, we (1) convert each color image to a single band luminance; (2) compute the binary edge image using a Canny edge operator; and (3) reduce the edge image to size 64 x 64. To do the reduction, for w the width of the image in pixels and h its height, we partition the image into blocks of size $w/64 \times h/64$, and if any pixel in a partition of the full size edge image is an edge pixel, the corresponding pixel in the reduced edge map is set to an edge pixel. Finally, we thin this reduced image. This gives the reduced edge map or "image abstraction" ([17]) on which the retrieval by sketch is performed. IBM0002390 at 2392-93 Query by Image and Video Content: The QBIC System (IBM 0000893–902): "We have developed the QBIC (Query by Image Content) system to explore content-based retrieval methods. QBIC allows queries on large image and video databases based on • example images, • user-constructed sketches and drawings, • selected color and texture patterns, • camera and object motion, and • other graphical information." *Id.* at pp. 1–3. "To achieve this functionality, QBIC has two main components: database population

Element	'529 Claim Recitation	Exemplary Citations in Reference
		(the process of creating an image database) and database query. During the population, images and videos are processed to extract features describing their content-colors, textures, shapes, and camera and object motion-and the features are stored in a database. During the query, the user composes a query graphically. Features are generated from the graphical query and then input to a matching engine that finds images or videos from the database with similar features. Figure 2 shows the system architecture." *Id.* at p. 3. 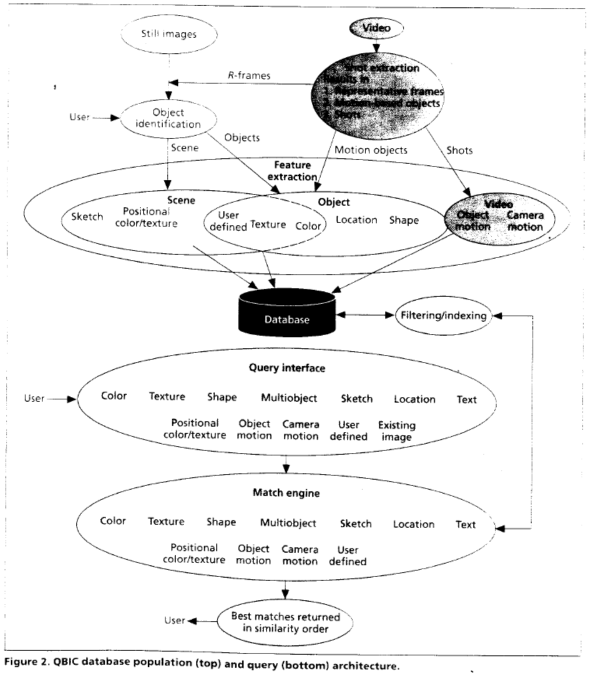 Figure 2. QBIC database population (top) and query (bottom) architecture. *Id.* at Fig. 2. Image, Audio, and Video Extenders Administration and Programming (IBM 000001–617): "Your applications can even search for images by content. Imagine an application that uses visual examples to search for images. With such an application, users could select an example image and have the application find other images that have colors or textures

Element	'529 Claim Recitation	Exemplary Citations in Reference
		similar to those in the example. With DB2 extenders' Query by Image Content (QBIC) capability, you can create applications that search for images in this visual way." *Id.* at p. 4. "To make an image available for searching by content, you catalog the image. When you catalog an image, the Image Extender analyzes the image, by computing the feature values for the image, and stores the values in a QBIC catalog. When you search for an image by content, your query identifies one or more features for the search (such as average color), a source for each feature (such as an example image), and a target set of cataloged images. The Image Extender computes the feature value of the source and compares it to the cataloged feature values for the target images. It then computes a score that indicates how similar the feature values of the target images are to the source. You can have the Image Extender return the images whose features are most similar to the source. The Image Extender will return the handle of each image and the image score. You can also have the Image Extender return only the score of a single image." *Id*. at pp. 20–21.

Element	'529 Claim Recitation	Exemplary Citations in Reference
		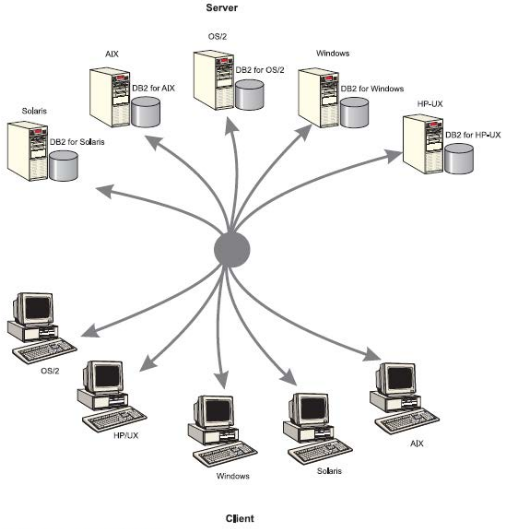 *Figure 6. DB2 extender platforms* *Id.* at p. 12 (Fig. 6). "The DB2 Extenders run in the DB2 client/server environment. This environment consists of a database server and one or more remote database clients. The DB2 extender services run on the server. Before you can access them, you have to start them." *Id.* at p. 47. Flickner Depo at 21:

Element	'529 Claim Recitation	Exemplary Citations in Reference
		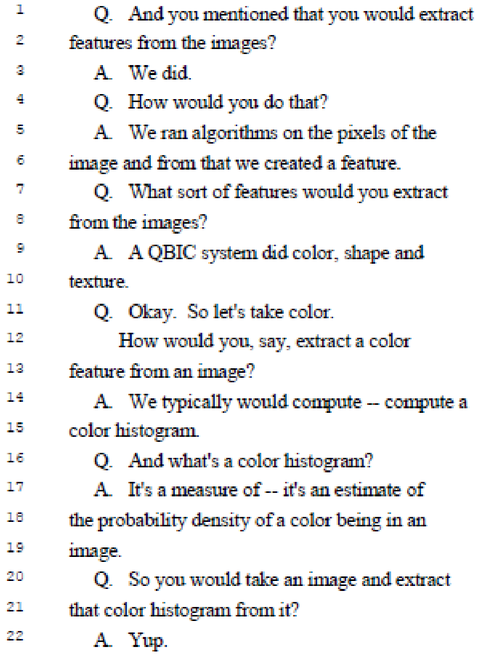 Q. And you mentioned that you would extract features from the images? A. We did. Q. How would you do that? A. We ran algorithms on the pixels of the image and from that we created a feature. Q. What sort of features would you extract from the images? A. A QBIC system did color, shape and texture. Q. Okay. So let's take color. How would you, say, extract a color feature from an image? A. We typically would compute -- compute a color histogram. Q. And what's a color histogram? A. It's a measure of -- it's an estimate of the probability density of a color being in an image. Q. So you would take an image and extract that color histogram from it? A. Yup. Flickner Depo at 212-213:

Element	'529 Claim Recitation	Exemplary Citations in Reference
		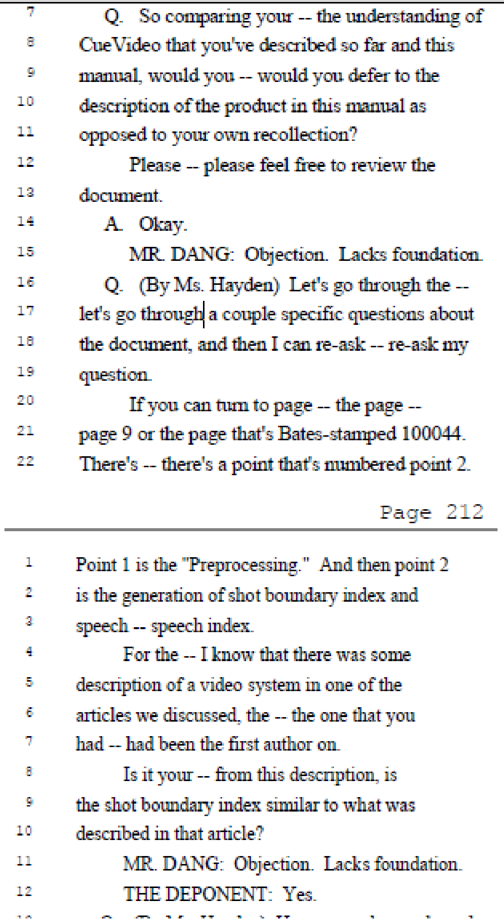 Q. So comparing your -- the understanding of CueVideo that you've described so far and this manual, would you -- would you defer to the description of the product in this manual as opposed to your own recollection? Please -- please feel free to review the document. A. Okay. MR. DANG: Objection. Lacks foundation. Q. (By Ms. Hayden) Let's go through the -- let's go through a couple specific questions about the document, and then I can re-ask -- re-ask my question. If you can turn to page -- the page -- page 9 or the page that's Bates-stamped 100044. There's -- there's a point that's numbered point 2. Page 212 Point 1 is the "Preprocessing." And then point 2 is the generation of shot boundary index and speech -- speech index. For the -- I know that there was some description of a video system in one of the articles we discussed, the -- the one that you had -- had been the first author on. Is it your -- from this description, is the shot boundary index similar to what was described in that article? MR. DANG: Objection. Lacks foundation. THE DEPONENT: Yes. Flickner Depo at 72-73:

Element	'529 Claim Recitation	Exemplary Citations in Reference
		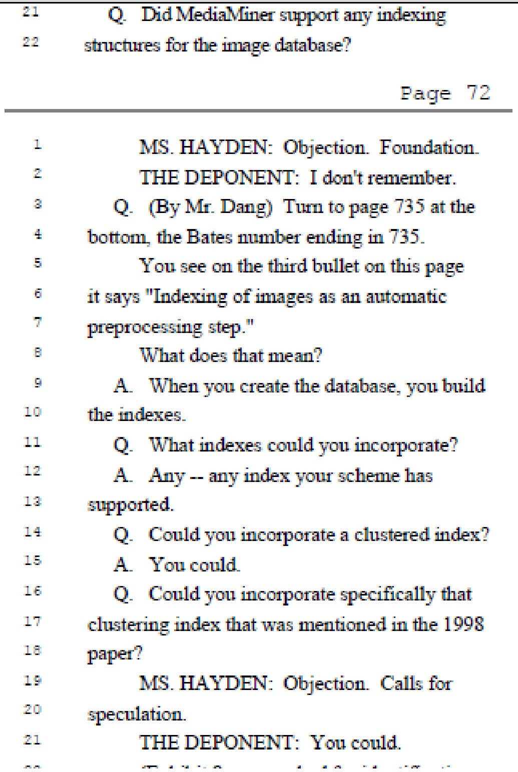 Q. Did MediaMiner support any indexing structures for the image database? Page 72 MS. HAYDEN: Objection. Foundation. THE DEPONENT: I don't remember. Q. (By Mr. Dang) Turn to page 735 at the bottom, the Bates number ending in 735. You see on the third bullet on this page it says "Indexing of images as an automatic preprocessing step." What does that mean? A. When you create the database, you build the indexes. Q. What indexes could you incorporate? A. Any -- any index your scheme has supported. Q. Could you incorporate a clustered index? A. You could. Q. Could you incorporate specifically that clustering index that was mentioned in the 1998 paper? MS. HAYDEN: Objection. Calls for speculation. THE DEPONENT: You could. Flickner Depo at 214-215:

Element	'529 Claim Recitation	Exemplary Citations in Reference
		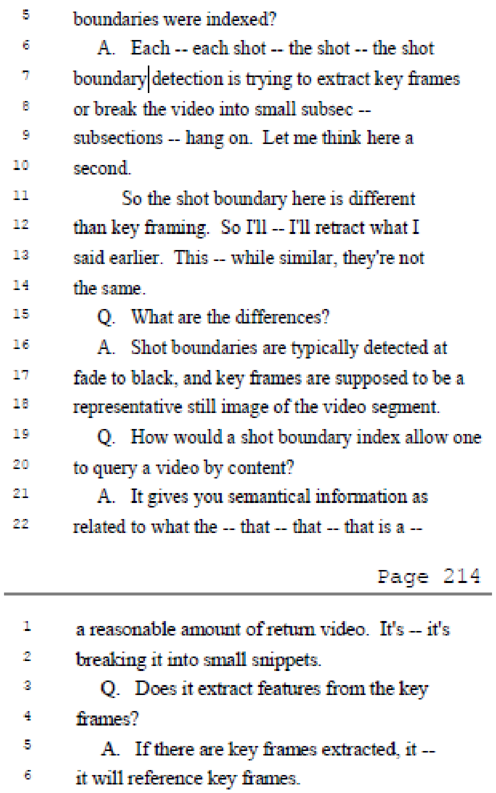 Q. What do you mean by the images of shot boundaries were indexed? A. Each -- each shot -- the shot -- the shot boundary detection is trying to extract key frames or break the video into small subsec -- subsections -- hang on. Let me think here a second. So the shot boundary here is different than key framing. So I'll -- I'll retract what I said earlier. This -- while similar, they're not the same. Q. What are the differences? A. Shot boundaries are typically detected at fade to black, and key frames are supposed to be a representative still image of the video segment. Q. How would a shot boundary index allow one to query a video by content? A. It gives you semantical information as related to what the -- that -- that -- that is a -- Page 214 a reasonable amount of return video. It's -- it's breaking it into small snippets. Q. Does it extract features from the key frames? A. If there are key frames extracted, it -- it will reference key frames. Flickner Depo at 22-25:

Element	'529 Claim Recitation	Exemplary Citations in Reference
		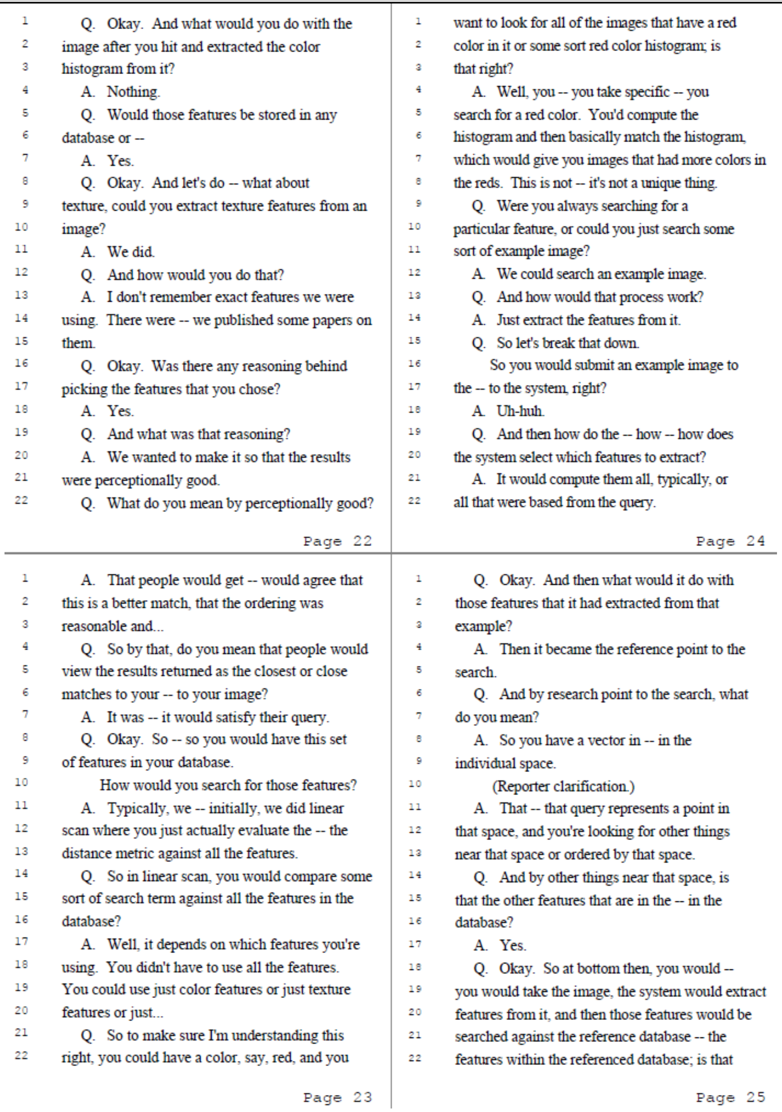

Q. Okay. And what would you do with the image after you hit and extracted the color histogram from it?

A. Nothing.

Q. Would those features be stored in any database or --

A. Yes.

Q. Okay. And let's do -- what about texture, could you extract texture features from an image?

A. We did.

Q. And how would you do that?

A. I don't remember exact features we were using. There were -- we published some papers on them.

Q. Okay. Was there any reasoning behind picking the features that you chose?

A. Yes.

Q. And what was that reasoning?

A. We wanted to make it so that the results were perceptionally good.

Q. What do you mean by perceptionally good?

Page 22

A. That people would get -- would agree that this is a better match, that the ordering was reasonable and...

Q. So by that, do you mean that people would view the results returned as the closest or close matches to your -- to your image?

A. It was -- it would satisfy their query.

Q. Okay. So -- so you would have this set of features in your database.

How would you search for those features?

A. Typically, we -- initially, we did linear scan where you just actually evaluate the -- the distance metric against all the features.

Q. So in linear scan, you would compare some sort of search term against all the features in the database?

A. Well, it depends on which features you're using. You didn't have to use all the features. You could use just color features or just texture features or just...

Q. So to make sure I'm understanding this right, you could have a color, say, red, and you

Page 23

want to look for all of the images that have a red color in it or some sort red color histogram; is that right?

A. Well, you -- you take specific -- you search for a red color. You'd compute the histogram and then basically match the histogram, which would give you images that had more colors in the reds. This is not -- it's not a unique thing.

Q. Were you always searching for a particular feature, or could you just search some sort of example image?

A. We could search an example image.

Q. And how would that process work?

A. Just extract the features from it.

Q. So let's break that down.

So you would submit an example image to the -- to the system, right?

A. Uh-huh.

Q. And then how do the -- how -- how does the system select which features to extract?

A. It would compute them all, typically, or all that were based from the query.

Page 24

Q. Okay. And then what would it do with those features that it had extracted from that example?

A. Then it became the reference point to the search.

Q. And by research point to the search, what do you mean?

A. So you have a vector in -- in the individual space.

(Reporter clarification.)

A. That -- that query represents a point in that space, and you're looking for other things near that space or ordered by that space.

Q. And by other things near that space, is that the other features that are in the -- in the database?

A. Yes.

Q. Okay. So at bottom then, you would -- you would take the image, the system would extract features from it, and then those features would be searched against the reference database -- the features within the referenced database; is that

Page 25

Element	'529 Claim Recitation	Exemplary Citations in Reference
		Flickner Depo at 88-89: 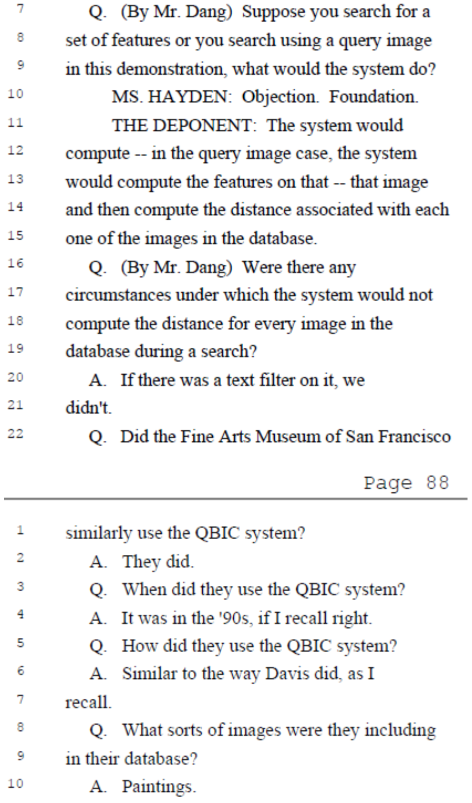 Q. (By Mr. Dang) Suppose you search for a set of features or you search using a query image in this demonstration, what would the system do? MS. HAYDEN: Objection. Foundation. THE DEPONENT: The system would compute -- in the query image case, the system would compute the features on that -- that image and then compute the distance associated with each one of the images in the database. Q. (By Mr. Dang) Were there any circumstances under which the system would not compute the distance for every image in the database during a search? A. If there was a text filter on it, we didn't. Q. Did the Fine Arts Museum of San Francisco Page 88 similarly use the QBIC system? A. They did. Q. When did they use the QBIC system? A. It was in the '90s, if I recall right. Q. How did they use the QBIC system? A. Similar to the way Davis did, as I recall. Q. What sorts of images were they including in their database? A. Paintings. Flickner Depo at 91-92:

Element	'529 Claim Recitation	Exemplary Citations in Reference
		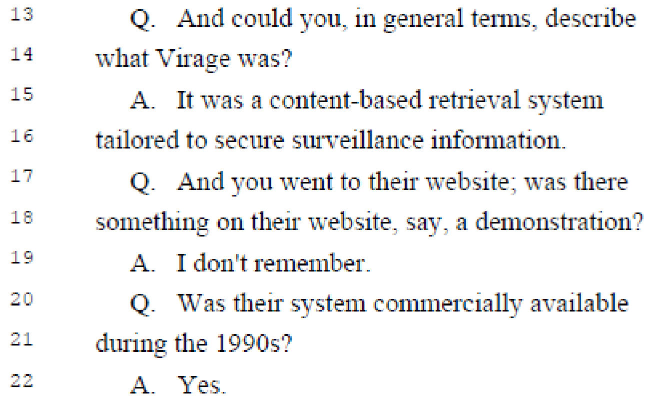 Q. And could you, in general terms, describe what Virage was? A. It was a content-based retrieval system tailored to secure surveillance information. Q. And you went to their website; was there something on their website, say, a demonstration? A. I don't remember. Q. Was their system commercially available during the 1990s? A. Yes. 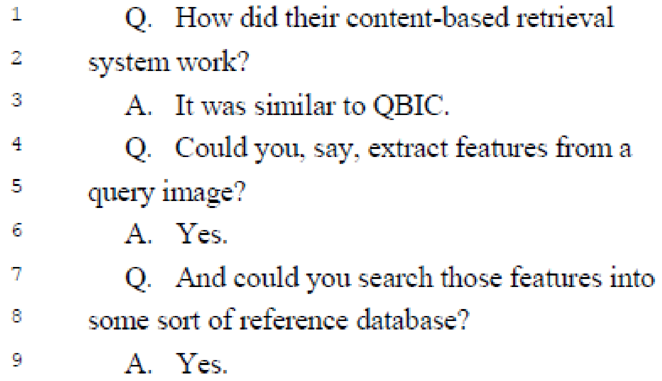 Q. How did their content-based retrieval system work? A. It was similar to QBIC. Q. Could you, say, extract features from a query image? A. Yes. Q. And could you search those features into some sort of reference database? A. Yes. *See also* Flickner Depo at 54-55. Flickner 8/29/23 Depo at 79–80:

Element	'529 Claim Recitation	Exemplary Citations in Reference
		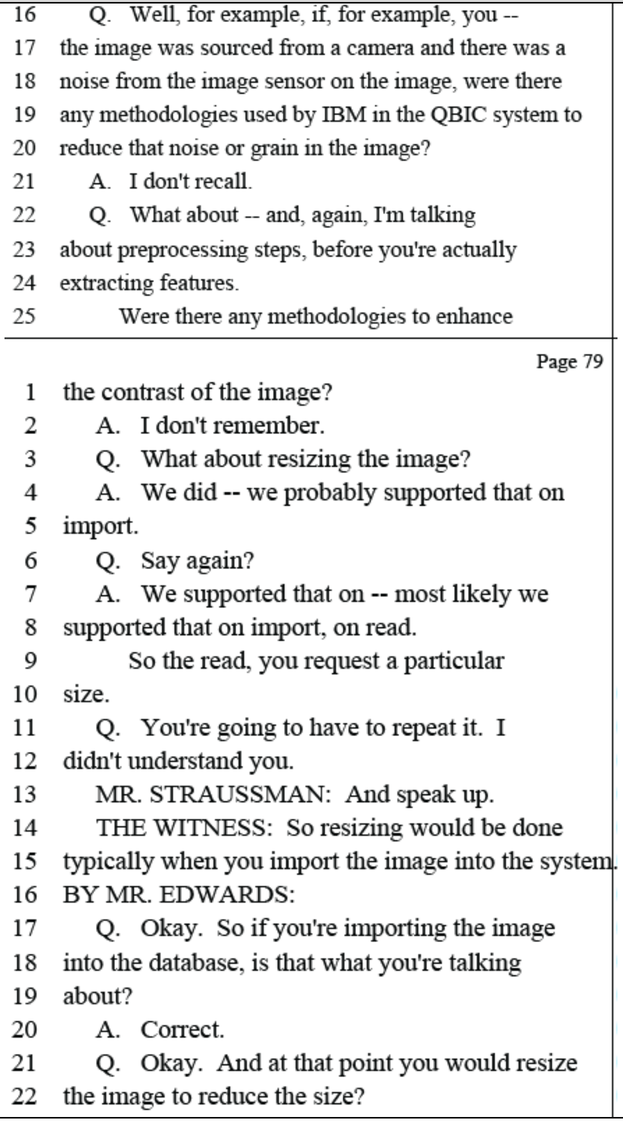 Q. Well, for example, if, for example, you -- the image was sourced from a camera and there was a noise from the image sensor on the image, were there any methodologies used by IBM in the QBIC system to reduce that noise or grain in the image? A. I don't recall. Q. What about -- and, again, I'm talking about preprocessing steps, before you're actually extracting features. Were there any methodologies to enhance Page 79 the contrast of the image? A. I don't remember. Q. What about resizing the image? A. We did -- we probably supported that on import. Q. Say again? A. We supported that on -- most likely we supported that on import, on read. So the read, you request a particular size. Q. You're going to have to repeat it. I didn't understand you. MR. STRAUSSMAN: And speak up. THE WITNESS: So resizing would be done typically when you import the image into the system. BY MR. EDWARDS: Q. Okay. So if you're importing the image into the database, is that what you're talking about? A. Correct. Q. Okay. And at that point you would resize the image to reduce the size?

Element	'529 Claim Recitation	Exemplary Citations in Reference
		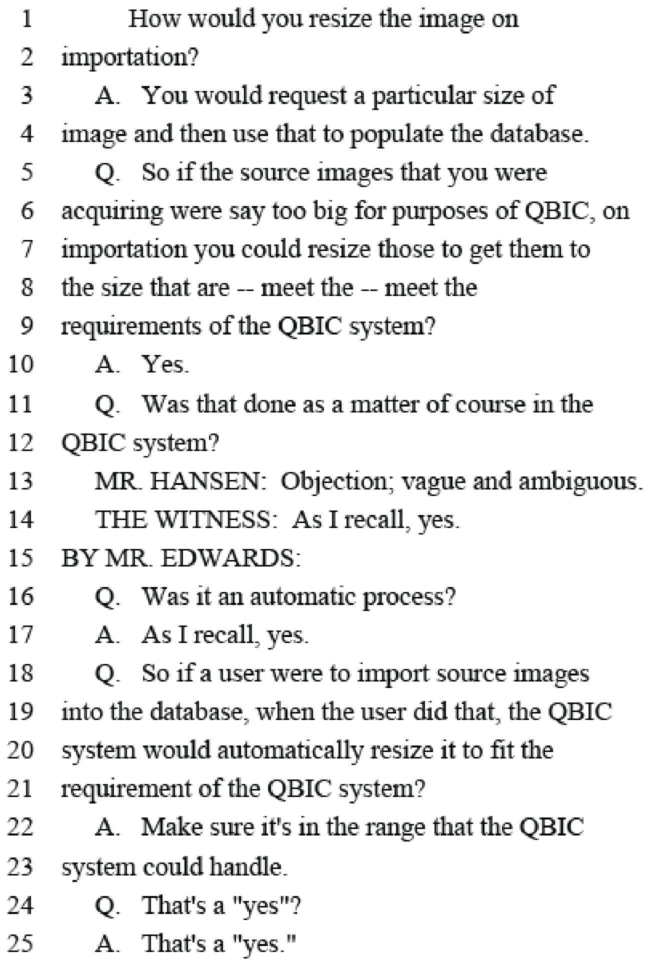 How would you resize the image on importation? A. You would request a particular size of image and then use that to populate the database. Q. So if the source images that you were acquiring were say too big for purposes of QBIC, on importation you could resize those to get them to the size that are -- meet the -- meet the requirements of the QBIC system? A. Yes. Q. Was that done as a matter of course in the QBIC system? MR. HANSEN: Objection; vague and ambiguous. THE WITNESS: As I recall, yes. BY MR. EDWARDS: Q. Was it an automatic process? A. As I recall, yes. Q. So if a user were to import source images into the database, when the user did that, the QBIC system would automatically resize it to fit the requirement of the QBIC system? A. Make sure it's in the range that the QBIC system could handle. Q. That's a "yes"? A. That's a "yes." *Id.* at 80–81:

Element	'529 Claim Recitation	Exemplary Citations in Reference
		Q. How about converting the images to binary? A. Threshold. There probably was a threshold control. Q. And what do you mean by "threshold control"? A. Just change the value of the threshold. You make a binary decision, is a pixel bigger than a -- is a pixel bigger than a particular number or is it smaller and you put a 01 based on that. Q. And that would have been done prior to feature extraction? A. For some features it would be done that way. Q. So, for example, for shape, was it required to convert the image to binary before you did your shape analysis? A. No. Q. Were there any other features that, as part of the -- as part of the extraction process, the image had to be converted to binary before?

Element	'529 Claim Recitation	Exemplary Citations in Reference
		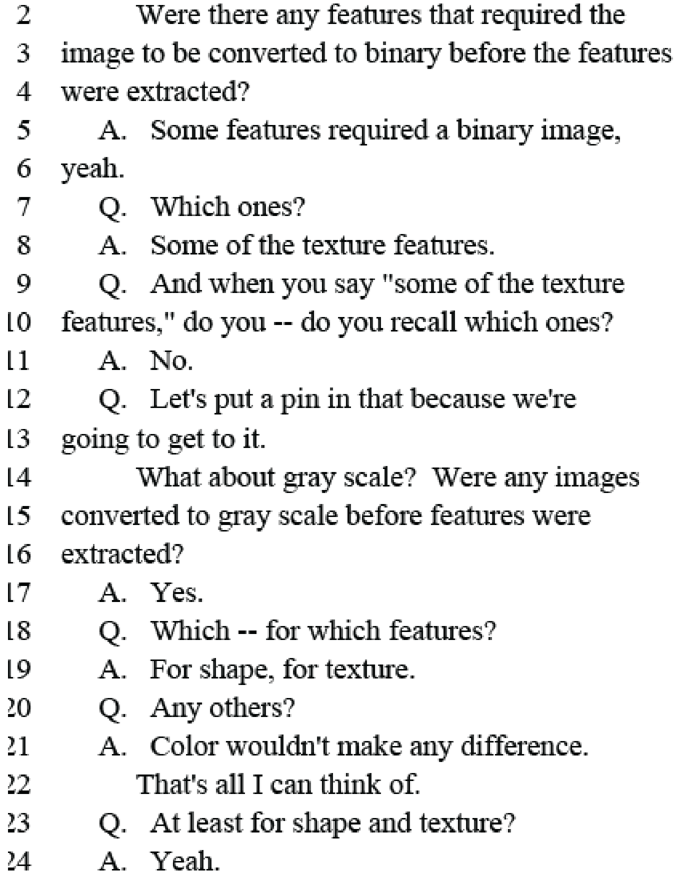 Were there any features that required the image to be converted to binary before the features were extracted? A. Some features required a binary image, yeah. Q. Which ones? A. Some of the texture features. Q. And when you say "some of the texture features," do you -- do you recall which ones? A. No. Q. Let's put a pin in that because we're going to get to it. What about gray scale? Were any images converted to gray scale before features were extracted? A. Yes. Q. Which -- for which features? A. For shape, for texture. Q. Any others? A. Color wouldn't make any difference. That's all I can think of. Q. At least for shape and texture? A. Yeah. *Id.* at 94–97: 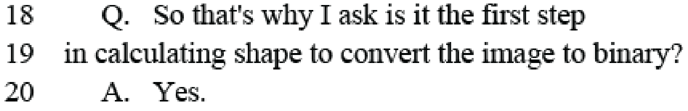 Q. So that's why I ask is it the first step in calculating shape to convert the image to binary? A. Yes. *Id.* at 101–102:

Element	'529 Claim Recitation	Exemplary Citations in Reference
		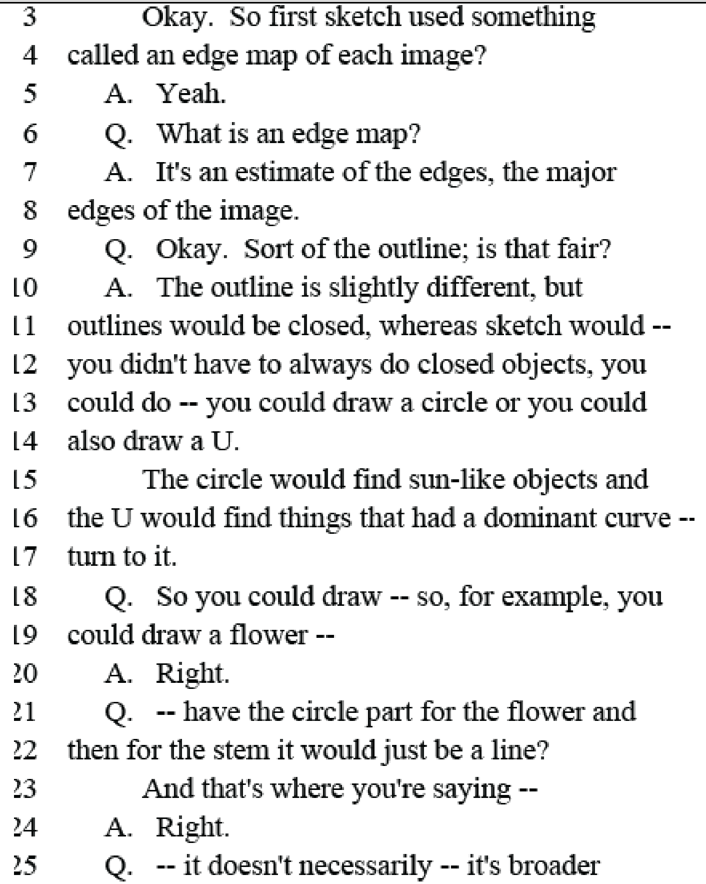 Okay. So first sketch used something called an edge map of each image? A. Yeah. Q. What is an edge map? A. It's an estimate of the edges, the major edges of the image. Q. Okay. Sort of the outline; is that fair? A. The outline is slightly different, but outlines would be closed, whereas sketch would -- you didn't have to always do closed objects, you could do -- you could draw a circle or you could also draw a U. The circle would find sun-like objects and the U would find things that had a dominant curve -- turn to it. Q. So you could draw -- so, for example, you could draw a flower -- A. Right. Q. -- have the circle part for the flower and then for the stem it would just be a line? And that's where you're saying -- A. Right. Q. -- it doesn't necessarily -- it's broader

Element	'529 Claim Recitation	Exemplary Citations in Reference
		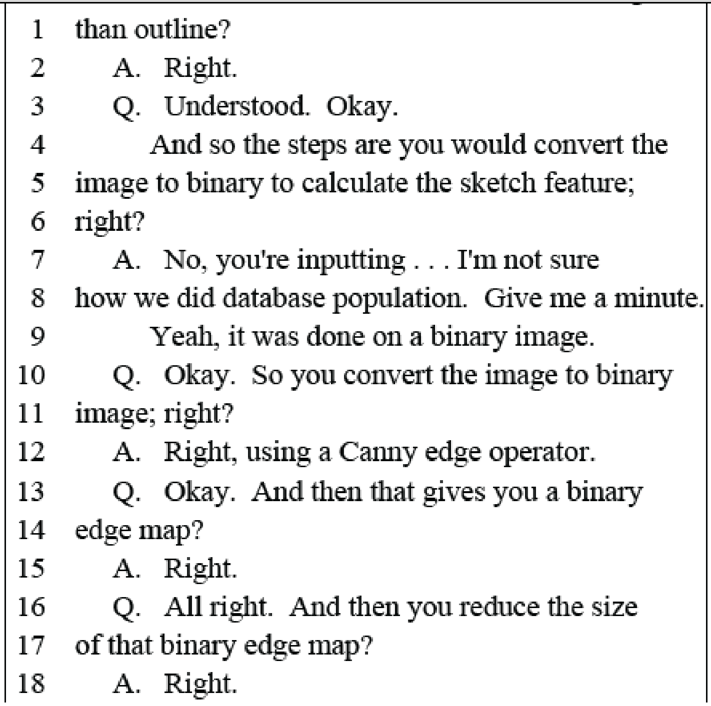 than outline? A. Right. Q. Understood. Okay. And so the steps are you would convert the image to binary to calculate the sketch feature; right? A. No, you're inputting . . . I'm not sure how we did database population. Give me a minute. Yeah, it was done on a binary image. Q. Okay. So you convert the image to binary image; right? A. Right, using a Canny edge operator. Q. Okay. And then that gives you a binary edge map? A. Right. Q. All right. And then you reduce the size of that binary edge map? A. Right. *Id.* at 140–141:

Element	'529 Claim Recitation	Exemplary Citations in Reference
		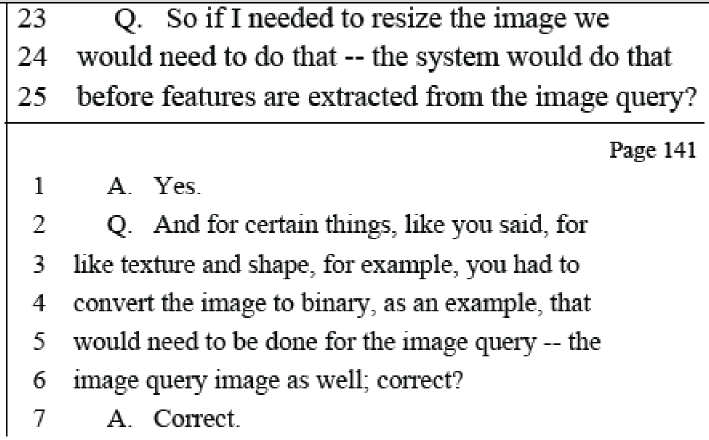 Q. So if I needed to resize the image we would need to do that -- the system would do that before features are extracted from the image query? Page 141 A. Yes. Q. And for certain things, like you said, for like texture and shape, for example, you had to convert the image to binary, as an example, that would need to be done for the image query -- the image query image as well; correct? A. Correct. *Id*. at 164–165:

Element	'529 Claim Recitation	Exemplary Citations in Reference
		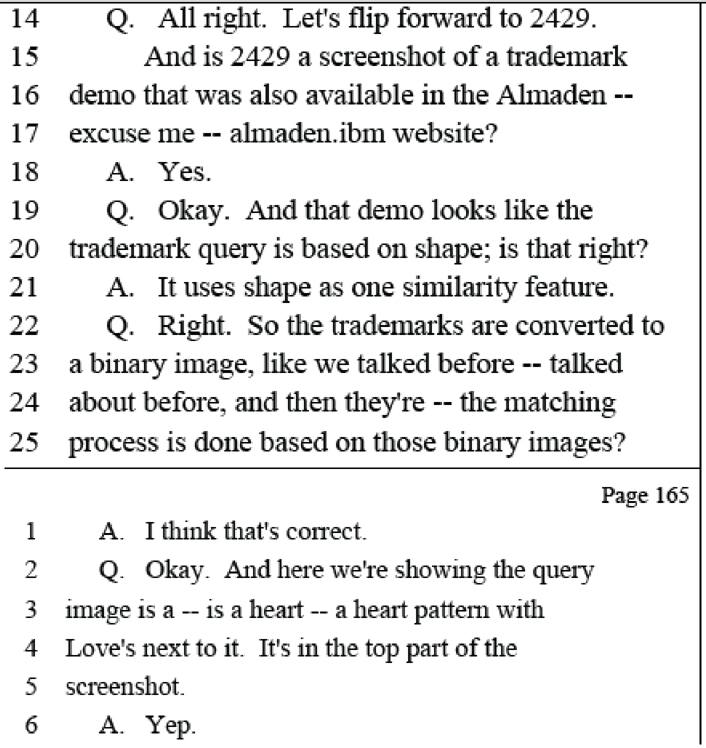 Q. All right. Let's flip forward to 2429. And is 2429 a screenshot of a trademark demo that was also available in the Almaden -- excuse me -- almaden.ibm website? A. Yes. Q. Okay. And that demo looks like the trademark query is based on shape; is that right? A. It uses shape as one similarity feature. Q. Right. So the trademarks are converted to a binary image, like we talked before -- talked about before, and then they're -- the matching process is done based on those binary images? Page 165 A. I think that's correct. Q. Okay. And here we're showing the query image is a -- is a heart -- a heart pattern with Love's next to it. It's in the top part of the screenshot. A. Yep. *Id*. at 82: 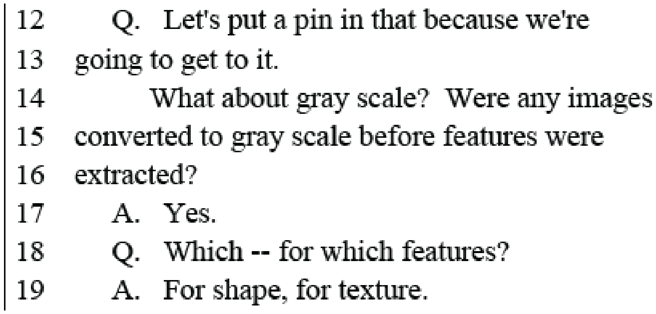 Q. Let's put a pin in that because we're going to get to it. What about gray scale? Were any images converted to gray scale before features were extracted? A. Yes. Q. Which -- for which features? A. For shape, for texture.

Element	'529 Claim Recitation	Exemplary Citations in Reference
		Id. at 83-85: 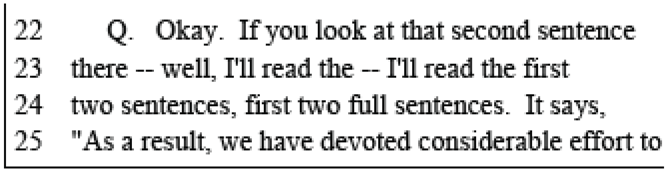 Q. Okay. If you look at that second sentence there -- well, I'll read the -- I'll read the first two sentences, first two full sentences. It says, "As a result, we have devoted considerable effort to

Element	'529 Claim Recitation	Exemplary Citations in Reference
		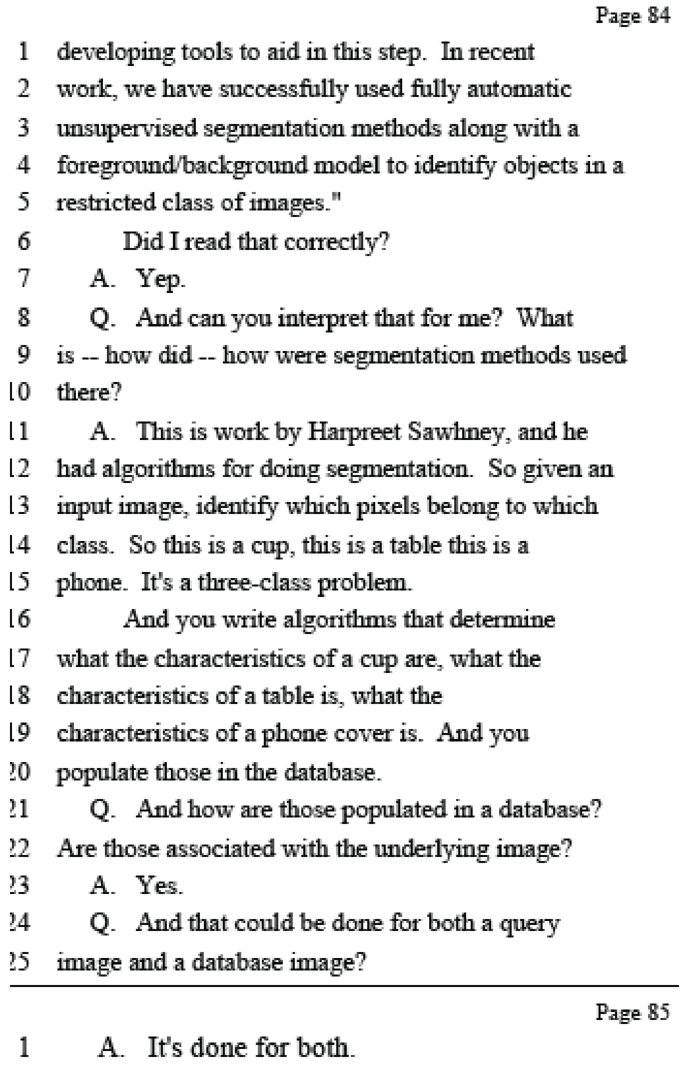Page 84 developing tools to aid in this step. In recent work, we have successfully used fully automatic unsupervised segmentation methods along with a foreground/background model to identify objects in a restricted class of images." Did I read that correctly? A. Yep. Q. And can you interpret that for me? What is -- how did -- how were segmentation methods used there? A. This is work by Harpreet Sawhney, and he had algorithms for doing segmentation. So given an input image, identify which pixels belong to which class. So this is a cup, this is a table this is a phone. It's a three-class problem. And you write algorithms that determine what the characteristics of a cup are, what the characteristics of a table is, what the characteristics of a phone cover is. And you populate those in the database. Q. And how are those populated in a database? Are those associated with the underlying image? A. Yes. Q. And that could be done for both a query image and a database image? Page 85 A. It's done for both.

Element	'529 Claim Recitation	Exemplary Citations in Reference
		Id. at 305: 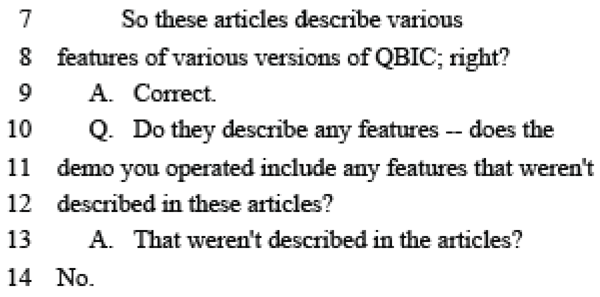 So these articles describe various features of various versions of QBIC; right? A. Correct. Q. Do they describe any features -- does the demo you operated include any features that weren't described in these articles? A. That weren't described in the articles? No. *Id.* at 35: 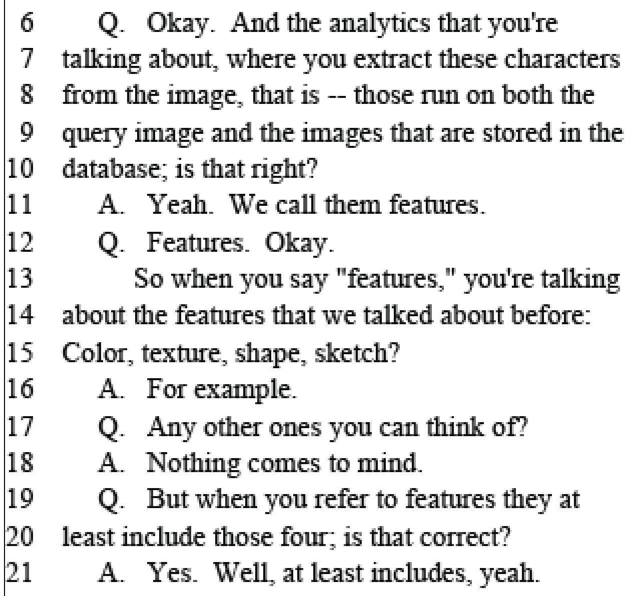 Q. Okay. And the analytics that you're talking about, where you extract these characters from the image, that is -- those run on both the query image and the images that are stored in the database; is that right? A. Yeah. We call them features. Q. Features. Okay. So when you say "features," you're talking about the features that we talked about before: Color, texture, shape, sketch? A. For example. Q. Any other ones you can think of? A. Nothing comes to mind. Q. But when you refer to features they at least include those four; is that correct? A. Yes. Well, at least includes, yeah.

Element	'529 Claim Recitation	Exemplary Citations in Reference
		Id. at 109-110: 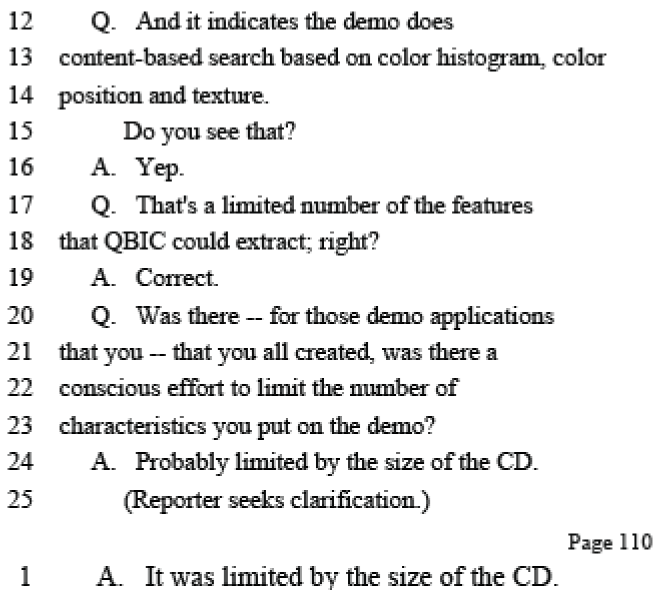 Q. And it indicates the demo does content-based search based on color histogram, color position and texture. Do you see that? A. Yep. Q. That's a limited number of the features that QBIC could extract; right? A. Correct. Q. Was there -- for those demo applications that you -- that you all created, was there a conscious effort to limit the number of characteristics you put on the demo? A. Probably limited by the size of the CD. (Reporter seeks clarification.) Page 110 A. It was limited by the size of the CD. *Id.* at 311: 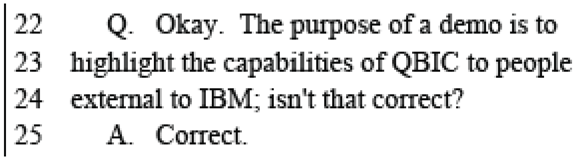 Q. Okay. The purpose of a demo is to highlight the capabilities of QBIC to people external to IBM; isn't that correct? A. Correct. *Id.* at 313-317:

Element	'529 Claim Recitation	Exemplary Citations in Reference
		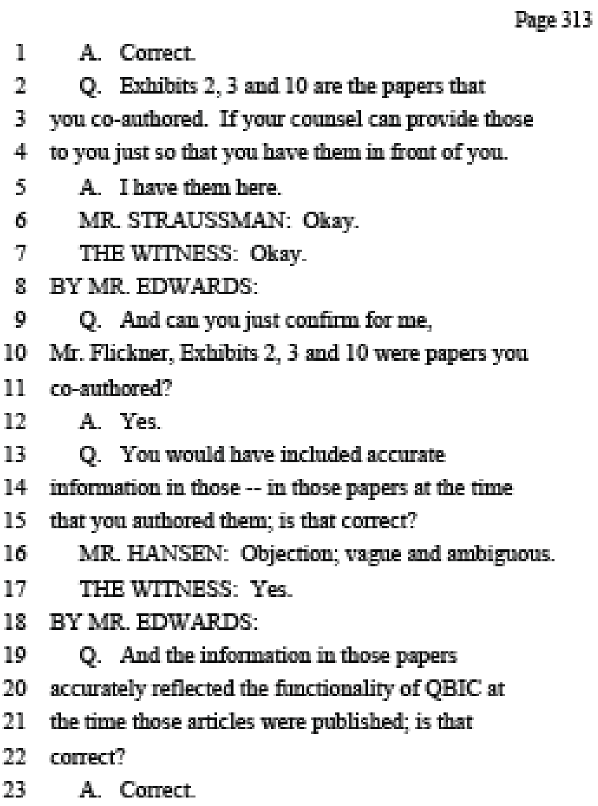Page 313 A. Correct. Q. Exhibits 2, 3 and 10 are the papers that you co-authored. If your counsel can provide those to you just so that you have them in front of you. A. I have them here. MR. STRAUSSMAN: Okay. THE WITNESS: Okay. BY MR. EDWARDS: Q. And can you just confirm for me, Mr. Flickner, Exhibits 2, 3 and 10 were papers you co-authored? A. Yes. Q. You would have included accurate information in those -- in those papers at the time that you authored them; is that correct? MR. HANSEN: Objection; vague and ambiguous. THE WITNESS: Yes. BY MR. EDWARDS: Q. And the information in those papers accurately reflected the functionality of QBIC at the time those articles were published; is that correct? A. Correct.

		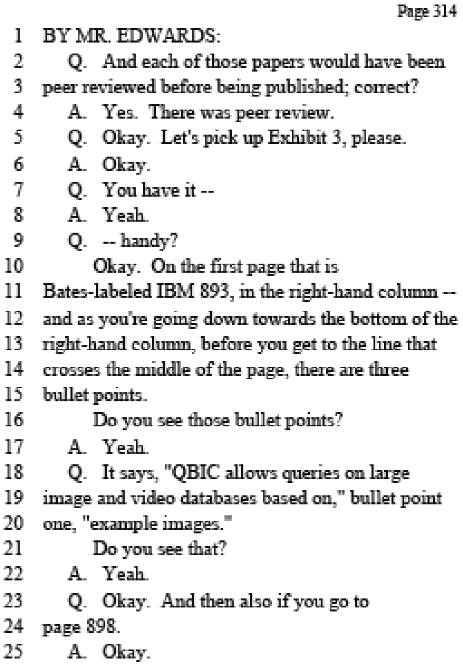 Page 314 BY MR. EDWARDS: Q. And each of those papers would have been peer reviewed before being published; correct? A. Yes. There was peer review. Q. Okay. Let's pick up Exhibit 3, please. A. Okay. Q. You have it -- A. Yeah. Q. -- handy? Okay. On the first page that is Bates-labeled IBM 893, in the right-hand column -- and as you're going down towards the bottom of the right-hand column, before you get to the line that crosses the middle of the page, there are three bullet points. Do you see those bullet points? A. Yeah. Q. It says, "QBIC allows queries on large image and video databases based on," bullet point one, "example images." Do you see that? A. Yeah. Q. Okay. And then also if you go to page 898. A. Okay. 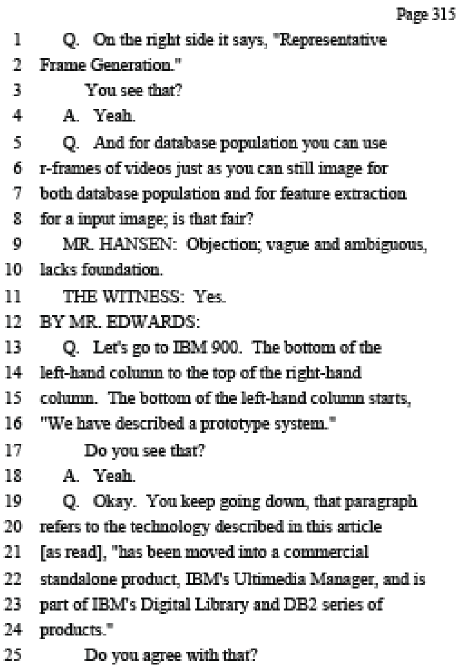 Page 315 Q. On the right side it says, "Representative Frame Generation." You see that? A. Yeah. Q. And for database population you can use r-frames of videos just as you can still image for both database population and for feature extraction for a input image; is that fair? MR. HANSEN: Objection; vague and ambiguous, lacks foundation. THE WITNESS: Yes. BY MR. EDWARDS: Q. Let's go to IBM 900. The bottom of the left-hand column to the top of the right-hand column. The bottom of the left-hand column starts, "We have described a prototype system." Do you see that? A. Yeah. Q. Okay. You keep going down, that paragraph refers to the technology described in this article [as read], "has been moved into a commercial standalone product, IBM's Ultimedia Manager, and is part of IBM's Digital Library and DB2 series of products." Do you agree with that?

Element	'529 Claim Recitation	Exemplary Citations in Reference
		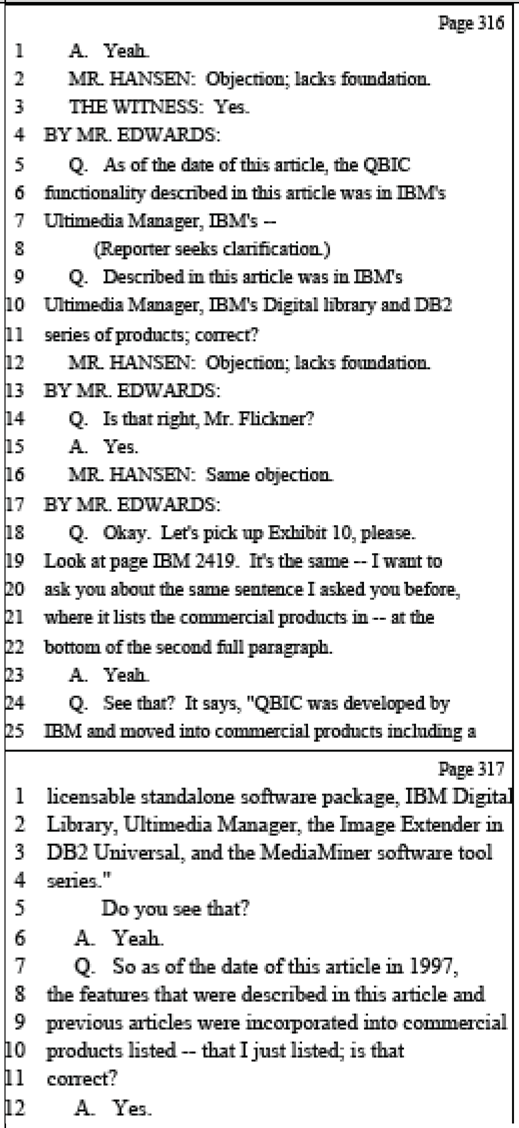 Page 316 A. Yeah. MR. HANSEN: Objection; lacks foundation. THE WITNESS: Yes. BY MR. EDWARDS: Q. As of the date of this article, the QBIC functionality described in this article was in IBM's Ultimedia Manager, IBM's -- (Reporter seeks clarification.) Q. Described in this article was in IBM's Ultimedia Manager, IBM's Digital library and DB2 series of products; correct? MR. HANSEN: Objection; lacks foundation. BY MR. EDWARDS: Q. Is that right, Mr. Flickner? A. Yes. MR. HANSEN: Same objection. BY MR. EDWARDS: Q. Okay. Let's pick up Exhibit 10, please. Look at page IBM 2419. It's the same -- I want to ask you about the same sentence I asked you before, where it lists the commercial products in -- at the bottom of the second full paragraph. A. Yeah. Q. See that? It says, "QBIC was developed by IBM and moved into commercial products including a Page 317 licensable standalone software package, IBM Digital Library, Ultimedia Manager, the Image Extender in DB2 Universal, and the MediaMiner software tool series." Do you see that? A. Yeah. Q. So as of the date of this article in 1997, the features that were described in this article and previous articles were incorporated into commercial products listed -- that I just listed; is that correct? A. Yes.

Element	'529 Claim Recitation	Exemplary Citations in Reference
		To the extent it is found that QBIC does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that QBIC does not anticipate this claim, QBIC renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine QBIC with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1C]	the service programmed to receive the data; identify an object within the image:	QBIC discloses the service programmed to receive the data; identify an object within the image. The QBIC Project: Querying Images by Content Using Color, Texture, and Shape (IBM 0002390–404) at IBM0002390 at 2390 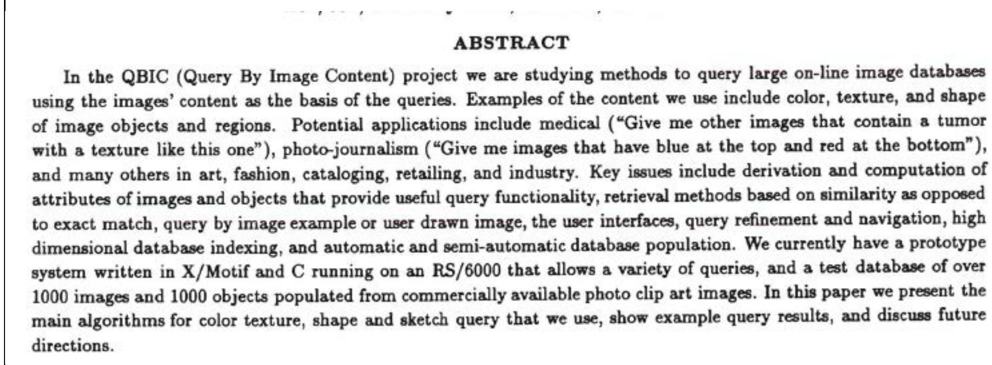 ABSTRACT In the QBIC (Query By Image Content) project we are studying methods to query large on-line image databases using the images' content as the basis of the queries. Examples of the content we use include color, texture, and shape of image objects and regions. Potential applications include medical ("Give me other images that contain a tumor with a texture like this one"), photo-journalism ("Give me images that have blue at the top and red at the bottom"), and many others in art, fashion, cataloging, retailing, and industry. Key issues include derivation and computation of attributes of images and objects that provide useful query functionality, retrieval methods based on similarity as opposed to exact match, query by image example or user drawn image, the user interfaces, query refinement and navigation, high dimensional database indexing, and automatic and semi-automatic database population. We currently have a prototype system written in X/Motif and C running on an RS/6000 that allows a variety of queries, and a test database of over 1000 images and 1000 objects populated from commercially available photo clip art images. In this paper we present the main algorithms for color texture, shape and sketch query that we use, show example query results, and discuss future directions. IBM0002390 at 2391:

Element	'529 Claim Recitation	Exemplary Citations in Reference
		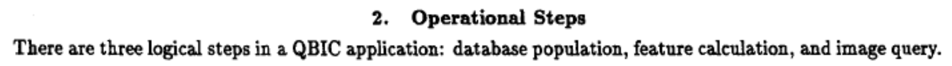 **2. Operational Steps** There are three logical steps in a QBIC application: database population, feature calculation, and image query. IBM0002390 at 2391-92 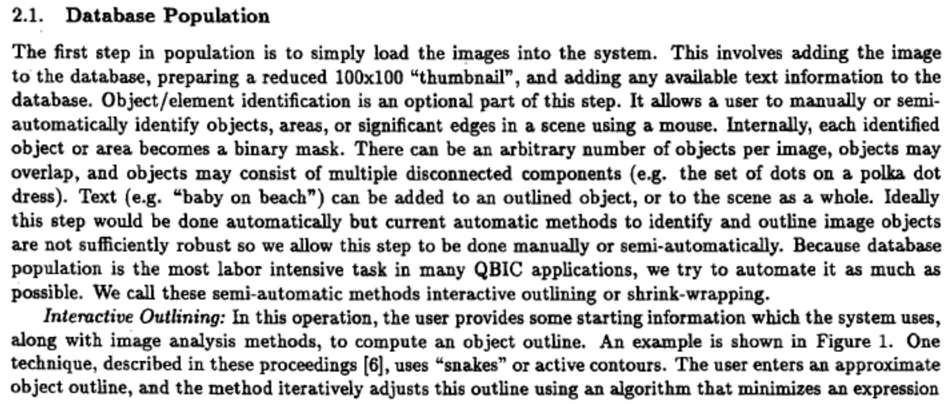 **2.1. Database Population** The first step in population is to simply load the images into the system. This involves adding the image to the database, preparing a reduced 100x100 "thumbnail", and adding any available text information to the database. Object/element identification is an optional part of this step. It allows a user to manually or semi-automatically identify objects, areas, or significant edges in a scene using a mouse. Internally, each identified object or area becomes a binary mask. There can be an arbitrary number of objects per image, objects may overlap, and objects may consist of multiple disconnected components (e.g. the set of dots on a polka dot dress). Text (e.g. "baby on beach") can be added to an outlined object, or to the scene as a whole. Ideally this step would be done automatically but current automatic methods to identify and outline image objects are not sufficiently robust so we allow this step to be done manually or semi-automatically. Because database population is the most labor intensive task in many QBIC applications, we try to automate it as much as possible. We call these semi-automatic methods interactive outlining or shrink-wrapping. *Interactive Outlining:* In this operation, the user provides some starting information which the system uses, along with image analysis methods, to compute an object outline. An example is shown in Figure 1. One technique, described in these proceedings [6], uses "snakes" or active contours. The user enters an approximate object outline, and the method iteratively adjusts this outline using an algorithm that minimizes an expression

Element	'529 Claim Recitation	Exemplary Citations in Reference
		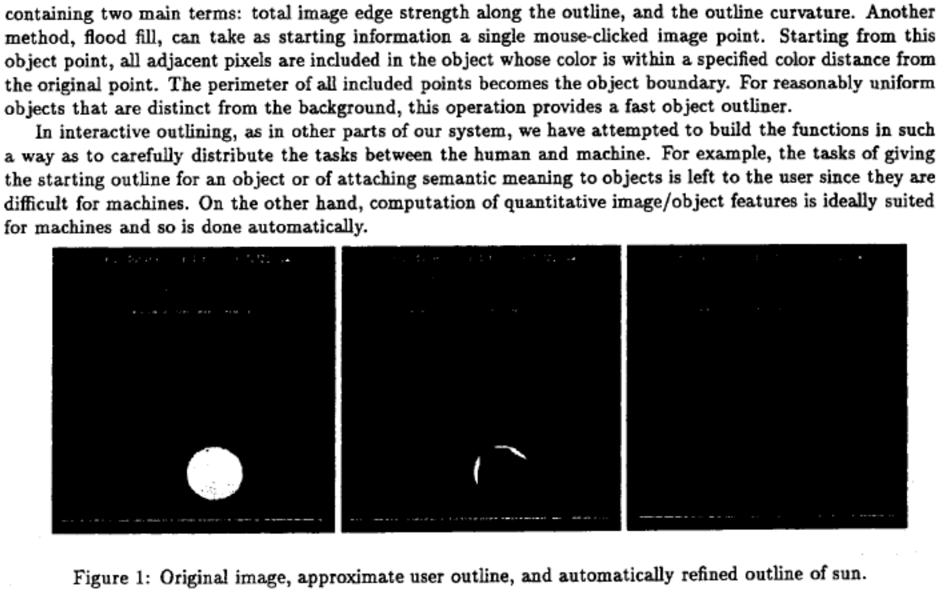 containing two main terms: total image edge strength along the outline, and the outline curvature. Another method, flood fill, can take as starting information a single mouse-clicked image point. Starting from this object point, all adjacent pixels are included in the object whose color is within a specified color distance from the original point. The perimeter of all included points becomes the object boundary. For reasonably uniform objects that are distinct from the background, this operation provides a fast object outliner. In interactive outlining, as in other parts of our system, we have attempted to build the functions in such a way as to carefully distribute the tasks between the human and machine. For example, the tasks of giving the starting outline for an object or of attaching semantic meaning to objects is left to the user since they are difficult for machines. On the other hand, computation of quantitative image/object features is ideally suited for machines and so is done automatically. Figure 1: Original image, approximate user outline, and automatically refined outline of sun. IBM0002390 at 2392-93

Element	'529 Claim Recitation	Exemplary Citations in Reference
		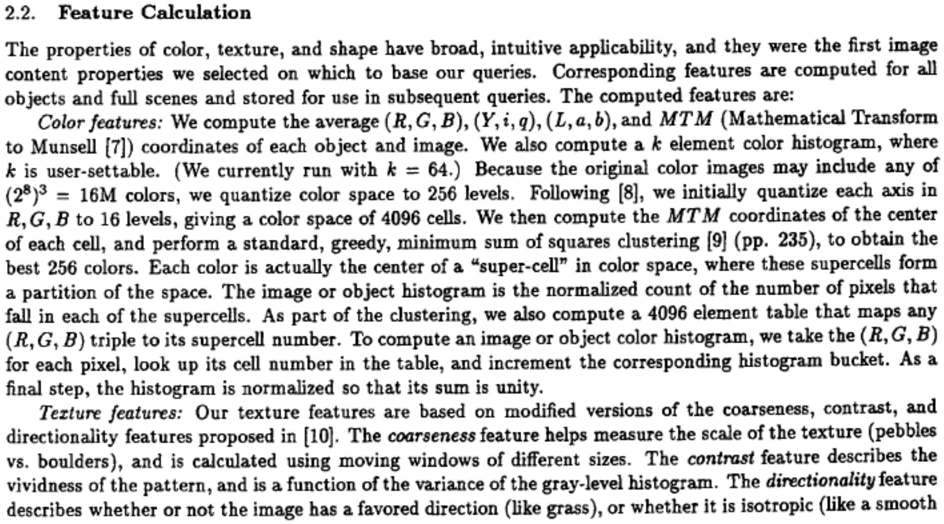 2.2. **Feature Calculation** The properties of color, texture, and shape have broad, intuitive applicability, and they were the first image content properties we selected on which to base our queries. Corresponding features are computed for all objects and full scenes and stored for use in subsequent queries. The computed features are: *Color features:* We compute the average (R,G,B), (Y,i,q), (L,a,b), and MTM (Mathematical Transform to Munsell [7]) coordinates of each object and image. We also compute a k element color histogram, where k is user-settable. (We currently run with $k = 64$.) Because the original color images may include any of $(2^8)^3 = 16$M colors, we quantize color space to 256 levels. Following [8], we initially quantize each axis in R, G, B to 16 levels, giving a color space of 4096 cells. We then compute the MTM coordinates of the center of each cell, and perform a standard, greedy, minimum sum of squares clustering [9] (pp. 235), to obtain the best 256 colors. Each color is actually the center of a "super-cell" in color space, where these supercells form a partition of the space. The image or object histogram is the normalized count of the number of pixels that fall in each of the supercells. As part of the clustering, we also compute a 4096 element table that maps any (R,G,B) triple to its supercell number. To compute an image or object color histogram, we take the (R,G,B) for each pixel, look up its cell number in the table, and increment the corresponding histogram bucket. As a final step, the histogram is normalized so that its sum is unity. *Texture features:* Our texture features are based on modified versions of the coarseness, contrast, and directionality features proposed in [10]. The *coarseness* feature helps measure the scale of the texture (pebbles vs. boulders), and is calculated using moving windows of different sizes. The *contrast* feature describes the vividness of the pattern, and is a function of the variance of the gray-level histogram. The *directionality* feature describes whether or not the image has a favored direction (like grass), or whether it is isotropic (like a smooth

Element	'529 Claim Recitation	Exemplary Citations in Reference
		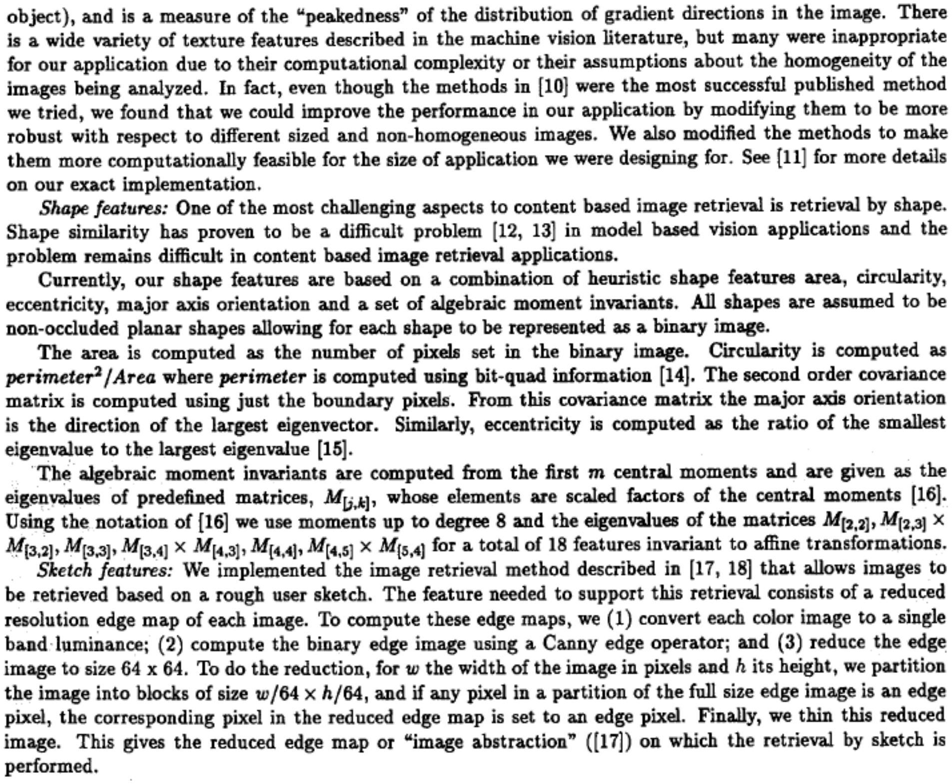object), and is a measure of the "peakedness" of the distribution of gradient directions in the image. There is a wide variety of texture features described in the machine vision literature, but many were inappropriate for our application due to their computational complexity or their assumptions about the homogeneity of the images being analyzed. In fact, even though the methods in [10] were the most successful published method we tried, we found that we could improve the performance in our application by modifying them to be more robust with respect to different sized and non-homogeneous images. We also modified the methods to make them more computationally feasible for the size of application we were designing for. See [11] for more details on our exact implementation. *Shape features:* One of the most challenging aspects to content based image retrieval is retrieval by shape. Shape similarity has proven to be a difficult problem [12, 13] in model based vision applications and the problem remains difficult in content based image retrieval applications. Currently, our shape features are based on a combination of heuristic shape features area, circularity, eccentricity, major axis orientation and a set of algebraic moment invariants. All shapes are assumed to be non-occluded planar shapes allowing for each shape to be represented as a binary image. The area is computed as the number of pixels set in the binary image. Circularity is computed as $perimeter^2/Area$ where $perimeter$ is computed using bit-quad information [14]. The second order covariance matrix is computed using just the boundary pixels. From this covariance matrix the major axis orientation is the direction of the largest eigenvector. Similarly, eccentricity is computed as the ratio of the smallest eigenvalue to the largest eigenvalue [15]. The algebraic moment invariants are computed from the first m central moments and are given as the eigenvalues of predefined matrices, $M_{[j,k]}$, whose elements are scaled factors of the central moments [16]. Using the notation of [16] we use moments up to degree 8 and the eigenvalues of the matrices $M_{[2,2]}, M_{[2,3]} \times M_{[3,2]}, M_{[3,3]}, M_{[3,4]} \times M_{[4,3]}, M_{[4,4]}, M_{[4,5]} \times M_{[5,4]}$ for a total of 18 features invariant to affine transformations. *Sketch features:* We implemented the image retrieval method described in [17, 18] that allows images to be retrieved based on a rough user sketch. The feature needed to support this retrieval consists of a reduced resolution edge map of each image. To compute these edge maps, we (1) convert each color image to a single band luminance; (2) compute the binary edge image using a Canny edge operator; and (3) reduce the edge image to size 64 x 64. To do the reduction, for w the width of the image in pixels and h its height, we partition the image into blocks of size $w/64 \times h/64$, and if any pixel in a partition of the full size edge image is an edge pixel, the corresponding pixel in the reduced edge map is set to an edge pixel. Finally, we thin this reduced image. This gives the reduced edge map or "image abstraction" ([17]) on which the retrieval by sketch is performed. Query by Image and Video Content: The QBIC System (IBM 0000893–902): "To achieve this functionality, QBIC has two main components: database population (the process of creating an image database) and database query. During the population, images and videos are processed to extract features describing their content-colors, textures, shapes, and camera and object motion-and the features are stored in a database. During the query, the user composes a query graphically. Features are generated from the graphical query and then input to a matching engine that finds images or videos from the database with similar features. Figure 2 shows the system architecture." *Id.* at p. 3.

Element	’529 Claim Recitation	Exemplary Citations in Reference
		“For both population and query, the QBIC data model has • still images or scenes (full images) that contain objects (subsets of an image), and • video shots that consist of sets of contiguous frames and contain motion objects. For still images, the QBIC data model distinguishes between ‘scenes’ (or images) and ‘objects.’ A scene is an image or single representative frame of video. An object is a part of a scene-for example, the fox in Figure 3-or a moving entity in a video. For still image database population, features are extracted from images and objects and stored in a database as shown in the top left part of Figure 2. Videos are broken into clips called shots. Representative frames, or r-frames, are generated for each extracted shot. R-frames are treated as still images, and features are extracted and stored in the database. Further processing of shots generates motion objects-for example, a car moving across the screen. Queries are allowed on objects (‘Find images with a red, round object’), scenes (‘Find images that have approximately 30-percent red and IS-percent blue colors’), shots (‘Find all shots panning from left to right’), or any combination (‘Find images that have 30 percent red and contain a blue textured object’).” *Id.* at p. 3. “**DATABASE POPULATION** In still image database population, the images are reduced to a standard-sized icon called a thumbnail and annotated with any available text information. Object identification is an optional but key part of this step. It lets users manually, semiautomatically, or fully automatically identify interesting regions—which we call objects—in the images. Internally, each object is represented as a binary mask. There may be an arbitrary number of objects per image. Objects can overlap and can consist of multiple disconnected components like the set of dots on a polka-dot dress. Text, like ‘baby on beach,’ can be associated with an outlined object or with the scene as a whole. **Object-outlining tools** Ideally, object identification would be automatic, but this is generally difficult. The alternative—manual identification—is tedious and can inhibit query-by-content applications. As a result, we have devoted considerable effort to developing tools to aid in this step. In recent work, we have successfully used fully automatic unsupervised

Element	'529 Claim Recitation	Exemplary Citations in Reference
		segmentation methods along with a foreground/background model to identify objects in a restricted class of images. The images, typical of museums and retail catalogs, have a small number of foreground objects on a generally separable background. Figure 4 shows example results. Even in this domain, robust algorithms are required because of the textured and variegated backgrounds. We also provide semiautomatic tools for identifying objects. One is an enhanced flood-fill technique. Flood-fill methods, found in most photo-editing programs, start from a single object pixel and repeatedly add adjacent pixels whose values are within some given threshold of the original pixel. Selecting the threshold, which must change from image to image and object to object, is tedious. We automatically calculate a dynamic threshold by having the user click on background as well as object points. For reasonably uniform objects that are distinct from the background, this operation allows fast object identification without manually adjusting a threshold. The example in Figure 3 shows an object, a fox, identified by using only a few clicks. We designed another outlining tool to help users track object edges. This tool takes a user-drawn curve and automatically aligns it with nearby image edges. Based on the 'snakes' concept developed in recent computer vision research, the tool finds the curve that maximizes the image gradient magnitude along the curve. The spline snake formulation we use allows for smooth solutions to the resulting nonlinear minimization problem. The computation is done at interactive speeds so that, as the user draws a curve, it is 'rubber-banded' to lie along object boundaries. . . . **Video data** For video data, database population has three major components: • shot detection, • representative frame creation for each shot, and • derivation of a layered representation of coherently moving structures/ objects. . . . Once the shot boundaries have been detected, each shot is represented using an r-frame. R-frames are used for several purposes. First, during database population, r-frames are treated as still images in which objects can be identified by using the previously described methods." *Id.* at pp. 3–6. "INTEGRATION WITH TEXT AND PARAMETRIC ANNOTATION. Query by

Element	'529 Claim Recitation	Exemplary Citations in Reference
		visual content complements and extends existing query methods. Systems must be able to integrate queries combining date, subject matter, price, and availability with content properties such as color, texture, and shape." *Id.* at p. 8. Image, Audio, and Video Extenders Administration and Programming (IBM 000001–617): "Your applications can even search for images by content. Imagine an application that uses visual examples to search for images. With such an application, users could select an example image and have the application find other images that have colors or textures similar to those in the example. With DB2 extenders' Query by Image Content (QBIC) capability, you can create applications that search for images in this visual way." *Id.* at p. 4. " A QBIC catalog is a set of files that hold data about the visual features of images. The Image Extender uses this data to search for images by content. You create a QBIC catalog for each column of images in a user table that you want to make available for searching by content. When you create a QBIC catalog you identify the features for which you want the Image Extender to analyze, store, and later query data. You can also add or drop features from a QBIC catalog after the catalog is created. A QBIC catalog can hold data for the following image features: **Average color** The sum of the color values for all pixels in an image divided by the number of pixels in the image. (A pixel is the smallest element of an image that can be assigned color and intensity.) For example, if 50% of an image consists of blue pixels and the other 50% red pixels, the image has an average color value of purple. Average color is used to search for images that have a predominant color. If an image has a predominant color, the average color will be similar to the predominant color. **Histogram color** Measures the distribution of colors in an image against a spectrum of 64 colors. For each of the 64 colors, histogram color identifies the percentage of pixels in an image that have that color. For example, the histogram color of an image might be 40% white

Element	'529 Claim Recitation	Exemplary Citations in Reference
		pixels, 50% blue, and 10% red; none of the pixels in the image have any of the remaining colors in the histogram spectrum. Histogram color is used to search for images that have a variety of colors. **Positional color** The average color value for the pixels in a specified area in an image. For example, the upper right-hand corner of an image might show a bright yellow sun; the positional color of this area of the image is bright yellow. Positional color is used to search for images that have a predominant color in a particular area. **Texture** Measures the coarseness, contrast, and directionality of an image. Coarseness indicates the size of repeating items in an image (for example, pebbles versus boulders). Contrast identifies the brightness variations in an image (light versus dark). Directionality indicates whether a direction predominates in an image (as in the vertical direction of a picket fence) or does not predominate (as in an image of sand). Texture is used to search for images that have a particular pattern. To make an image available for searching by content, you catalog the image. When you catalog an image, the Image Extender analyzes the image, by computing the feature values for the image, and stores the values in a QBIC catalog. When you search for an image by content, your query identifies one or more features for the search (such as average color), a source for each feature (such as an example image), and a target set of cataloged images. The Image Extender computes the feature value of the source and compares it to the cataloged feature values for the target images. It then computes a score that indicates how similar the feature values of the target images are to the source. You can have the Image Extender return the images whose features are most similar to the source. The Image Extender will return the handle of each image and the image score. You can also have the Image Extender return only the score of a single image." *Id*. at pp. 20–21. QBIC Demo (IBM000747):

Element	'529 Claim Recitation	Exemplary Citations in Reference
		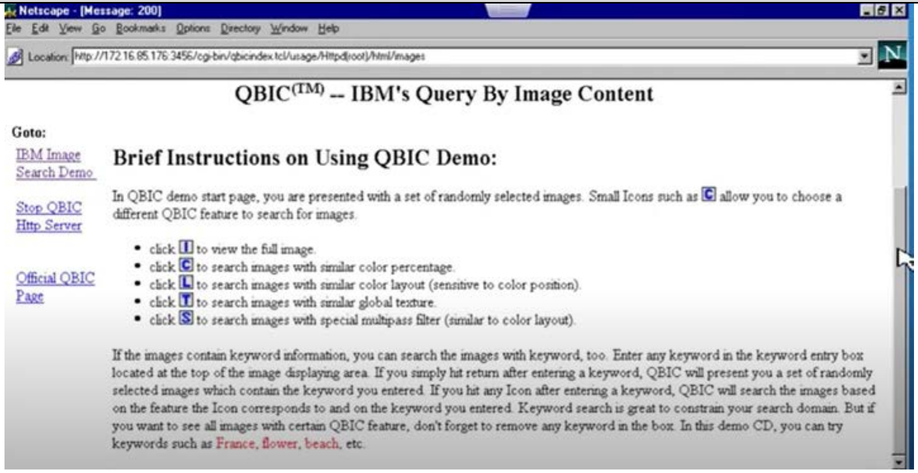 QBIC$^{(TM)}$ -- IBM's Query By Image Content **Brief Instructions on Using QBIC Demo:** In QBIC demo start page, you are presented with a set of randomly selected images. Small Icons such as C allow you to choose a different QBIC feature to search for images. - click I to view the full image. - click C to search images with similar color percentage. - click L to search images with similar color layout (sensitive to color position). - click T to search images with similar global texture. - click S to search images with special multipass filter (similar to color layout). If the images contain keyword information, you can search the images with keyword, too. Enter any keyword in the keyword entry box located at the top of the image displaying area. If you simply hit return after entering a keyword, QBIC will present you a set of randomly selected images which contain the keyword you entered. If you hit any Icon after entering a keyword, QBIC will search the images based on the feature the Icon corresponds to and on the keyword you entered. Keyword search is great to constrain your search domain. But if you want to see all images with certain QBIC feature, don't forget to remove any keyword in the box. In this demo CD, you can try keywords such as France, flower, beach, etc. QBIC Demo (IBM000747): 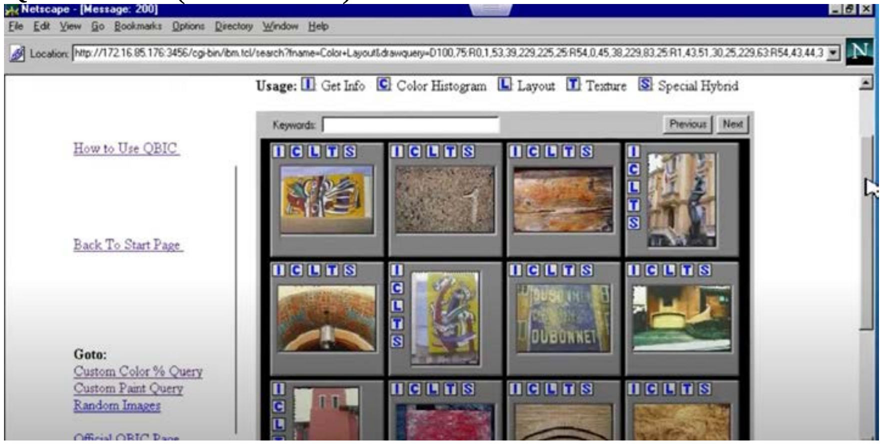 QBIC Demo (IBM000747):

Element	'529 Claim Recitation	Exemplary Citations in Reference
		 Location: http://172.16.85.176:3456/cgi-bin/ibm.tcl/histpicker Red: 59 Green: 131 Blue: 59 9% (255,255,255) 8% (59,131,59) 4% (106,182,106) To specify your color histogram query, "paint" a graphical representation of your pattern of colors and approximate percentages that you want. First select a *percentage* by clicking in the painting area to block off a rectangle of the desired size (percentage of the full area). All new regions are painted white, initially. Then select a color from the QBIC Demo (IBM000747): 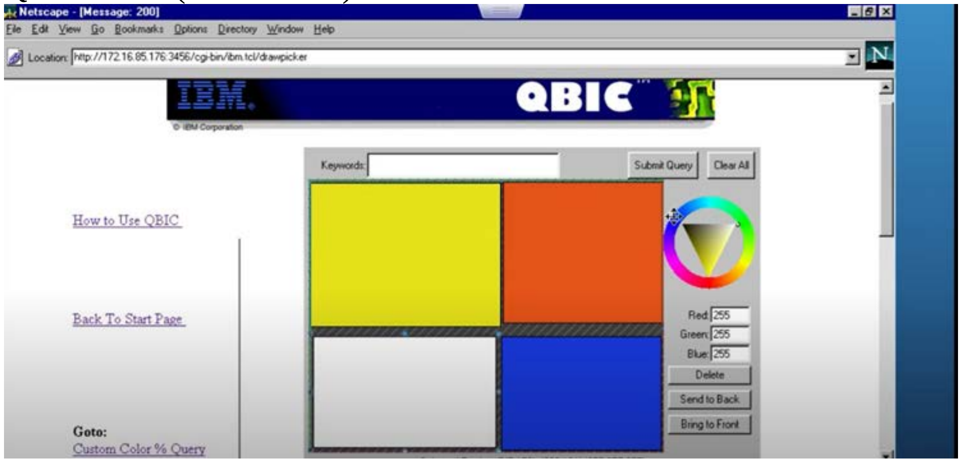 Location: http://172.16.85.176:3456/cgi-bin/ibm.tcl/drawpicker Red: 255 Green: 255 Blue: 255

Element	'529 Claim Recitation	Exemplary Citations in Reference
		QBIC Demo (IBM000747):  http://172.16.85.176:3456/cgi-bin/ibm.tcl/search?fname=HybridColor&france10%2Ejpg.x=0 Goto: Custom Color % Query Custom Paint Query Random Images Official QBIC Page Query was: Example: =france10.jpg Query Type: QbHybridColorFeatureClass Flickner Depo at 25-26:

Element	'529 Claim Recitation	Exemplary Citations in Reference
		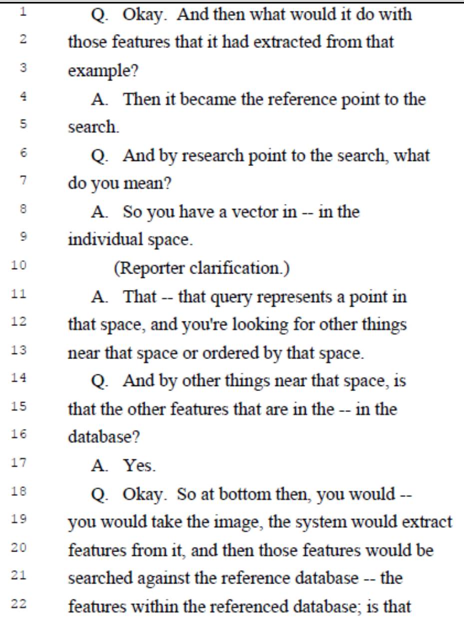 Q. Okay. And then what would it do with those features that it had extracted from that example? A. Then it became the reference point to the search. Q. And by research point to the search, what do you mean? A. So you have a vector in -- in the individual space. (Reporter clarification.) A. That -- that query represents a point in that space, and you're looking for other things near that space or ordered by that space. Q. And by other things near that space, is that the other features that are in the -- in the database? A. Yes. Q. Okay. So at bottom then, you would -- you would take the image, the system would extract features from it, and then those features would be searched against the reference database -- the features within the referenced database; is that 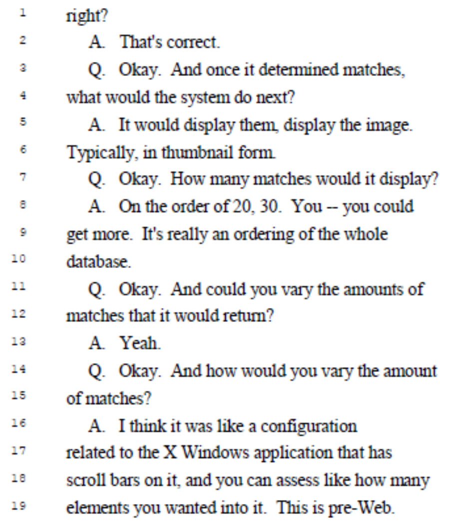 right? A. That's correct. Q. Okay. And once it determined matches, what would the system do next? A. It would display them, display the image. Typically, in thumbnail form. Q. Okay. How many matches would it display? A. On the order of 20, 30. You -- you could get more. It's really an ordering of the whole database. Q. Okay. And could you vary the amounts of matches that it would return? A. Yeah. Q. Okay. And how would you vary the amount of matches? A. I think it was like a configuration related to the X Windows application that has scroll bars on it, and you can assess like how many elements you wanted into it. This is pre-Web.

Element	'529 Claim Recitation	Exemplary Citations in Reference
		Flickner Depo at 62-63: 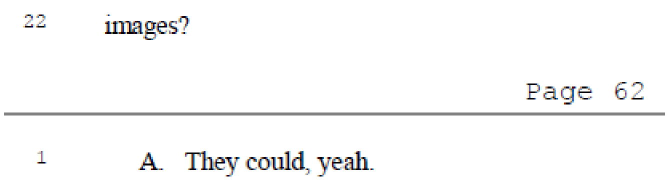 Q. So would the user themselves supply the images? Page 62 A. They could, yeah. Flickner Depo at 78-79:

Element	'529 Claim Recitation	Exemplary Citations in Reference
		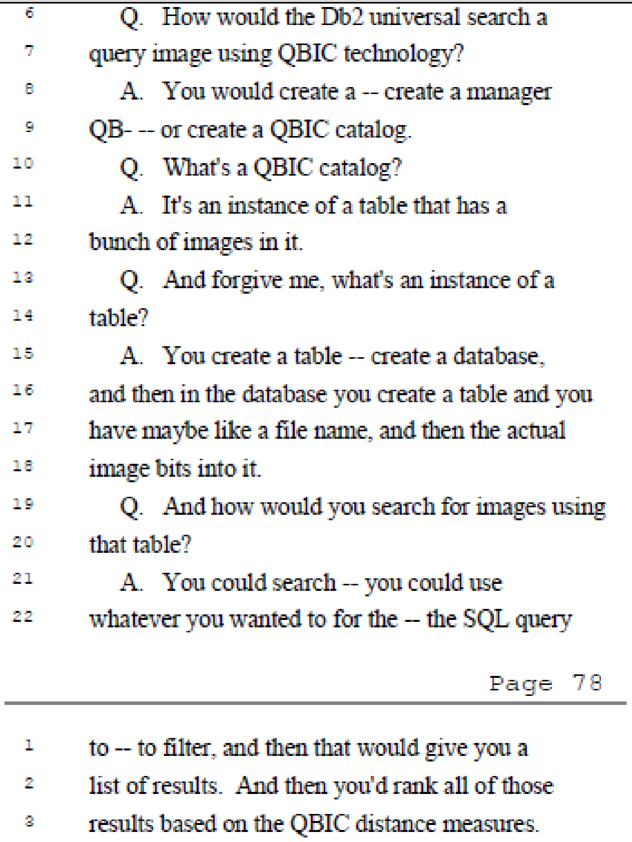 Q. How would the Db2 universal search a query image using QBIC technology? A. You would create a -- create a manager QB- -- or create a QBIC catalog. Q. What's a QBIC catalog? A. It's an instance of a table that has a bunch of images in it. Q. And forgive me, what's an instance of a table? A. You create a table -- create a database, and then in the database you create a table and you have maybe like a file name, and then the actual image bits into it. Q. And how would you search for images using that table? A. You could search -- you could use whatever you wanted to for the -- the SQL query Page 78 to -- to filter, and then that would give you a list of results. And then you'd rank all of those results based on the QBIC distance measures. Flickner Depo at 201:

Element	'529 Claim Recitation	Exemplary Citations in Reference
		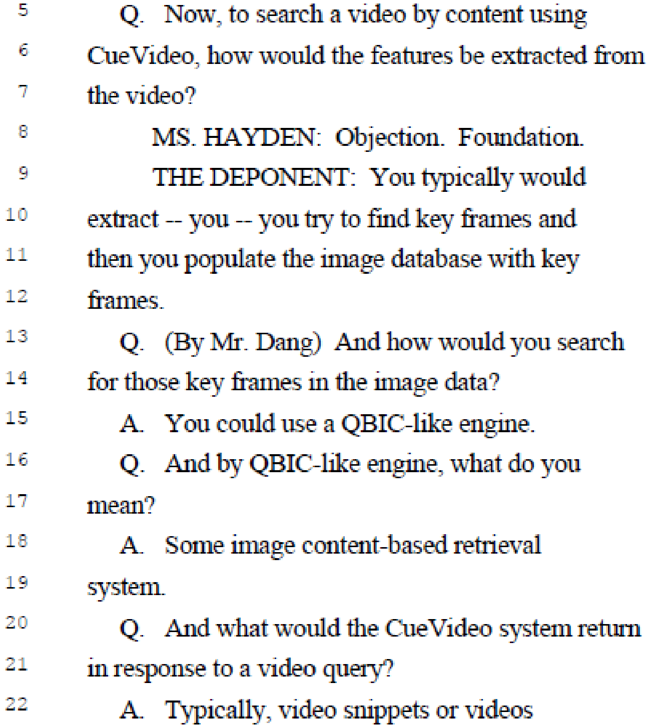 Q. Now, to search a video by content using CueVideo, how would the features be extracted from the video? MS. HAYDEN: Objection. Foundation. THE DEPONENT: You typically would extract -- you -- you try to find key frames and then you populate the image database with key frames. Q. (By Mr. Dang) And how would you search for those key frames in the image data? A. You could use a QBIC-like engine. Q. And by QBIC-like engine, what do you mean? A. Some image content-based retrieval system. Q. And what would the CueVideo system return in response to a video query? A. Typically, video snippets or videos Flickner Depo at 159: 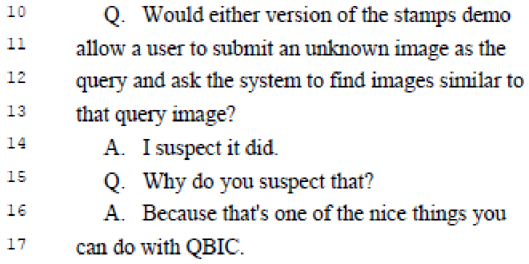 Q. Would either version of the stamps demo allow a user to submit an unknown image as the query and ask the system to find images similar to that query image? A. I suspect it did. Q. Why do you suspect that? A. Because that's one of the nice things you can do with QBIC.

Element	'529 Claim Recitation	Exemplary Citations in Reference
		Flickner Depo at 209-210: 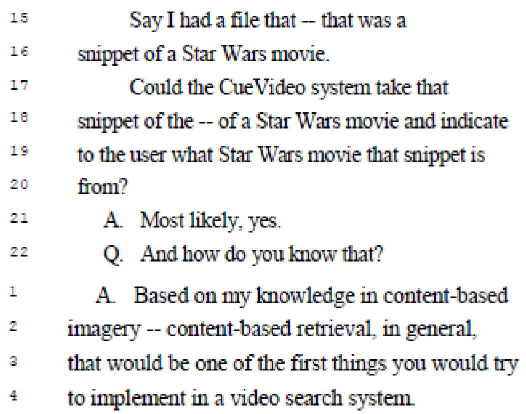 Q. Okay. Let me ask -- let me ask that question using -- using an example. Say I had a file that -- that was a snippet of a Star Wars movie. Could the CueVideo system take that snippet of the -- of a Star Wars movie and indicate to the user what Star Wars movie that snippet is from? A. Most likely, yes. Q. And how do you know that? A. Based on my knowledge in content-based imagery -- content-based retrieval, in general, that would be one of the first things you would try to implement in a video search system. Flickner Depo at 75-76: 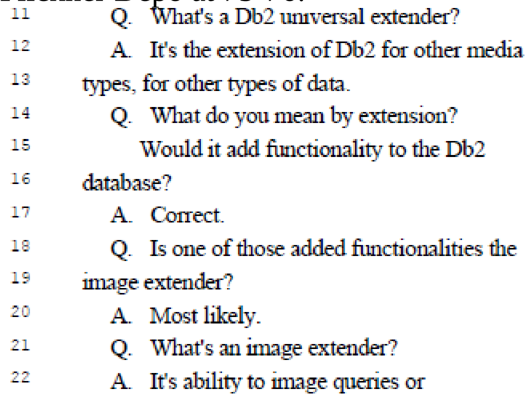 Q. What's a Db2 universal extender? A. It's the extension of Db2 for other media types, for other types of data. Q. What do you mean by extension? Would it add functionality to the Db2 database? A. Correct. Q. Is one of those added functionalities the image extender? A. Most likely. Q. What's an image extender? A. It's ability to image queries or

Element	'529 Claim Recitation	Exemplary Citations in Reference
		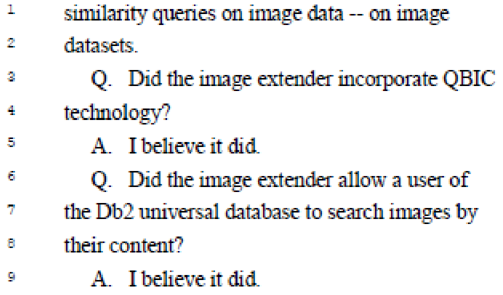 similarity queries on image data -- on image datasets. Q. Did the image extender incorporate QBIC technology? A. I believe it did. Q. Did the image extender allow a user of the Db2 universal database to search images by their content? A. I believe it did. *See also* Flickner Depo at 21, 22-25, 54-55, 89-90. Flickner 8/29/23 Depo at 55–61: 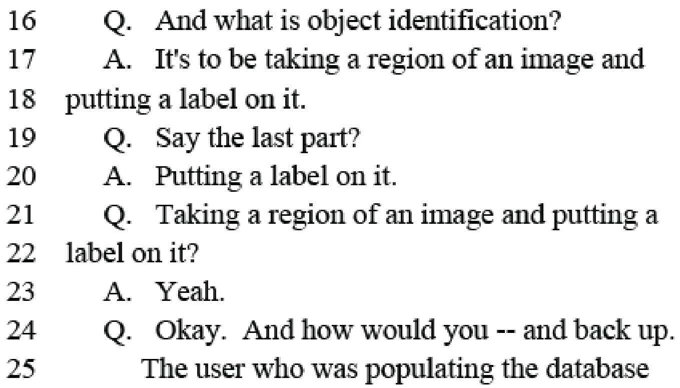 Q. And what is object identification? A. It's to be taking a region of an image and putting a label on it. Q. Say the last part? A. Putting a label on it. Q. Taking a region of an image and putting a label on it? A. Yeah. Q. Okay. And how would you -- and back up. The user who was populating the database

		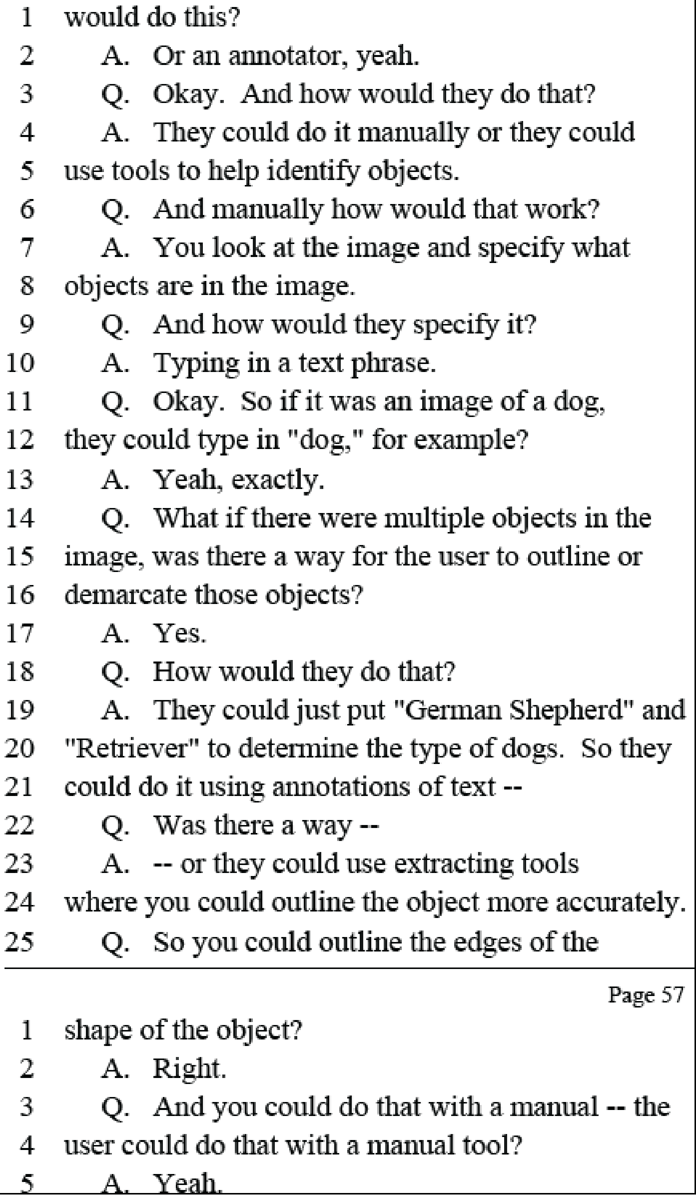 would do this? A. Or an annotator, yeah. Q. Okay. And how would they do that? A. They could do it manually or they could use tools to help identify objects. Q. And manually how would that work? A. You look at the image and specify what objects are in the image. Q. And how would they specify it? A. Typing in a text phrase. Q. Okay. So if it was an image of a dog, they could type in "dog," for example? A. Yeah, exactly. Q. What if there were multiple objects in the image, was there a way for the user to outline or demarcate those objects? A. Yes. Q. How would they do that? A. They could just put "German Shepherd" and "Retriever" to determine the type of dogs. So they could do it using annotations of text -- Q. Was there a way -- A. -- or they could use extracting tools where you could outline the object more accurately. Q. So you could outline the edges of the Page 57 shape of the object? A. Right. Q. And you could do that with a manual -- the user could do that with a manual tool? A. Yeah.	

Element	'529 Claim Recitation	Exemplary Citations in Reference
		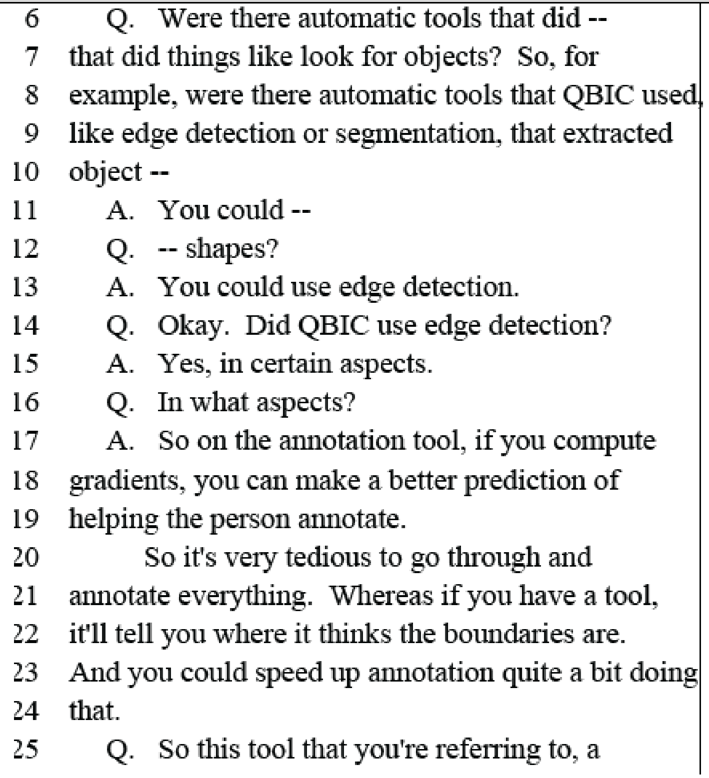 Q. Were there automatic tools that did -- that did things like look for objects? So, for example, were there automatic tools that QBIC used, like edge detection or segmentation, that extracted object -- A. You could -- Q. -- shapes? A. You could use edge detection. Q. Okay. Did QBIC use edge detection? A. Yes, in certain aspects. Q. In what aspects? A. So on the annotation tool, if you compute gradients, you can make a better prediction of helping the person annotate. So it's very tedious to go through and annotate everything. Whereas if you have a tool, it'll tell you where it thinks the boundaries are. And you could speed up annotation quite a bit doing that. Q. So this tool that you're referring to, a

Page 58

user would provide a rough outline of the object and then the tool would essentially conform that outline better to the actual object in the image?
A. That was one way to do it. The other way to do it is to have the next prediction of the outline to be driven by the tool with the user input.
Q. And what was the -- I assume an algorithm did that?
A. Yeah.
Q. Do you know the -- what algorithm was used?
A. We were doing a lot of stuff with splines at the time.
(Reporter seeks clarification.)
A. Splines. S-p-l-i-n-e-s.
Q. Okay. And what are splines?
A. Splines are curves, mathematical representation of curves -- or a mathematical representation of curves.
Q. And so was also any automatic tools, like Canny edge detection or any other edge detection used that didn't require assistance from the user?
A. There were -- yeah, there were some that were more automatic. Some tools were more

Page 59

automatic.
Q. And those were in the form of algorithms?
A. Yeah.
Q. Do you recall any of the algorithms, the names?
A. We did various edge detection stuff.
Q. Okay.
(Reporter seeks clarification.)
A. Just various types. There were things like Sobel operators that we used --
(Reporter seeks clarification.)
A. Sobel.
Q. How do you spell Sobel?
A. S-o-b-e-l, I think.
Q. Okay.
A. That's one form of an edge detector.
Q. Any others that you can recall?
A. This is kind of image processing 101. There were Gabor filters.
(Reporter seeks clarification.)
Q. How do you spell that?
A. G-a-b-o-r.
There was . . . it's been so long. I don't remember what they call it. Very simple filters, but there's a simple convolution filter

Page 60

that was --
(Reporter seeks clarification.)
A. Simple convolution filter that was used for edge detection. Sobel was an advanced version of it.
Q. And what about segmentation, do you know what segmentation is?
A. Yeah.
Q. Okay. Segmentation, as I understand it, essentially separates foreground objects from the background. Is that a high level of explaining it?
A. That's a two-class segmentation, yeah.
Q. Okay. Did you do that type of segmentation to differentiate foreground objects from the background in this database population step?
A. We did research prototypes of it. I don't know what all went out.
Q. You don't know whether those made it into a commercial product?
A. I would say they probably didn't.
Q. Okay. Were they in the -- were they used in the -- in demo applications that you referred to?
A. They might have been.
Q. Okay. And so let's say that objects were

Page 61

identified within a particular image, what were done with those objects? Were they stored, for example, as a separate image or were they somehow linked to or associated with the image from which they were extracted?
A. Typically they were linked.
Q. They were linked to the source image?
A. Yeah.
Q. Okay. And this was done in the database?
A. Could be done in the database.
Q. How was it done?
A. It was probably done in a text database or some sort of key index pair database. So database has a lot of different meanings.
MR. HANSEN: Text database.
BY MR. EDWARDS:
Q. And, Mr. Flickner, you said probably done in a text database, is that what you said?
A. Yeah.
MR. STRAUSSMAN: Make sure to speak up.
THE WITNESS: Okay.
BY MR. EDWARDS:
Q. Or some sort of key index database, is that the other one?
A. Yes.

Element	'529 Claim Recitation	Exemplary Citations in Reference
		Id. at 83 – 85: 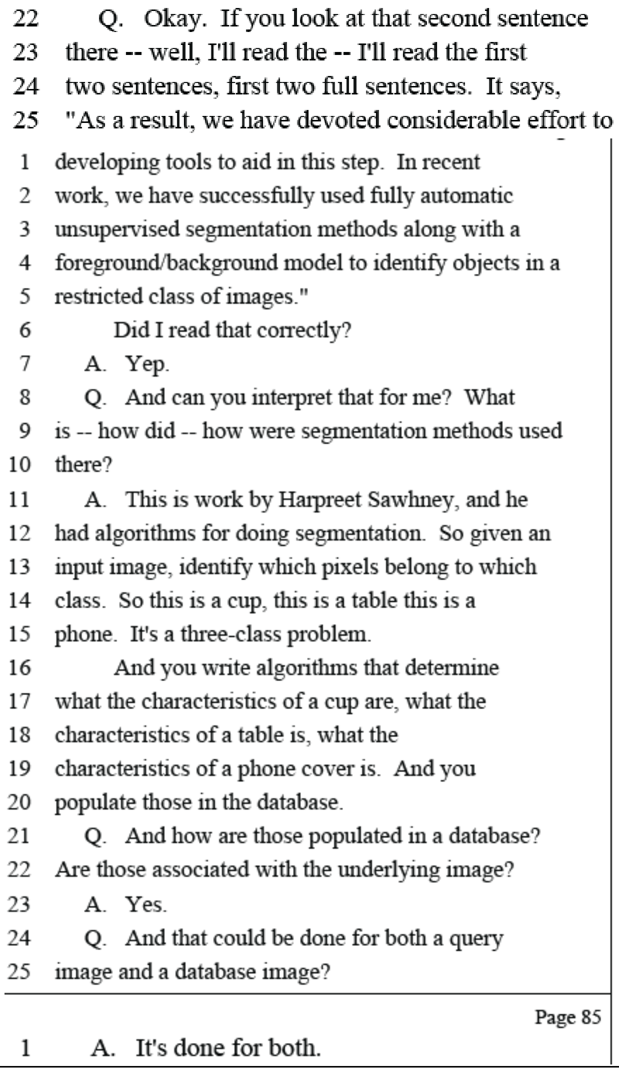 Q. Okay. If you look at that second sentence there -- well, I'll read the -- I'll read the first two sentences, first two full sentences. It says, "As a result, we have devoted considerable effort to developing tools to aid in this step. In recent work, we have successfully used fully automatic unsupervised segmentation methods along with a foreground/background model to identify objects in a restricted class of images." Did I read that correctly? A. Yep. Q. And can you interpret that for me? What is -- how did -- how were segmentation methods used there? A. This is work by Harpreet Sawhney, and he had algorithms for doing segmentation. So given an input image, identify which pixels belong to which class. So this is a cup, this is a table this is a phone. It's a three-class problem. And you write algorithms that determine what the characteristics of a cup are, what the characteristics of a table is, what the characteristics of a phone cover is. And you populate those in the database. Q. And how are those populated in a database? Are those associated with the underlying image? A. Yes. Q. And that could be done for both a query image and a database image? Page 85 A. It's done for both.

Element	'529 Claim Recitation	Exemplary Citations in Reference
		Id. at 305: 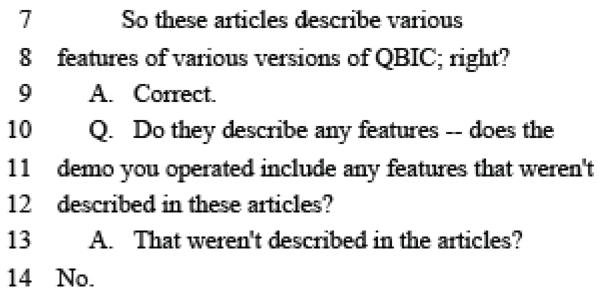 So these articles describe various features of various versions of QBIC; right? A. Correct. Q. Do they describe any features -- does the demo you operated include any features that weren't described in these articles? A. That weren't described in the articles? No. *Id.* at 35: 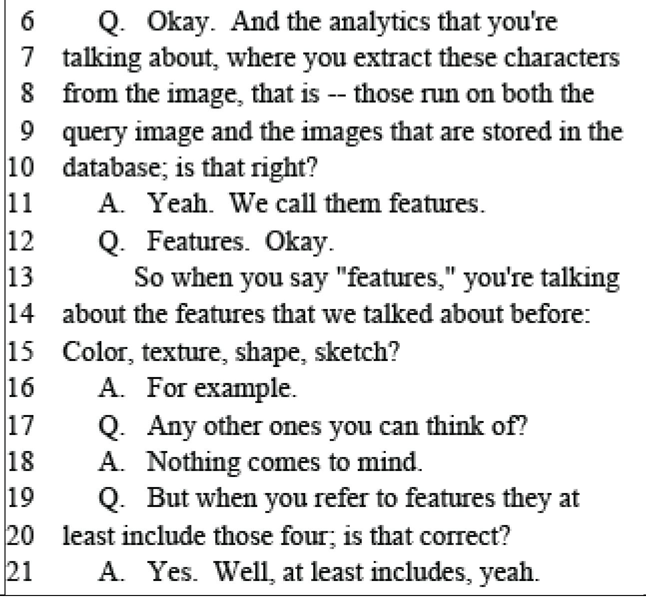 Q. Okay. And the analytics that you're talking about, where you extract these characters from the image, that is -- those run on both the query image and the images that are stored in the database; is that right? A. Yeah. We call them features. Q. Features. Okay. So when you say "features," you're talking about the features that we talked about before: Color, texture, shape, sketch? A. For example. Q. Any other ones you can think of? A. Nothing comes to mind. Q. But when you refer to features they at least include those four; is that correct? A. Yes. Well, at least includes, yeah.

Element	'529 Claim Recitation	Exemplary Citations in Reference
		Id. at 109-110: 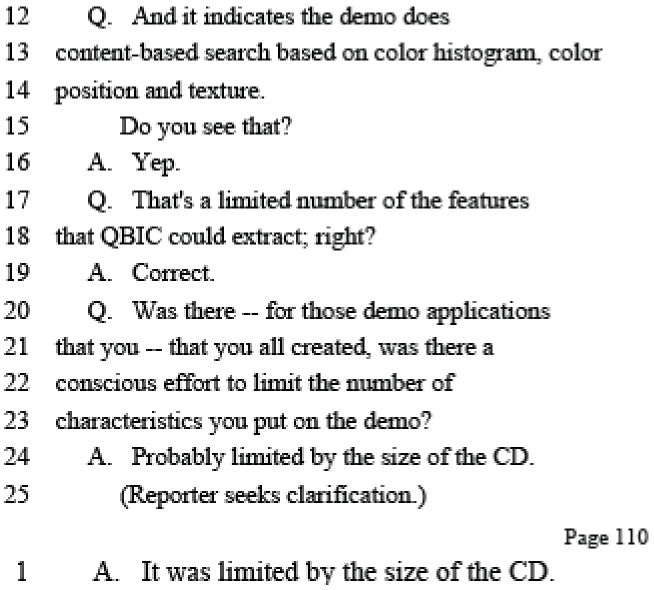 Q. And it indicates the demo does content-based search based on color histogram, color position and texture. Do you see that? A. Yep. Q. That's a limited number of the features that QBIC could extract; right? A. Correct. Q. Was there -- for those demo applications that you -- that you all created, was there a conscious effort to limit the number of characteristics you put on the demo? A. Probably limited by the size of the CD. (Reporter seeks clarification.) Page 110 A. It was limited by the size of the CD. *Id.* at 311: 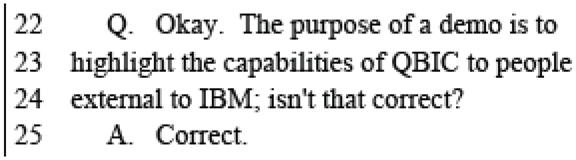 Q. Okay. The purpose of a demo is to highlight the capabilities of QBIC to people external to IBM; isn't that correct? A. Correct. *Id.* at 312-313:

Element	'529 Claim Recitation	Exemplary Citations in Reference
		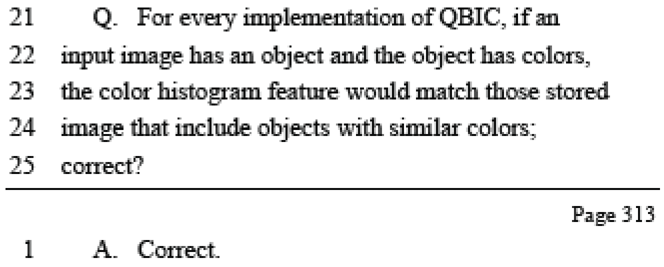 Q. For every implementation of QBIC, if an input image has an object and the object has colors, the color histogram feature would match those stored image that include objects with similar colors; correct? Page 313 A. Correct. *Id.* at 313-317:

Element	'529 Claim Recitation	Exemplary Citations in Reference
		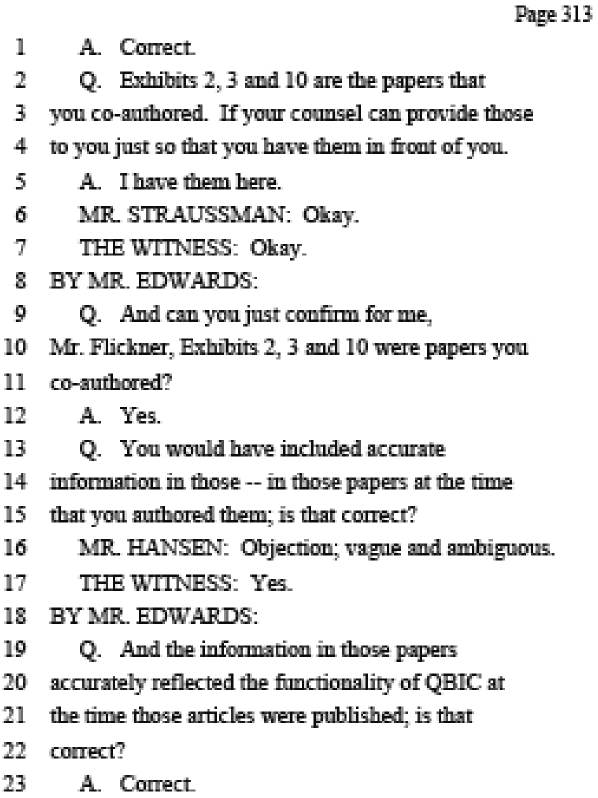 Page 313 A. Correct. Q. Exhibits 2, 3 and 10 are the papers that you co-authored. If your counsel can provide those to you just so that you have them in front of you. A. I have them here. MR. STRAUSSMAN: Okay. THE WITNESS: Okay. BY MR. EDWARDS: Q. And can you just confirm for me, Mr. Flickner, Exhibits 2, 3 and 10 were papers you co-authored? A. Yes. Q. You would have included accurate information in those -- in those papers at the time that you authored them; is that correct? MR. HANSEN: Objection; vague and ambiguous. THE WITNESS: Yes. BY MR. EDWARDS: Q. And the information in those papers accurately reflected the functionality of QBIC at the time those articles were published; is that correct? A. Correct.

		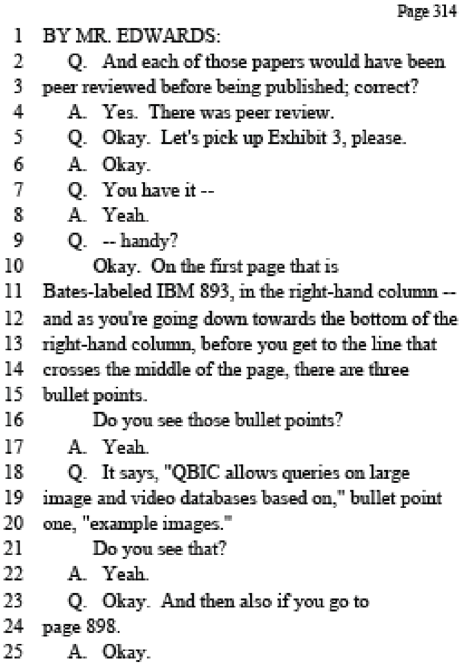 Page 314 BY MR. EDWARDS: Q. And each of those papers would have been peer reviewed before being published; correct? A. Yes. There was peer review. Q. Okay. Let's pick up Exhibit 3, please. A. Okay. Q. You have it -- A. Yeah. Q. -- handy? Okay. On the first page that is Bates-labeled IBM 893, in the right-hand column -- and as you're going down towards the bottom of the right-hand column, before you get to the line that crosses the middle of the page, there are three bullet points. Do you see those bullet points? A. Yeah. Q. It says, "QBIC allows queries on large image and video databases based on," bullet point one, "example images." Do you see that? A. Yeah. Q. Okay. And then also if you go to page 898. A. Okay. 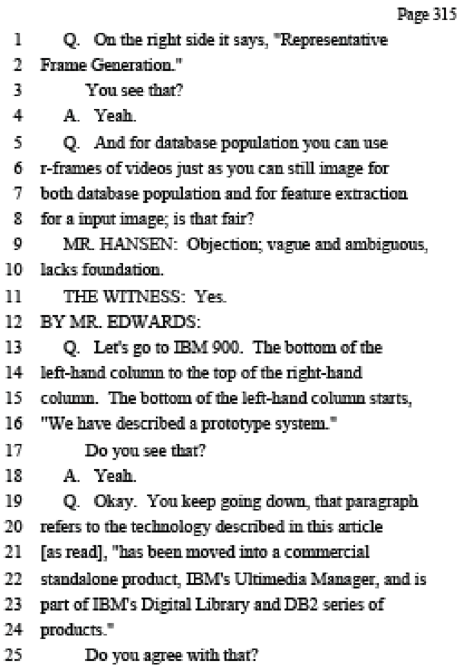 Page 315 Q. On the right side it says, "Representative Frame Generation." You see that? A. Yeah. Q. And for database population you can use r-frames of videos just as you can still image for both database population and for feature extraction for a input image; is that fair? MR. HANSEN: Objection; vague and ambiguous, lacks foundation. THE WITNESS: Yes. BY MR. EDWARDS: Q. Let's go to IBM 900. The bottom of the left-hand column to the top of the right-hand column. The bottom of the left-hand column starts, "We have described a prototype system." Do you see that? A. Yeah. Q. Okay. You keep going down, that paragraph refers to the technology described in this article [as read], "has been moved into a commercial standalone product, IBM's Ultimedia Manager, and is part of IBM's Digital Library and DB2 series of products." Do you agree with that?

Element	'529 Claim Recitation	Exemplary Citations in Reference
		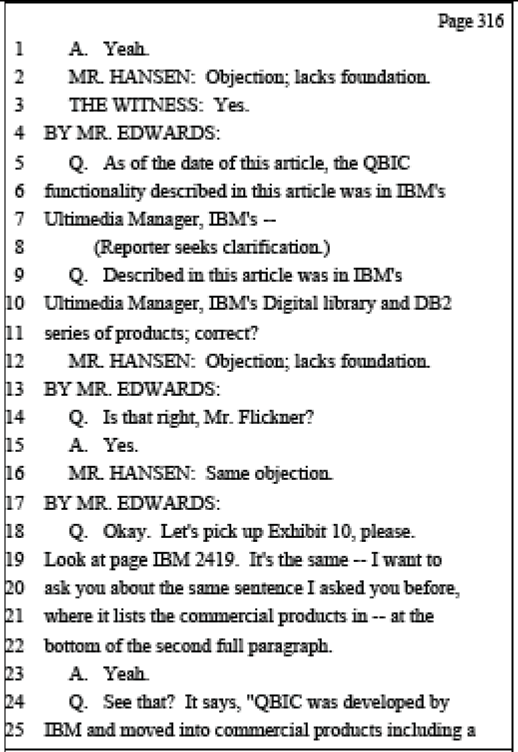 Page 316 A. Yeah. MR. HANSEN: Objection; lacks foundation. THE WITNESS: Yes. BY MR. EDWARDS: Q. As of the date of this article, the QBIC functionality described in this article was in IBM's Ultimedia Manager, IBM's -- (Reporter seeks clarification.) Q. Described in this article was in IBM's Ultimedia Manager, IBM's Digital library and DB2 series of products; correct? MR. HANSEN: Objection; lacks foundation. BY MR. EDWARDS: Q. Is that right, Mr. Flickner? A. Yes. MR. HANSEN: Same objection. BY MR. EDWARDS: Q. Okay. Let's pick up Exhibit 10, please. Look at page IBM 2419. It's the same -- I want to ask you about the same sentence I asked you before, where it lists the commercial products in -- at the bottom of the second full paragraph. A. Yeah. Q. See that? It says, "QBIC was developed by IBM and moved into commercial products including a 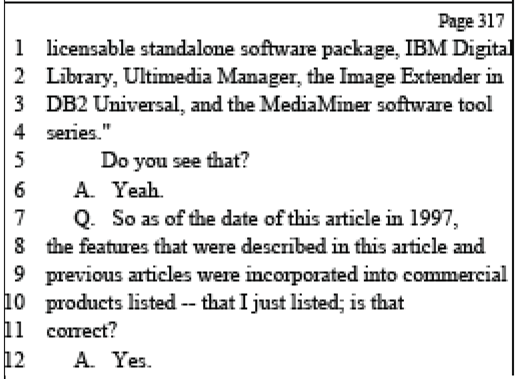 Page 317 licensable standalone software package, IBM Digital Library, Ultimedia Manager, the Image Extender in DB2 Universal, and the MediaMiner software tool series." Do you see that? A. Yeah. Q. So as of the date of this article in 1997, the features that were described in this article and previous articles were incorporated into commercial products listed -- that I just listed; is that correct? A. Yes.

Element	'529 Claim Recitation	Exemplary Citations in Reference
		To the extent it is found that QBIC does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that QBIC does not anticipate this claim, QBIC renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine QBIC with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1D]	distinguish an object present in the image from others using a database that stores data characteristics of target objects;	QBIC discloses distinguish an object present in the image from others using a database that stores data characteristics of target objects. The QBIC Project: Querying Images by Content Using Color, Texture, and Shape (IBM 0002390–404) at 2390 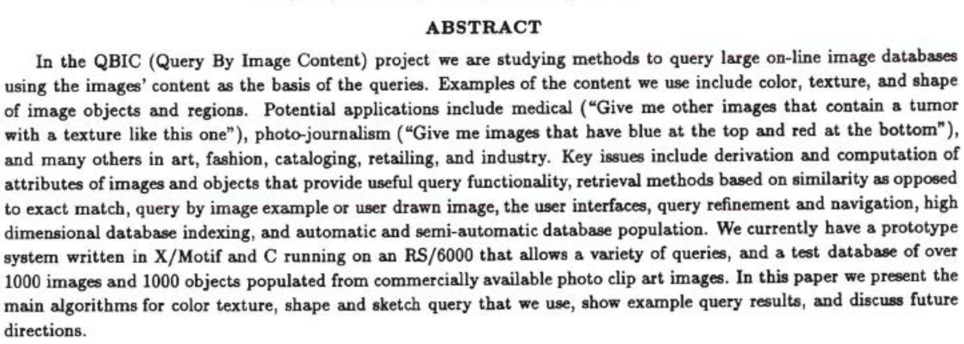 ABSTRACT In the QBIC (Query By Image Content) project we are studying methods to query large on-line image databases using the images' content as the basis of the queries. Examples of the content we use include color, texture, and shape of image objects and regions. Potential applications include medical ("Give me other images that contain a tumor with a texture like this one"), photo-journalism ("Give me images that have blue at the top and red at the bottom"), and many others in art, fashion, cataloging, retailing, and industry. Key issues include derivation and computation of attributes of images and objects that provide useful query functionality, retrieval methods based on similarity as opposed to exact match, query by image example or user drawn image, the user interfaces, query refinement and navigation, high dimensional database indexing, and automatic and semi-automatic database population. We currently have a prototype system written in X/Motif and C running on an RS/6000 that allows a variety of queries, and a test database of over 1000 images and 1000 objects populated from commercially available photo clip art images. In this paper we present the main algorithms for color texture, shape and sketch query that we use, show example query results, and discuss future directions. IBM0002390 at 2391

Element	'529 Claim Recitation	Exemplary Citations in Reference
		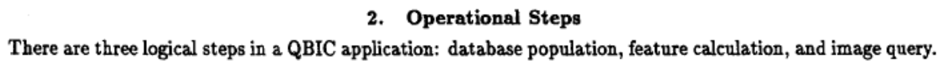 **2. Operational Steps** There are three logical steps in a QBIC application: database population, feature calculation, and image query. IBM0002390 at 2391-92 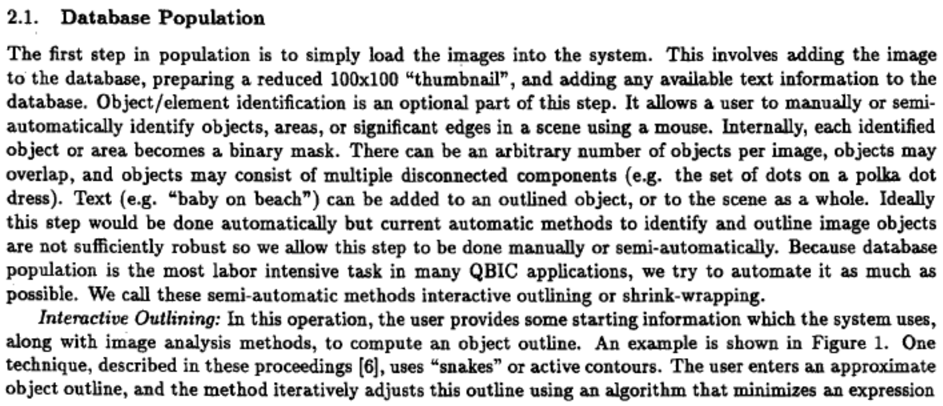 **2.1. Database Population** The first step in population is to simply load the images into the system. This involves adding the image to the database, preparing a reduced 100x100 "thumbnail", and adding any available text information to the database. Object/element identification is an optional part of this step. It allows a user to manually or semi-automatically identify objects, areas, or significant edges in a scene using a mouse. Internally, each identified object or area becomes a binary mask. There can be an arbitrary number of objects per image, objects may overlap, and objects may consist of multiple disconnected components (e.g. the set of dots on a polka dot dress). Text (e.g. "baby on beach") can be added to an outlined object, or to the scene as a whole. Ideally this step would be done automatically but current automatic methods to identify and outline image objects are not sufficiently robust so we allow this step to be done manually or semi-automatically. Because database population is the most labor intensive task in many QBIC applications, we try to automate it as much as possible. We call these semi-automatic methods interactive outlining or shrink-wrapping. *Interactive Outlining:* In this operation, the user provides some starting information which the system uses, along with image analysis methods, to compute an object outline. An example is shown in Figure 1. One technique, described in these proceedings [6], uses "snakes" or active contours. The user enters an approximate object outline, and the method iteratively adjusts this outline using an algorithm that minimizes an expression

Element	'529 Claim Recitation	Exemplary Citations in Reference
		containing two main terms: total image edge strength along the outline, and the outline curvature. Another method, flood fill, can take as starting information a single mouse-clicked image point. Starting from this object point, all adjacent pixels are included in the object whose color is within a specified color distance from the original point. The perimeter of all included points becomes the object boundary. For reasonably uniform objects that are distinct from the background, this operation provides a fast object outliner. In interactive outlining, as in other parts of our system, we have attempted to build the functions in such a way as to carefully distribute the tasks between the human and machine. For example, the tasks of giving the starting outline for an object or of attaching semantic meaning to objects is left to the user since they are difficult for machines. On the other hand, computation of quantitative image/object features is ideally suited for machines and so is done automatically.  Figure 1: Original image, approximate user outline, and automatically refined outline of sun. IBM0002390 at 2392-93

Element	’529 Claim Recitation	Exemplary Citations in Reference
		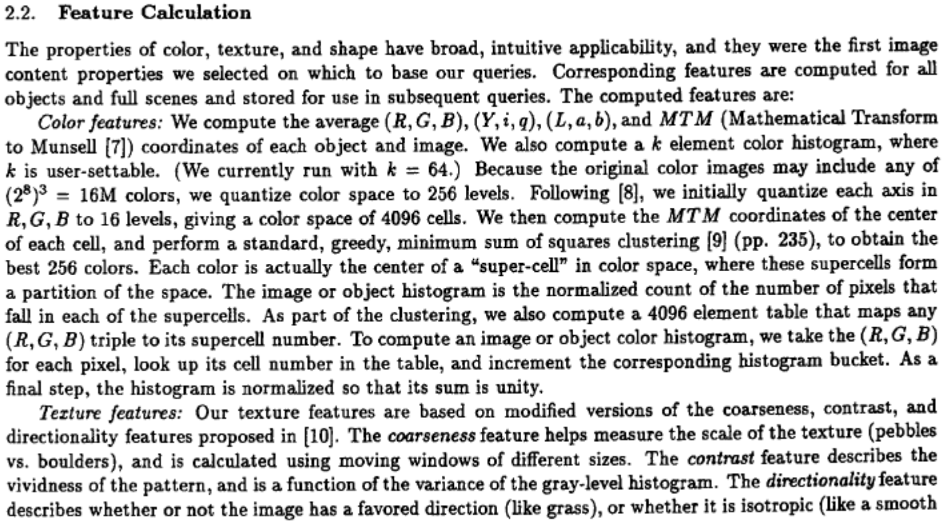 2.2. **Feature Calculation** The properties of color, texture, and shape have broad, intuitive applicability, and they were the first image content properties we selected on which to base our queries. Corresponding features are computed for all objects and full scenes and stored for use in subsequent queries. The computed features are: *Color features:* We compute the average (R, G, B), (Y, i, q), (L, a, b), and MTM (Mathematical Transform to Munsell [7]) coordinates of each object and image. We also compute a k element color histogram, where k is user-settable. (We currently run with $k = 64$.) Because the original color images may include any of $(2^8)^3 = 16$M colors, we quantize color space to 256 levels. Following [8], we initially quantize each axis in R, G, B to 16 levels, giving a color space of 4096 cells. We then compute the MTM coordinates of the center of each cell, and perform a standard, greedy, minimum sum of squares clustering [9] (pp. 235), to obtain the best 256 colors. Each color is actually the center of a "super-cell" in color space, where these supercells form a partition of the space. The image or object histogram is the normalized count of the number of pixels that fall in each of the supercells. As part of the clustering, we also compute a 4096 element table that maps any (R, G, B) triple to its supercell number. To compute an image or object color histogram, we take the (R, G, B) for each pixel, look up its cell number in the table, and increment the corresponding histogram bucket. As a final step, the histogram is normalized so that its sum is unity. *Texture features:* Our texture features are based on modified versions of the coarseness, contrast, and directionality features proposed in [10]. The *coarseness* feature helps measure the scale of the texture (pebbles vs. boulders), and is calculated using moving windows of different sizes. The *contrast* feature describes the vividness of the pattern, and is a function of the variance of the gray-level histogram. The *directionality* feature describes whether or not the image has a favored direction (like grass), or whether it is isotropic (like a smooth

Element	'529 Claim Recitation	Exemplary Citations in Reference
		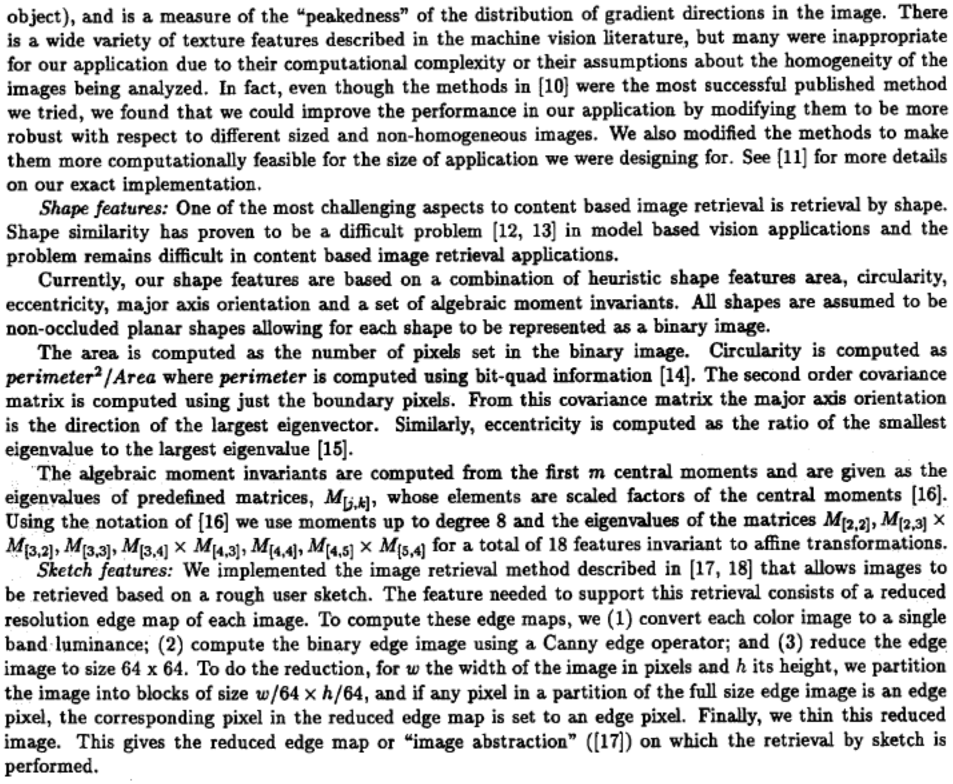 object), and is a measure of the "peakedness" of the distribution of gradient directions in the image. There is a wide variety of texture features described in the machine vision literature, but many were inappropriate for our application due to their computational complexity or their assumptions about the homogeneity of the images being analyzed. In fact, even though the methods in [10] were the most successful published method we tried, we found that we could improve the performance in our application by modifying them to be more robust with respect to different sized and non-homogeneous images. We also modified the methods to make them more computationally feasible for the size of application we were designing for. See [11] for more details on our exact implementation. *Shape features:* One of the most challenging aspects to content based image retrieval is retrieval by shape. Shape similarity has proven to be a difficult problem [12, 13] in model based vision applications and the problem remains difficult in content based image retrieval applications. Currently, our shape features are based on a combination of heuristic shape features area, circularity, eccentricity, major axis orientation and a set of algebraic moment invariants. All shapes are assumed to be non-occluded planar shapes allowing for each shape to be represented as a binary image. The area is computed as the number of pixels set in the binary image. Circularity is computed as $perimeter^2/Area$ where $perimeter$ is computed using bit-quad information [14]. The second order covariance matrix is computed using just the boundary pixels. From this covariance matrix the major axis orientation is the direction of the largest eigenvector. Similarly, eccentricity is computed as the ratio of the smallest eigenvalue to the largest eigenvalue [15]. The algebraic moment invariants are computed from the first m central moments and are given as the eigenvalues of predefined matrices, $M_{[j,k]}$, whose elements are scaled factors of the central moments [16]. Using the notation of [16] we use moments up to degree 8 and the eigenvalues of the matrices $M_{[2,2]}, M_{[2,3]} \times M_{[3,2]}, M_{[3,3]}, M_{[3,4]} \times M_{[4,3]}, M_{[4,4]}, M_{[4,5]} \times M_{[5,4]}$ for a total of 18 features invariant to affine transformations. *Sketch features:* We implemented the image retrieval method described in [17, 18] that allows images to be retrieved based on a rough user sketch. The feature needed to support this retrieval consists of a reduced resolution edge map of each image. To compute these edge maps, we (1) convert each color image to a single band luminance; (2) compute the binary edge image using a Canny edge operator; and (3) reduce the edge image to size 64 x 64. To do the reduction, for w the width of the image in pixels and h its height, we partition the image into blocks of size $w/64 \times h/64$, and if any pixel in a partition of the full size edge image is an edge pixel, the corresponding pixel in the reduced edge map is set to an edge pixel. Finally, we thin this reduced image. This gives the reduced edge map or "image abstraction" ([17]) on which the retrieval by sketch is performed. IBM0002390 at 2393-95:

Element	'529 Claim Recitation	Exemplary Citations in Reference
		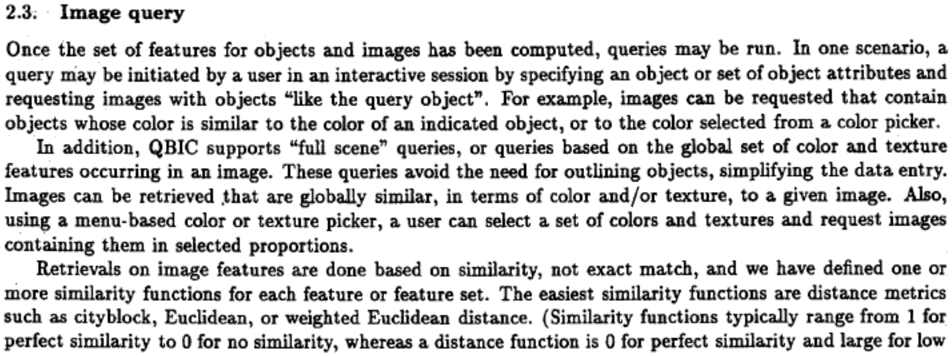 **2.3. Image query** Once the set of features for objects and images has been computed, queries may be run. In one scenario, a query may be initiated by a user in an interactive session by specifying an object or set of object attributes and requesting images with objects "like the query object". For example, images can be requested that contain objects whose color is similar to the color of an indicated object, or to the color selected from a color picker. In addition, QBIC supports "full scene" queries, or queries based on the global set of color and texture features occurring in an image. These queries avoid the need for outlining objects, simplifying the data entry. Images can be retrieved that are globally similar, in terms of color and/or texture, to a given image. Also, using a menu-based color or texture picker, a user can select a set of colors and textures and request images containing them in selected proportions. Retrievals on image features are done based on similarity, not exact match, and we have defined one or more similarity functions for each feature or feature set. The easiest similarity functions are distance metrics such as cityblock, Euclidean, or weighted Euclidean distance. (Similarity functions typically range from 1 for perfect similarity to 0 for no similarity, whereas a distance function is 0 for perfect similarity and large for low

Element	'529 Claim Recitation	Exemplary Citations in Reference
		similarity, so we must convert as necessary to obtain consistent measures.) Different similarity measures may be used, and they may vary from application to application. The matching methods and similarity functions for the different feature types are described below: *Color:* For average color, the distance between a query object and database object is weighted Euclidean distance. The weights are user adjustable so that we may, for example, downweight the L component of an (L, a, b) query. Best results, as judged by our subjective evaluation in on-going experiments, are obtained by representing the average color in the MTM color space, and by inversely weighting each component by its standard deviation over the samples in the database. Improved results are obtained when color matching is done not on average color, but on the distribution of colors that occur in an object or image. In this way, a red and blue object matches other red and blue objects better than to a purple object. To match color histograms, we use a method similar to the one described in [8]. Let X be the query histogram and Y the histogram of an item in the database, both normalized. We compute the element by element difference histogram Z. Then the similarity between X and Y is given by $\|Z\| = Z^T A Z$, where A is a symmetric color similarity matrix with $$a(i,j) \equiv 1 - d(c_i, c_j)/d_{max}$$ where c_i and c_j are the ith and jth colors in the color histograms, and $d(c_i, c_j)$ is the MTM color distance, and d_{max} is the maximum distance between any two colors ([8] used Luv distance.) This metric gives the (weighted) length of the difference vector between X and Y, weighted by A, which accounts for the perceptual distance between different pairs of colors. The result is a measure that accounts for both the difference between the amounts in any given color (e.g. a particular shade of red), as well as differences between similar colors (e.g. red and orange). Based on our experience, we also found it useful for the user to be able to request images with x% of color 1, y% of color 2, etc. where both the percentages and the colors are explicitly specified by the user. This allows searches such as for a beach scene (25% white (for the sand) and 50% blue (for the sky and water)), and for many scenes for which the user has seen before, and for which he/she can remember approximate color percentages. The color picker that allows this type of query is shown in Figure 2. 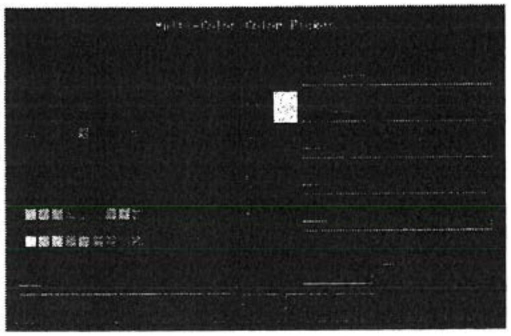 Figure 2: Multi-Color Color Picker with two selected colors. The palette on the left displays the set of selectable colors. Once a color is selected, it appears in a box on the right, and the amount of the color can be specified by the sliders. In this case, 25% light blue and 15% white have been selected. *Texture:* Texture distance is computed as weighted Euclidean distance in the three dimensional texture space. The most common weighting factors are the inverse variances for each component, computed over the

Element	'529 Claim Recitation	Exemplary Citations in Reference
		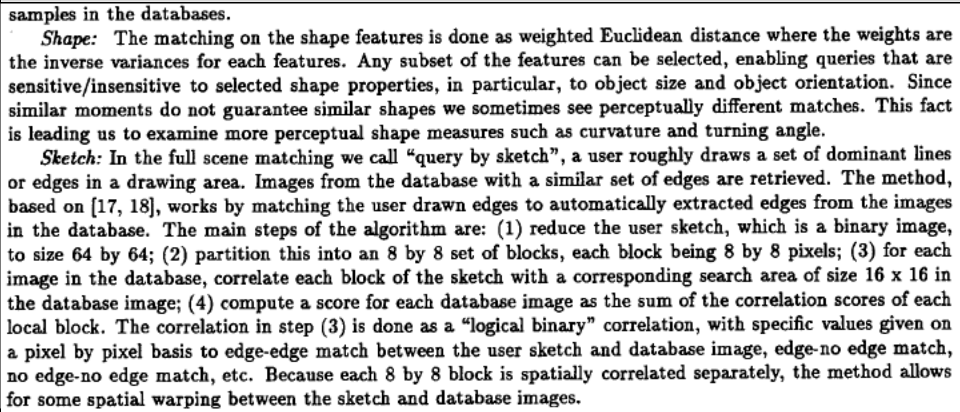 samples in the databases. *Shape:* The matching on the shape features is done as weighted Euclidean distance where the weights are the inverse variances for each features. Any subset of the features can be selected, enabling queries that are sensitive/insensitive to selected shape properties, in particular, to object size and object orientation. Since similar moments do not guarantee similar shapes we sometimes see perceptually different matches. This fact is leading us to examine more perceptual shape measures such as curvature and turning angle. *Sketch:* In the full scene matching we call "query by sketch", a user roughly draws a set of dominant lines or edges in a drawing area. Images from the database with a similar set of edges are retrieved. The method, based on [17, 18], works by matching the user drawn edges to automatically extracted edges from the images in the database. The main steps of the algorithm are: (1) reduce the user sketch, which is a binary image, to size 64 by 64; (2) partition this into an 8 by 8 set of blocks, each block being 8 by 8 pixels; (3) for each image in the database, correlate each block of the sketch with a corresponding search area of size 16 x 16 in the database image; (4) compute a score for each database image as the sum of the correlation scores of each local block. The correlation in step (3) is done as a "logical binary" correlation, with specific values given on a pixel by pixel basis to edge-edge match between the user sketch and database image, edge-no edge match, no edge-no edge match, etc. Because each 8 by 8 block is spatially correlated separately, the method allows for some spatial warping between the sketch and database images. IBM0002390 at 2395 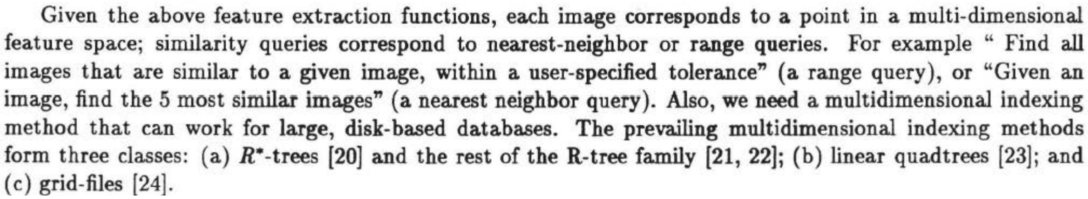 Given the above feature extraction functions, each image corresponds to a point in a multi-dimensional feature space; similarity queries correspond to nearest-neighbor or range queries. For example " Find all images that are similar to a given image, within a user-specified tolerance" (a range query), or "Given an image, find the 5 most similar images" (a nearest neighbor query). Also, we need a multidimensional indexing method that can work for large, disk-based databases. The prevailing multidimensional indexing methods form three classes: (a) R^*-trees [20] and the rest of the R-tree family [21, 22]; (b) linear quadtrees [23]; and (c) grid-files [24]. QBIC (Query by Image Content) at 1

Element	'529 Claim Recitation	Exemplary Citations in Reference
		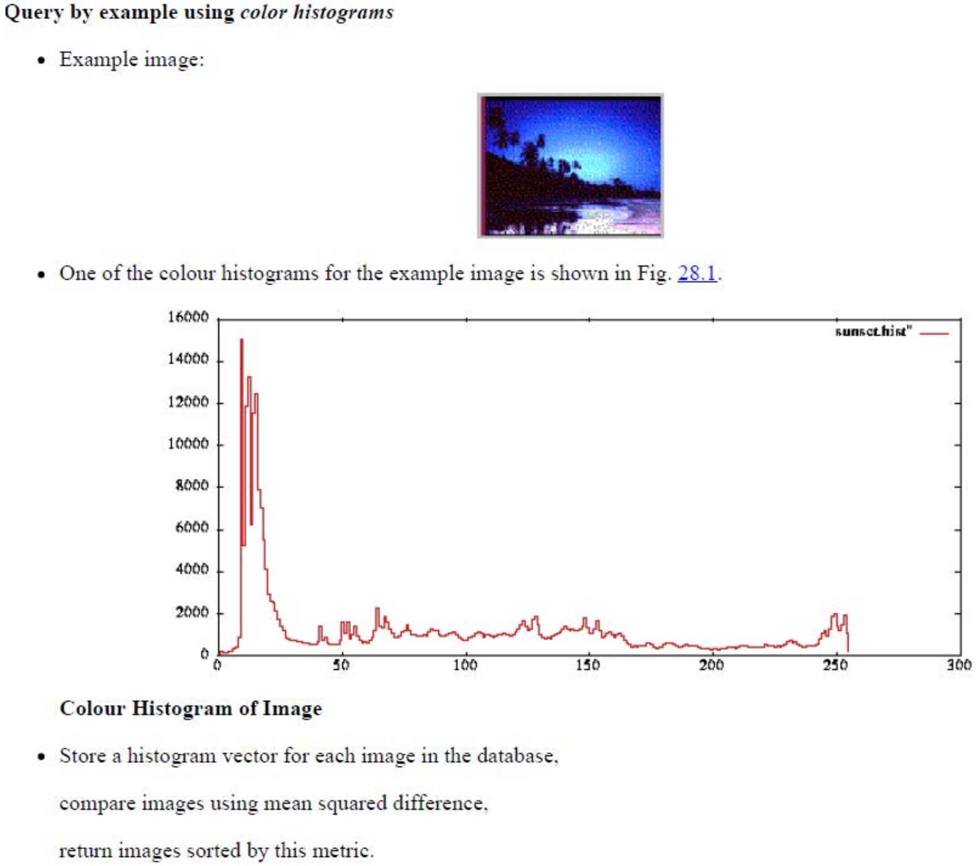 **Query by example using *color histograms*** - Example image: - One of the colour histograms for the example image is shown in Fig. 28.1. **Colour Histogram of Image** - Store a histogram vector for each image in the database, compare images using mean squared difference, return images sorted by this metric. Query by Image and Video Content: The QBIC System (IBM 0000893–902): "To achieve this functionality, QBIC has two main components: database population (the process of creating an image database) and database query. During the population, images and videos are processed to extract features describing their content-colors, textures, shapes, and camera and object motion-and the features are stored in a database. During the query, the user composes a query graphically. Features are generated from the graphical query and then input to a matching engine that finds images or videos from

Element	'529 Claim Recitation	Exemplary Citations in Reference
		the database with similar features. Figure 2 shows the system architecture." *Id.* at p. 3. "For both population and query, the QBIC data model has • still images or scenes (full images) that contain objects (subsets of an image), and • video shots that consist of sets of contiguous frames and contain motion objects. For still images, the QBIC data model distinguishes between 'scenes' (or images) and 'objects.' A scene is an image or single representative frame of video. An object is a part of a scene-for example, the fox in Figure 3-or a moving entity in a video. For still image database population, features are extracted from images and objects and stored in a database as shown in the top left part of Figure 2. Videos are broken into clips called shots. Representative frames, or r-frames, are generated for each extracted shot. R-frames are treated as still images, and features are extracted and stored in the database. Further processing of shots generates motion objects-for example, a car moving across the screen. Queries are allowed on objects ('Find images with a red, round object'), scenes ('Find images that have approximately 30-percent red and IS-percent blue colors'), shots ('Find all shots panning from left to right'), or any combination ('Find images that have 30 percent red and contain a blue textured object')." *Id.* at p. 3. "**SAMPLE QUERIES** For each full-scene image, identified image object, r-frame, and identified video object resulting from the above processing, a set of features is computed to allow content-based queries. The features are computed and stored during database population. We present a brief description of the features and the associated queries. Mathematical details on the features and matching methods can be found in Ashley et al. and Niblack et al. Average color queries let users find images or objects that are similar to a selected color, say from a color wheel, or to the color of an object. The feature used in the query is a 3D vector of Munsell color coordinates. Histogram color queries return items with matching color distributions—say, a fabric pattern with approximately 40 percent red and 20 percent blue. For this case, the underlying feature is a 256-element histogram computed over a quantized version of the color space.

Element	'529 Claim Recitation	Exemplary Citations in Reference
		Figure 7 shows a histogram query on still images and a color query on video r-frames. Note that in the query specification for the histogram query of Figure 7, the user has selected percentages of two colors (blue and white) by adjusting sliders. Using such a query, an advertising agent could, for example, search for a picture of a beach scene, one predominantly blue (for sky and water) and white (for sand and clouds); or find images with similar color spreads for a uniform ad campaign. The average color query demonstrates a query against a video shot database where the user is searching for red r-frames. Again, the query specification is on the left and the best hits are on the right. Figure 8 shows an example texture query. In this case, the query is specified by selecting from a sampler—a set of prestored example images. The underlying texture features are mathematical representations of coarseness, contrast, and directionality features. Coarseness measures the scale of a texture (pebbles versus boulders), contrast describes its vividness, and directionality describes whether it has a favored direction (like grass) or not (like a smooth object). An object shape query is shown in Figure 8. In this case, the query specification is the drawn shape on the left. Area, circularity, eccentricity, major-axis direction, features derived from the object moments, and a set of tangent angles around the object perimeter are the features used to characterize and match shapes. Figure 9 illustrates query by sketch. In this case, the query specification is a freehand drawing of the dominant lines and edges in the image. The sketch feature is an automatically extracted reduced-resolution 'edge map.' Matching is done by using a template-matching technique. A multiobject query asking for images that contain both a red round object and a green textured object is shown in the bottom of Figure 9. The features are standard color and texture. The matching is done by combining the color and texture distances. Combining distances is applied to arbitrary sets of objects and features to implement logical And semantics. . . . . **ANNOTATION AND DATABASE POPULATION TOOLS**. Automatic methods (such as our Positional Color query) that don't rely on object identification, methods that identify objects automatically as in the museum image example, fast and easy-to-use

Element	'529 Claim Recitation	Exemplary Citations in Reference
		semiautomatic outlining tools, and motion-based segmentation algorithms will enable additional application areas. **FEATURE EXTRACTION AND MATCHING METHODS**. New mathematical representations of video, image, and object attributes that capture 'interesting' features for retrieval are needed. Features that describe new image properties such as alternate texture measures or that are based on fractals or wavelet representations, for example, may offer advantages of representation, indexability, and ease of similarity matching. **INTEGRATION WITH TEXT AND PARAMETRIC ANNOTATION**. Query by visual content complements and extends existing query methods. Systems must be able to integrate queries combining date, subject matter, price, and availability with content properties such as color, texture, and shape." *Id.* at pp. 7–8. Image, Audio, and Video Extenders Administration and Programming (IBM 000001–617): "Your applications can even search for images by content. Imagine an application that uses visual examples to search for images. With such an application, users could select an example image and have the application find other images that have colors or textures similar to those in the example. With DB2 extenders' Query by Image Content (QBIC) capability, you can create applications that search for images in this visual way." *Id.* at p. 4. "To make an image available for searching by content, you catalog the image. When you catalog an image, the Image Extender analyzes the image, by computing the feature values for the image, and stores the values in a QBIC catalog. When you search for an image by content, your query identifies one or more features for the search (such as average color), a source for each feature (such as an example image), and a target set of cataloged images. The Image Extender computes the feature value of the source and compares it to the cataloged feature values for the target images. It then computes a score that indicates how similar the feature values of the target images are to the source. You can have the Image Extender return the images whose features are most similar to

Element	'529 Claim Recitation	Exemplary Citations in Reference
		the source. The Image Extender will return the handle of each image and the image score. You can also have the Image Extender return only the score of a single image." *Id*. at pp. 20–21. QBIC Demo (IBM000747): 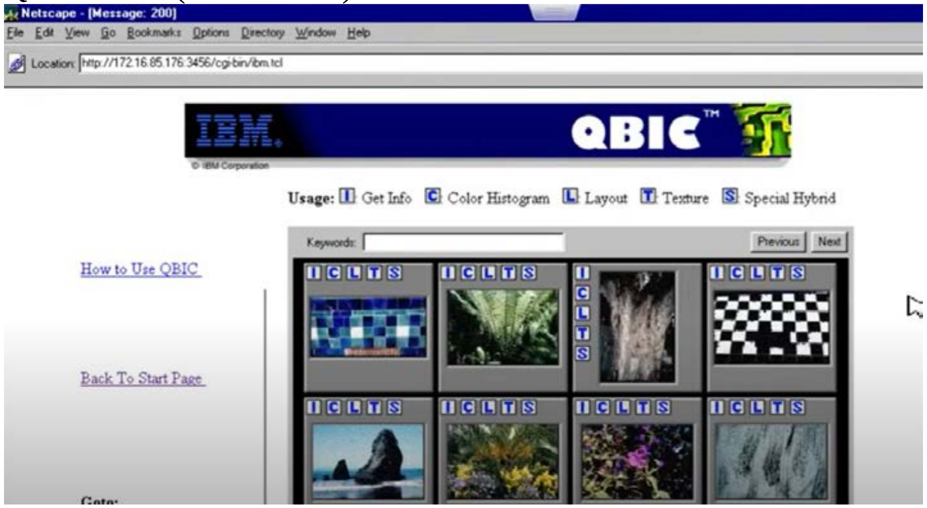 QBIC Demo (IBM000747):

Element	'529 Claim Recitation	Exemplary Citations in Reference
		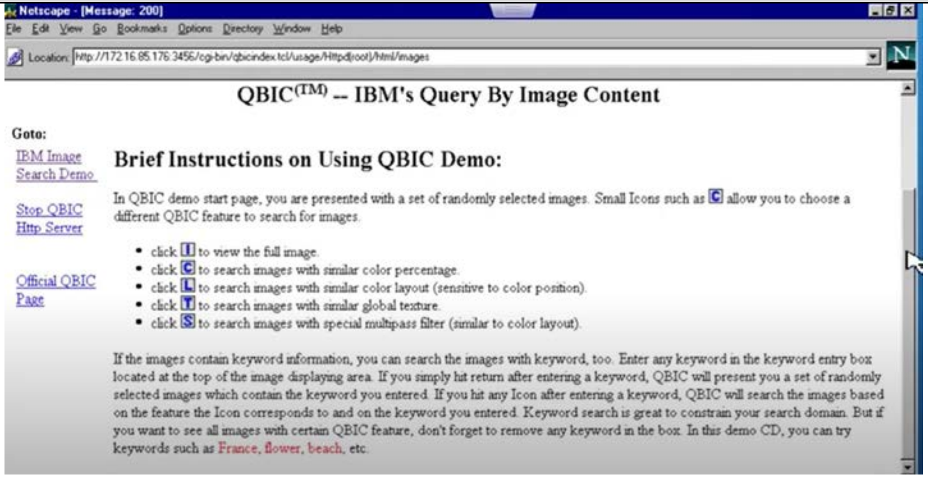 QBIC(TM) -- IBM's Query By Image Content Brief Instructions on Using QBIC Demo: In QBIC demo start page, you are presented with a set of randomly selected images. Small Icons such as C allow you to choose a different QBIC feature to search for images. - click I to view the full image. - click C to search images with similar color percentage. - click L to search images with similar color layout (sensitive to color position). - click T to search images with similar global texture. - click S to search images with special multipass filter (similar to color layout). If the images contain keyword information, you can search the images with keyword, too. Enter any keyword in the keyword entry box located at the top of the image displaying area. If you simply hit return after entering a keyword, QBIC will present you a set of randomly selected images which contain the keyword you entered. If you hit any Icon after entering a keyword, QBIC will search the images based on the feature the Icon corresponds to and on the keyword you entered. Keyword search is great to constrain your search domain. But if you want to see all images with certain QBIC feature, don't forget to remove any keyword in the box. In this demo CD, you can try keywords such as France, flower, beach, etc. QBIC Demo (IBM000747): 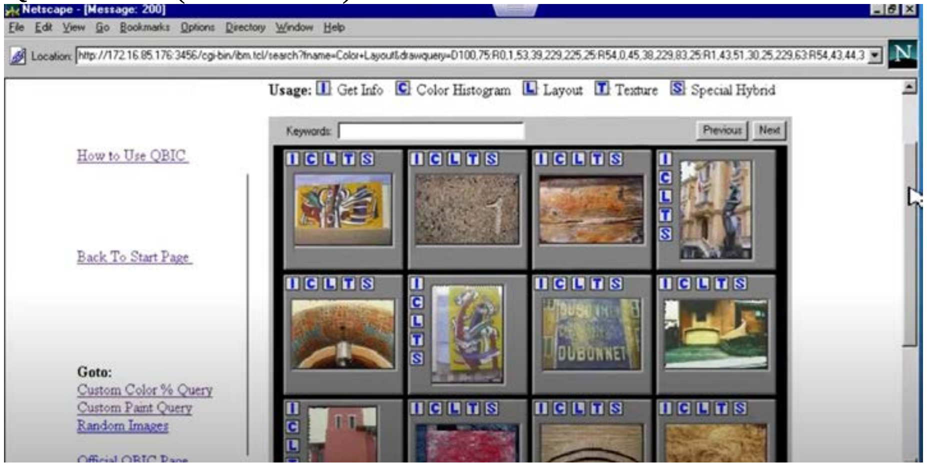

Element	'529 Claim Recitation	Exemplary Citations in Reference
		QBIC Demo (IBM000747): 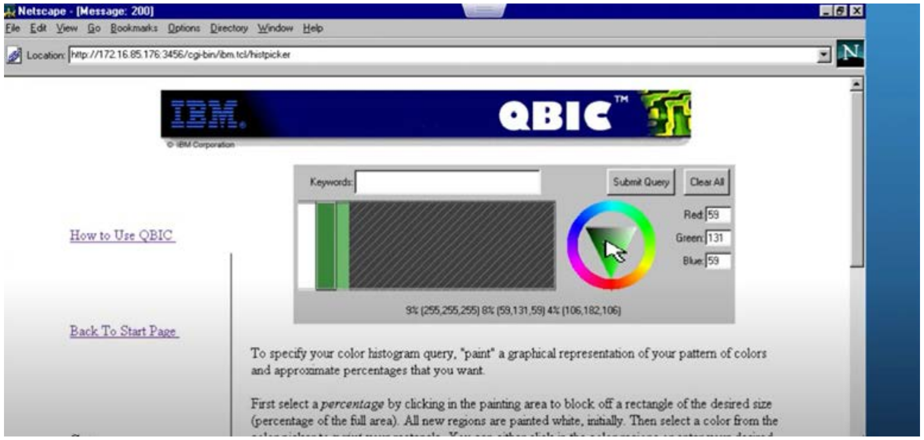 Location: http://172.16.85.176:3456/cgi-bin/ibm.tcl/histpicker Keywords: Red: 59 Green: 131 Blue: 59 9% (255,255,255) 8% (59,131,59) 4% (106,182,106) To specify your color histogram query, "paint" a graphical representation of your pattern of colors and approximate percentages that you want. First select a *percentage* by clicking in the painting area to block off a rectangle of the desired size (percentage of the full area). All new regions are painted white, initially. Then select a color from the QBIC Demo (IBM000747): 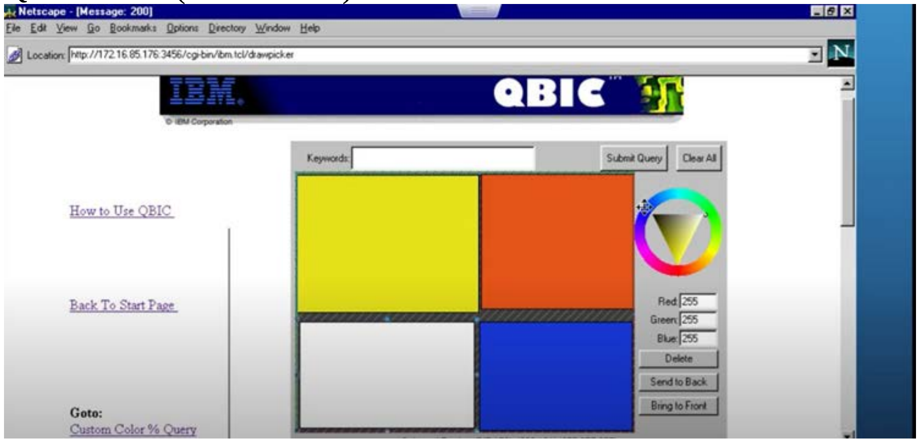 Location: http://172.16.85.176:3456/cgi-bin/ibm.tcl/drawpicker Keywords: Red: 255 Green: 255 Blue: 255

Element	'529 Claim Recitation	Exemplary Citations in Reference
		QBIC Demo (IBM000747):  Location: http://172.16.85.176:3456/cgi-bin/ibm.tcl/search?fname=HybridColor&france10%2Ejpg.x=0 **Goto:** Custom Color % Query Custom Paint Query Random Images Official QBIC Page **Query was:** Example: =france10.jpg Query Type: QbHybridColorFeatureClass Flickner Depo at 20: 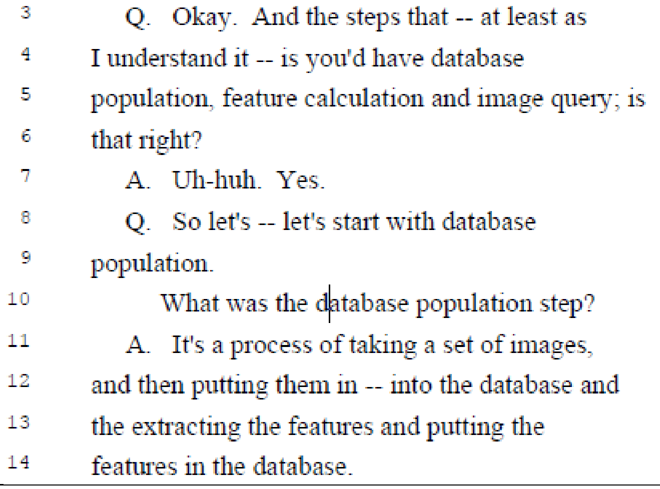 Q. Okay. And the steps that -- at least as I understand it -- is you'd have database population, feature calculation and image query; is that right? A. Uh-huh. Yes. Q. So let's -- let's start with database population. What was the database population step? A. It's a process of taking a set of images, and then putting them in -- into the database and the extracting the features and putting the features in the database.

Element	'529 Claim Recitation	Exemplary Citations in Reference
		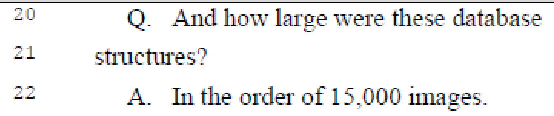 Q. And how large were these database structures? A. In the order of 15,000 images. Flickner Depo at 18: 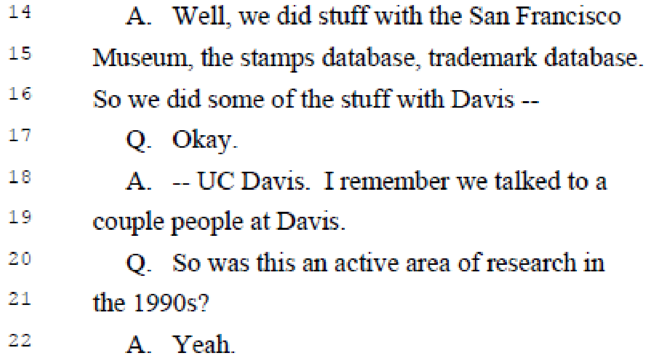 A. Well, we did stuff with the San Francisco Museum, the stamps database, trademark database. So we did some of the stuff with Davis -- Q. Okay. A. -- UC Davis. I remember we talked to a couple people at Davis. Q. So was this an active area of research in the 1990s? A. Yeah. Flickner Depo at 36: 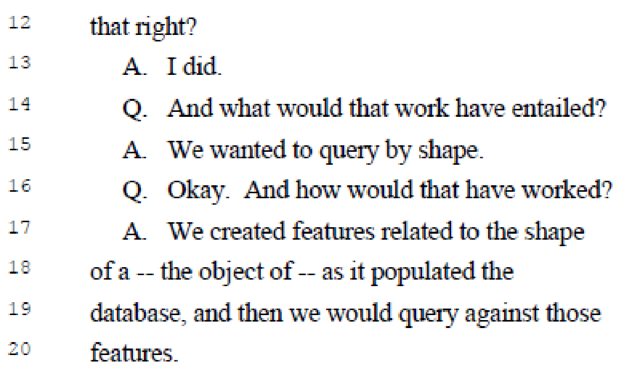 Q. So underlying this paper, you mentioned you did more work on the shape side of things; is that right? A. I did. Q. And what would that work have entailed? A. We wanted to query by shape. Q. Okay. And how would that have worked? A. We created features related to the shape of a -- the object of -- as it populated the database, and then we would query against those features.

Element	'529 Claim Recitation	Exemplary Citations in Reference
		Flickner Depo 38-41: 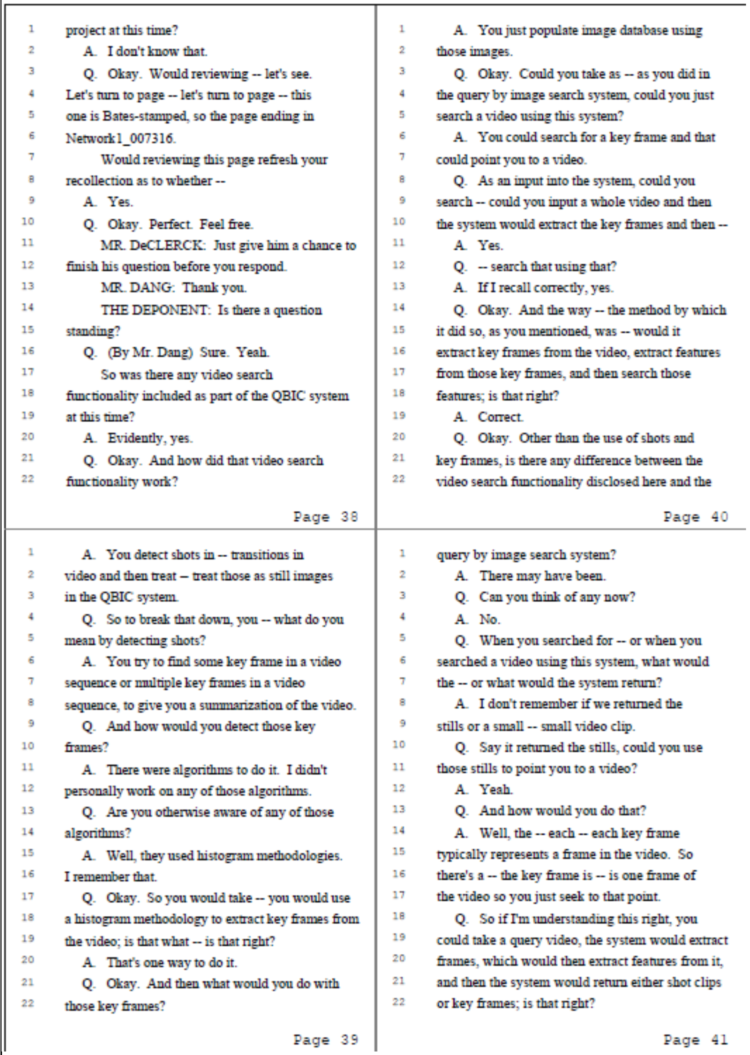 project at this time? A. I don't know that. Q. Okay. Would reviewing -- let's see. Let's turn to page -- let's turn to page -- this one is Bates-stamped, so the page ending in Network1_007316. Would reviewing this page refresh your recollection as to whether -- A. Yes. Q. Okay. Perfect. Feel free. MR. DeCLERCK: Just give him a chance to finish his question before you respond. MR. DANG: Thank you. THE DEPONENT: Is there a question standing? Q. (By Mr. Dang) Sure. Yeah. So was there any video search functionality included as part of the QBIC system at this time? A. Evidently, yes. Q. Okay. And how did that video search functionality work? Page 38 A. You detect shots in -- transitions in video and then treat -- treat those as still images in the QBIC system. Q. So to break that down, you -- what do you mean by detecting shots? A. You try to find some key frame in a video sequence or multiple key frames in a video sequence, to give you a summarization of the video. Q. And how would you detect those key frames? A. There were algorithms to do it. I didn't personally work on any of those algorithms. Q. Are you otherwise aware of any of those algorithms? A. Well, they used histogram methodologies. I remember that. Q. Okay. So you would take -- you would use a histogram methodology to extract key frames from the video; is that what -- is that right? A. That's one way to do it. Q. Okay. And then what would you do with those key frames? Page 39 A. You just populate image database using those images. Q. Okay. Could you take as -- as you did in the query by image search system, could you just search a video using this system? A. You could search for a key frame and that could point you to a video. Q. As an input into the system, could you search -- could you input a whole video and then the system would extract the key frames and then -- A. Yes. Q. -- search that using that? A. If I recall correctly, yes. Q. Okay. And the way -- the method by which it did so, as you mentioned, was -- would it extract key frames from the video, extract features from those key frames, and then search those features; is that right? A. Correct. Q. Okay. Other than the use of shots and key frames, is there any difference between the video search functionality disclosed here and the Page 40 query by image search system? A. There may have been. Q. Can you think of any now? A. No. Q. When you searched for -- or when you searched a video using this system, what would the -- or what would the system return? A. I don't remember if we returned the stills or a small -- small video clip. Q. Say it returned the stills, could you use those stills to point you to a video? A. Yeah. Q. And how would you do that? A. Well, the -- each -- each key frame typically represents a frame in the video. So there's a -- the key frame is -- is one frame of the video so you just seek to that point. Q. So if I'm understanding this right, you could take a query video, the system would extract frames, which would then extract features from it, and then the system would return either shot clips or key frames; is that right? Page 41

Element	'529 Claim Recitation	Exemplary Citations in Reference
		Flickner Depo at 142: 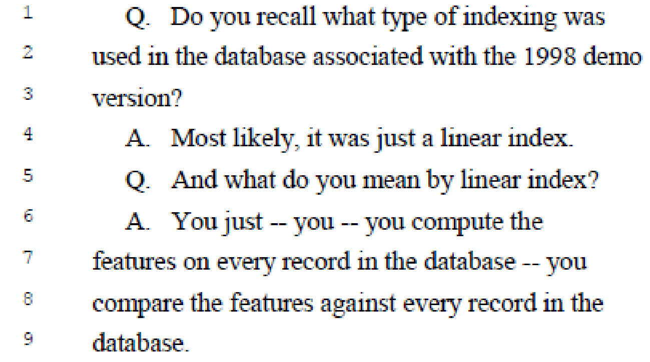 Q. Do you recall what type of indexing was used in the database associated with the 1998 demo version? A. Most likely, it was just a linear index. Q. And what do you mean by linear index? A. You just -- you -- you compute the features on every record in the database -- you compare the features against every record in the database. Flickner Depo at 168-169: 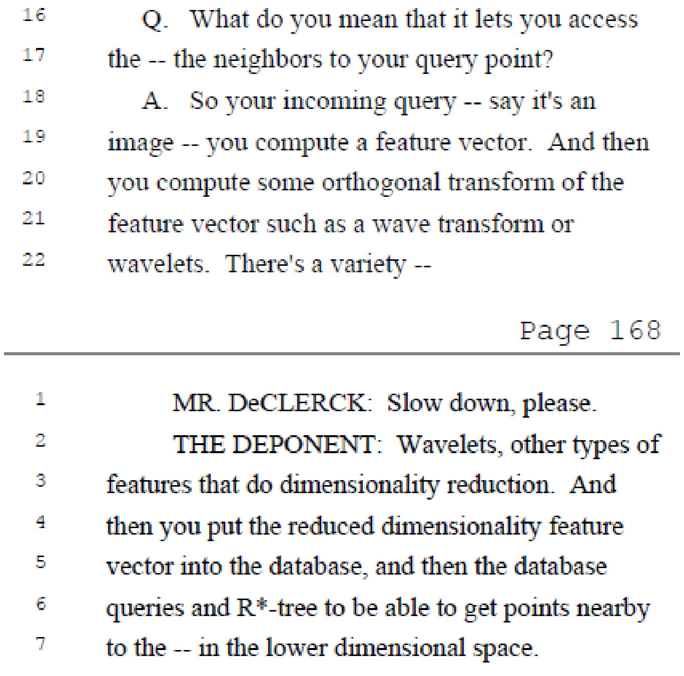 Q. What do you mean that it lets you access the -- the neighbors to your query point? A. So your incoming query -- say it's an image -- you compute a feature vector. And then you compute some orthogonal transform of the feature vector such as a wave transform or wavelets. There's a variety -- Page 168 MR. DeCLERCK: Slow down, please. THE DEPONENT: Wavelets, other types of features that do dimensionality reduction. And then you put the reduced dimensionality feature vector into the database, and then the database queries and R*-tree to be able to get points nearby to the -- in the lower dimensional space.

Element	'529 Claim Recitation	Exemplary Citations in Reference
		Flickner Depo at 181: 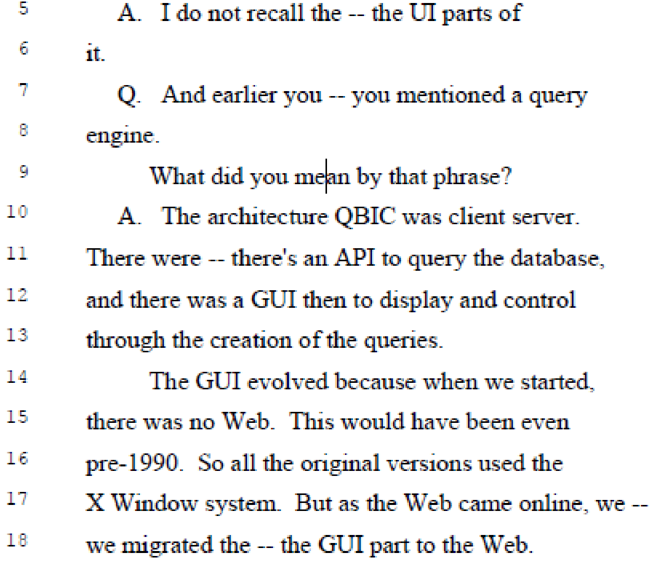 A. I do not recall the -- the UI parts of it. Q. And earlier you -- you mentioned a query engine. What did you mean by that phrase? A. The architecture QBIC was client server. There were -- there's an API to query the database, and there was a GUI then to display and control through the creation of the queries. The GUI evolved because when we started, there was no Web. This would have been even pre-1990. So all the original versions used the X Window system. But as the Web came online, we -- we migrated the -- the GUI part to the Web. Flickner Depo at 48-49: 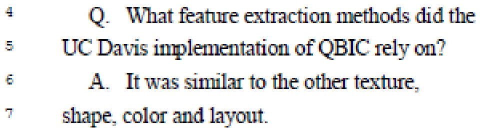 Q. What feature extraction methods did the UC Davis implementation of QBIC rely on? A. It was similar to the other texture, shape, color and layout. *See also* Flickner Depo at 22-25, 25-26, 54-55,75-76, 78-79, 88-89, 91-92, 159, 201-202, 209-210. Flickner 8/29/23 Depo at 33–37:

Element	'529 Claim Recitation	Exemplary Citations in Reference
		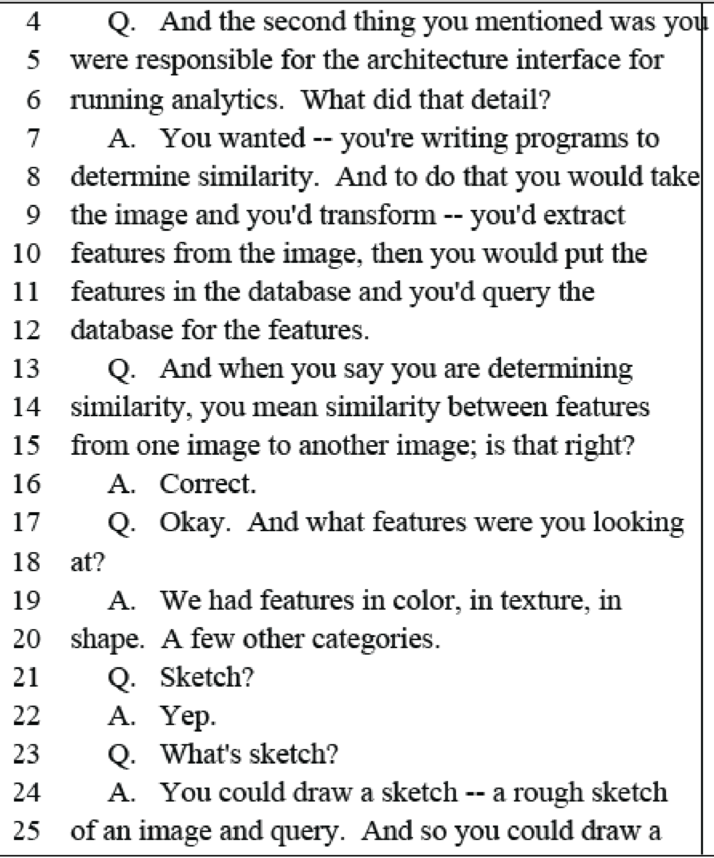 Q. And the second thing you mentioned was you were responsible for the architecture interface for running analytics. What did that detail? A. You wanted -- you're writing programs to determine similarity. And to do that you would take the image and you'd transform -- you'd extract features from the image, then you would put the features in the database and you'd query the database for the features. Q. And when you say you are determining similarity, you mean similarity between features from one image to another image; is that right? A. Correct. Q. Okay. And what features were you looking at? A. We had features in color, in texture, in shape. A few other categories. Q. Sketch? A. Yep. Q. What's sketch? A. You could draw a sketch -- a rough sketch of an image and query. And so you could draw a

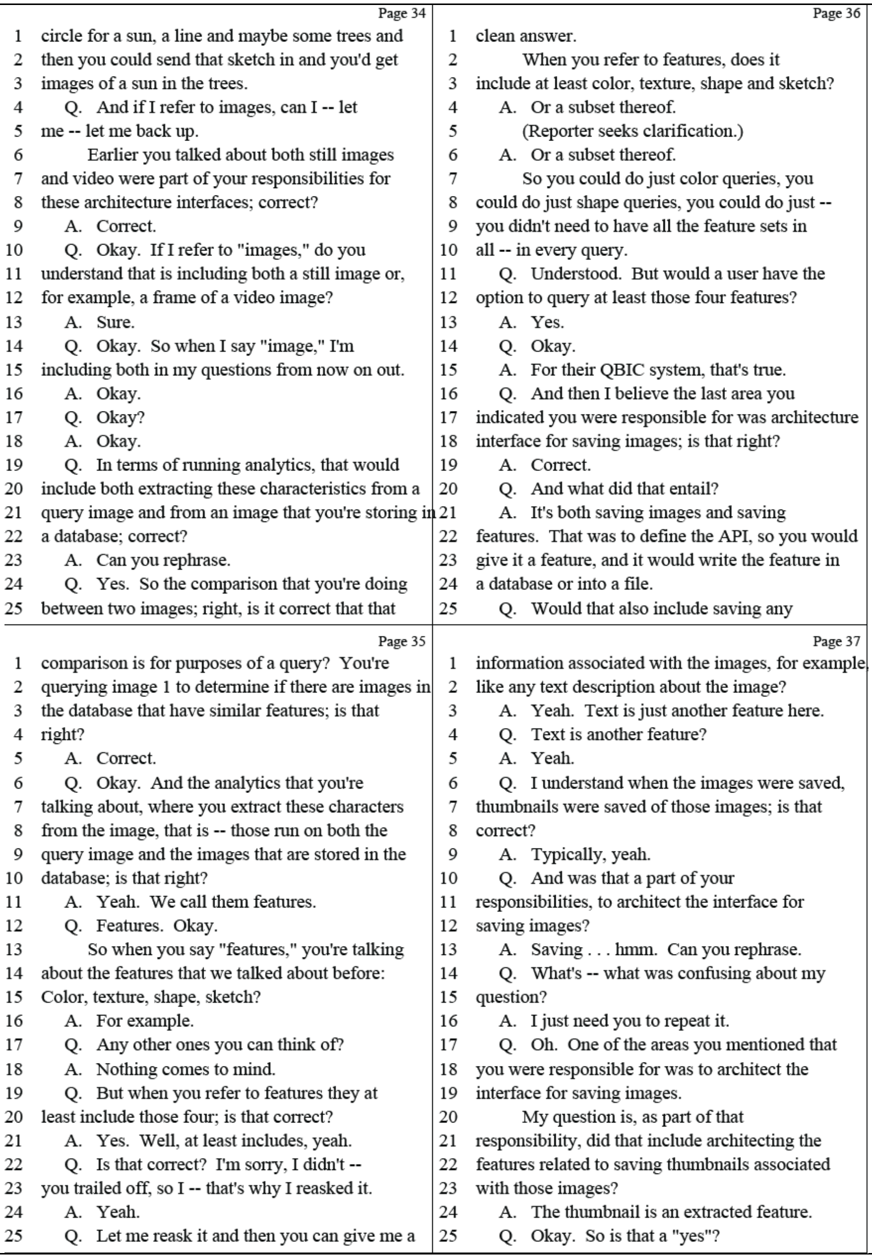

Page 34

circle for a sun, a line and maybe some trees and then you could send that sketch in and you'd get images of a sun in the trees.
Q. And if I refer to images, can I -- let me -- let me back up.
Earlier you talked about both still images and video were part of your responsibilities for these architecture interfaces; correct?
A. Correct.
Q. Okay. If I refer to "images," do you understand that is including both a still image or, for example, a frame of a video image?
A. Sure.
Q. Okay. So when I say "image," I'm including both in my questions from now on out.
A. Okay.
Q. Okay?
A. Okay.
Q. In terms of running analytics, that would include both extracting these characteristics from a query image and from an image that you're storing in a database; correct?
A. Can you rephrase.
Q. Yes. So the comparison that you're doing between two images; right, is it correct that that

Page 35

comparison is for purposes of a query? You're querying image 1 to determine if there are images in the database that have similar features; is that right?
A. Correct.
Q. Okay. And the analytics that you're talking about, where you extract these characters from the image, that is -- those run on both the query image and the images that are stored in the database; is that right?
A. Yeah. We call them features.
Q. Features. Okay.
So when you say "features," you're talking about the features that we talked about before: Color, texture, shape, sketch?
A. For example.
Q. Any other ones you can think of?
A. Nothing comes to mind.
Q. But when you refer to features they at least include those four; is that correct?
A. Yes. Well, at least includes, yeah.
Q. Is that correct? I'm sorry, I didn't -- you trailed off, so I -- that's why I reasked it.
A. Yeah.
Q. Let me reask it and then you can give me a

Page 36

clean answer.
When you refer to features, does it include at least color, texture, shape and sketch?
A. Or a subset thereof.
(Reporter seeks clarification.)
A. Or a subset thereof.
So you could do just color queries, you could do just shape queries, you could do just -- you didn't need to have all the feature sets in all -- in every query.
Q. Understood. But would a user have the option to query at least those four features?
A. Yes.
Q. Okay.
A. For their QBIC system, that's true.
Q. And then I believe the last area you indicated you were responsible for was architecture interface for saving images; is that right?
A. Correct.
Q. And what did that entail?
A. It's both saving images and saving features. That was to define the API, so you would give it a feature, and it would write the feature in a database or into a file.
Q. Would that also include saving any

Page 37

information associated with the images, for example, like any text description about the image?
A. Yeah. Text is just another feature here.
Q. Text is another feature?
A. Yeah.
Q. I understand when the images were saved, thumbnails were saved of those images; is that correct?
A. Typically, yeah.
Q. And was that a part of your responsibilities, to architect the interface for saving images?
A. Saving . . . hmm. Can you rephrase.
Q. What's -- what was confusing about my question?
A. I just need you to repeat it.
Q. Oh. One of the areas you mentioned that you were responsible for was to architect the interface for saving images.
My question is, as part of that responsibility, did that include architecting the features related to saving thumbnails associated with those images?
A. The thumbnail is an extracted feature.
Q. Okay. So is that a "yes"?

Element	'529 Claim Recitation	Exemplary Citations in Reference
		Id. at 52–53: 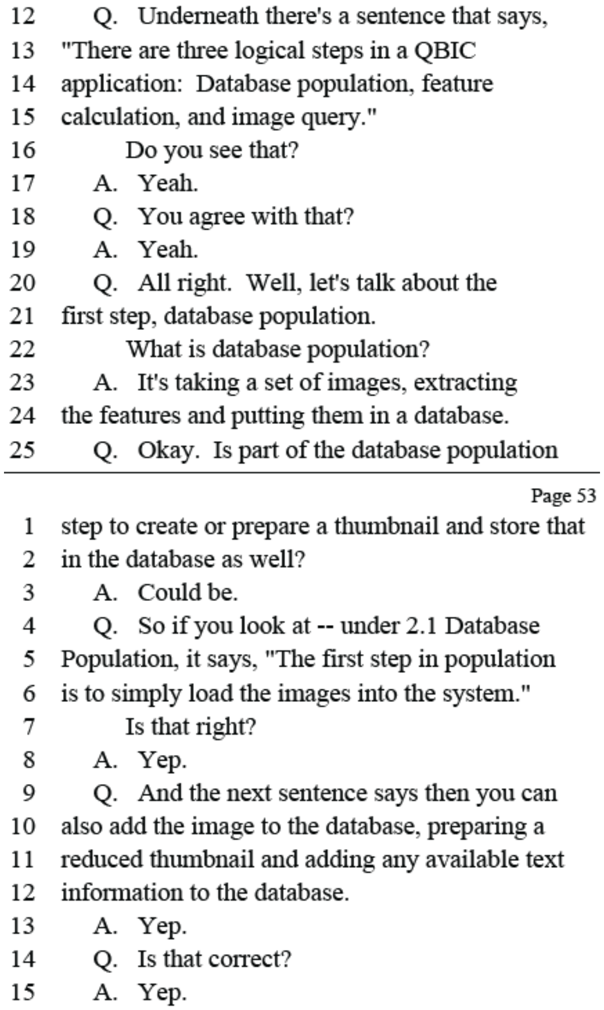 Q. Underneath there's a sentence that says, "There are three logical steps in a QBIC application: Database population, feature calculation, and image query." Do you see that? A. Yeah. Q. You agree with that? A. Yeah. Q. All right. Well, let's talk about the first step, database population. What is database population? A. It's taking a set of images, extracting the features and putting them in a database. Q. Okay. Is part of the database population Page 53 step to create or prepare a thumbnail and store that in the database as well? A. Could be. Q. So if you look at -- under 2.1 Database Population, it says, "The first step in population is to simply load the images into the system." Is that right? A. Yep. Q. And the next sentence says then you can also add the image to the database, preparing a reduced thumbnail and adding any available text information to the database. A. Yep. Q. Is that correct? A. Yep. *Id.* at 77–78:

Element	'529 Claim Recitation	Exemplary Citations in Reference
		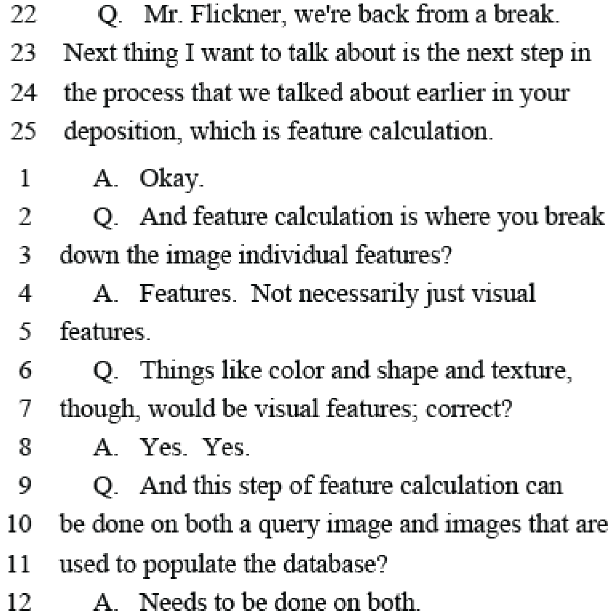 Q. Mr. Flickner, we're back from a break. Next thing I want to talk about is the next step in the process that we talked about earlier in your deposition, which is feature calculation. A. Okay. Q. And feature calculation is where you break down the image individual features? A. Features. Not necessarily just visual features. Q. Things like color and shape and texture, though, would be visual features; correct? A. Yes. Yes. Q. And this step of feature calculation can be done on both a query image and images that are used to populate the database? A. Needs to be done on both. *Id*. at 85–99: 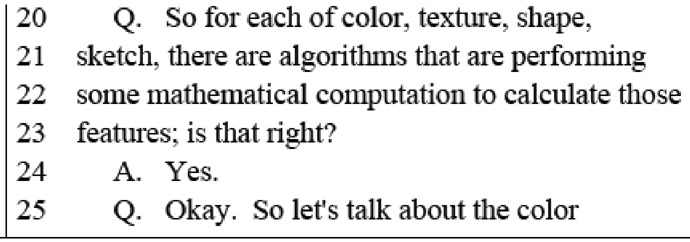 Q. So for each of color, texture, shape, sketch, there are algorithms that are performing some mathematical computation to calculate those features; is that right? A. Yes. Q. Okay. So let's talk about the color

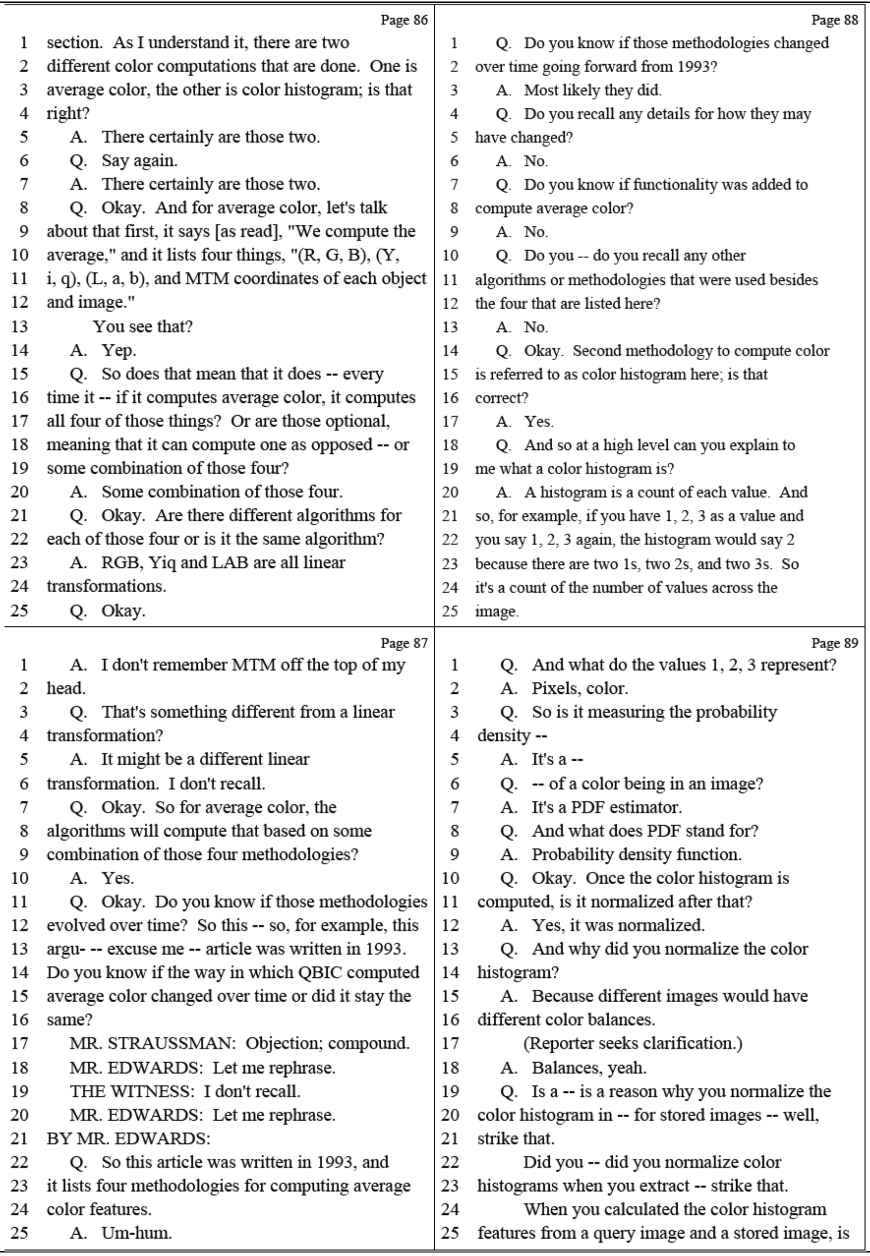

Page 86

section. As I understand it, there are two different color computations that are done. One is average color, the other is color histogram; is that right?
A. There certainly are those two.
Q. Say again.
A. There certainly are those two.
Q. Okay. And for average color, let's talk about that first, it says [as read], "We compute the average," and it lists four things, "(R, G, B), (Y, i, q), (L, a, b), and MTM coordinates of each object and image."
You see that?
A. Yep.
Q. So does that mean that it does -- every time it -- if it computes average color, it computes all four of those things? Or are those optional, meaning that it can compute one as opposed -- or some combination of those four?
A. Some combination of those four.
Q. Okay. Are there different algorithms for each of those four or is it the same algorithm?
A. RGB, Yiq and LAB are all linear transformations.
Q. Okay.

Page 87

A. I don't remember MTM off the top of my head.
Q. That's something different from a linear transformation?
A. It might be a different linear transformation. I don't recall.
Q. Okay. So for average color, the algorithms will compute that based on some combination of those four methodologies?
A. Yes.
Q. Okay. Do you know if those methodologies evolved over time? So this -- so, for example, this argu- -- excuse me -- article was written in 1993. Do you know if the way in which QBIC computed average color changed over time or did it stay the same?
MR. STRAUSSMAN: Objection; compound.
MR. EDWARDS: Let me rephrase.
THE WITNESS: I don't recall.
MR. EDWARDS: Let me rephrase.
BY MR. EDWARDS:
Q. So this article was written in 1993, and it lists four methodologies for computing average color features.
A. Um-hum.

Page 88

Q. Do you know if those methodologies changed over time going forward from 1993?
A. Most likely they did.
Q. Do you recall any details for how they may have changed?
A. No.
Q. Do you know if functionality was added to compute average color?
A. No.
Q. Do you -- do you recall any other algorithms or methodologies that were used besides the four that are listed here?
A. No.
Q. Okay. Second methodology to compute color is referred to as color histogram here; is that correct?
A. Yes.
Q. And so at a high level can you explain to me what a color histogram is?
A. A histogram is a count of each value. And so, for example, if you have 1, 2, 3 as a value and you say 1, 2, 3 again, the histogram would say 2 because there are two 1s, two 2s, and two 3s. So it's a count of the number of values across the image.

Page 89

Q. And what do the values 1, 2, 3 represent?
A. Pixels, color.
Q. So is it measuring the probability density --
A. It's a --
Q. -- of a color being in an image?
A. It's a PDF estimator.
Q. And what does PDF stand for?
A. Probability density function.
Q. Okay. Once the color histogram is computed, is it normalized after that?
A. Yes, it was normalized.
Q. And why did you normalize the color histogram?
A. Because different images would have different color balances.
(Reporter seeks clarification.)
A. Balances, yeah.
Q. Is a -- is a reason why you normalize the color histogram in -- for stored images -- well, strike that.
Did you -- did you normalize color histograms when you extract -- strike that.
When you calculated the color histogram features from a query image and a stored image, is

Page 90

one reason why you normalize it so you're comparing apples to apples?
A. Yes.
Q. And then going back to average color, just for the record, what does RGB stand for?
A. Red, green, blue.
Q. And Yiq?
A. Luminance and chrominance. Y is --
(Reporter seeks clarification.)
A. The Y is luminance and the IQ is chrominance.
Q. And LAB?
A. I think LAB -- the L is luminance. AB are different chrominance transformations.
Q. And AB, you say chrominance?
A. Chrominance.
Q. C-h-r-o-m-a-n-c-e [sic]?
A. Sounds about right.
Q. Okay. And then the MTM, mathematical transform to Munsell, is that what that stands for?
A. Yeah. That was a particular color space that was visually appeasing.
Q. Okay. And then for -- so, for example, for RGB, that's red, green, blue; correct?
A. Correct.

Page 91

Q. Those are different color features that are -- that are decomposed or broken down from the image?
A. Yeah, the value of each pixel.
Q. Okay. And then going back to the color histogram, so the color histogram's calculated, it's normalized. After it's normalized, is there anything else done to it?
A. Sounds like they're computing MTM coordinates of each cell and then clustering. This is to help quantize the color space.
Q. Okay. And that would be -- you would do that both for a query image and a stored image?
A. Yes.
Q. Are there any details about color histogramming that you recall that are not disclosed in this paragraph?
A. Not that I recall.
Q. Do you recall if the methodology changed over time?
A. Most likely it did, but I don't recall specifics.
Q. Okay. Okay. Let's go on to texture, and that's at the bottom of 2392 and goes to the top of 2393. Take a look at that and then let me know when

Page 92

you're done.
(Witness reviews document.)
A. Okay.
Q. So for texture, it looks like the calculations were based on coarseness, contrast and directionality; is that right?
A. Correct.
Q. Are those three different algorithms that are used to calculate those, or is it the same algorithm?
A. It's three different algorithms.
Q. And I think earlier you mentioned that gray scaling was used in the texture calculation. For which of those three features is the image converted to gray scale?
A. Probably all three.
Q. Okay. So coarseness is a measure of the scale of the textures, such as pebbles versus boulders?
A. Correct.
Q. Contrast is the vividness of a pattern?
A. Correct.
Q. And then this says that contrast is a function of the variance of the gray-level histogram; is that right?

Page 93

A. Correct.
Q. So at least for contrast that's converted to a gray scale histogram; correct?
A. Correct.
Q. And directionality describes whether or not the image has a direction, like grass; is that right?
A. Correct.
(Reporter seeks clarification.)
Q. Do you recall any other details about the texture feature calculation that is not disclosed in this paragraph?
A. No.
Q. Do you recall if the texture feature calculation changed over time after this article?
A. Most likely it did.
Q. Do you recall any specifics?
A. No.
Q. All right. Let's go on to shape. Take a look at that paragraph and then let me know when you're done. It's on 2393.
(Witness reviews document.)
Q. And while you're looking at that, Mr. Flickner, for any of these four features that we're talking about, did you -- did your work on the

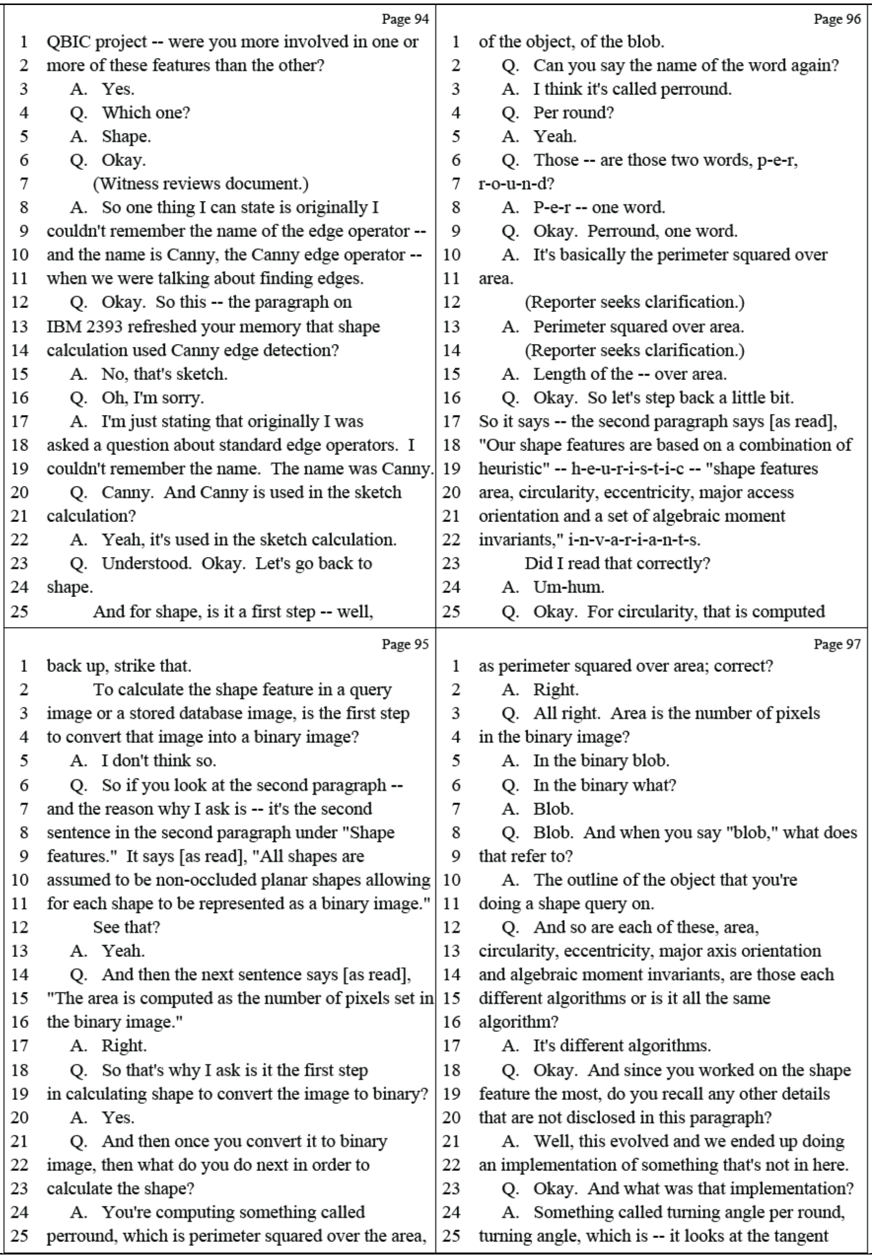

Page 94

QBIC project -- were you more involved in one or more of these features than the other?
A. Yes.
Q. Which one?
A. Shape.
Q. Okay.
(Witness reviews document.)
A. So one thing I can state is originally I couldn't remember the name of the edge operator -- and the name is Canny, the Canny edge operator -- when we were talking about finding edges.
Q. Okay. So this -- the paragraph on IBM 2393 refreshed your memory that shape calculation used Canny edge detection?
A. No, that's sketch.
Q. Oh, I'm sorry.
A. I'm just stating that originally I was asked a question about standard edge operators. I couldn't remember the name. The name was Canny.
Q. Canny. And Canny is used in the sketch calculation?
A. Yeah, it's used in the sketch calculation.
Q. Understood. Okay. Let's go back to shape.
And for shape, is it a first step -- well,

Page 95

back up, strike that.
To calculate the shape feature in a query image or a stored database image, is the first step to convert that image into a binary image?
A. I don't think so.
Q. So if you look at the second paragraph -- and the reason why I ask is -- it's the second sentence in the second paragraph under "Shape features." It says [as read], "All shapes are assumed to be non-occluded planar shapes allowing for each shape to be represented as a binary image."
See that?
A. Yeah.
Q. And then the next sentence says [as read], "The area is computed as the number of pixels set in the binary image."
A. Right.
Q. So that's why I ask is it the first step in calculating shape to convert the image to binary?
A. Yes.
Q. And then once you convert it to binary image, then what do you do next in order to calculate the shape?
A. You're computing something called perround, which is perimeter squared over the area,

Page 96

of the object, of the blob.
Q. Can you say the name of the word again?
A. I think it's called perround.
Q. Per round?
A. Yeah.
Q. Those -- are those two words, p-e-r, r-o-u-n-d?
A. P-e-r -- one word.
Q. Okay. Perround, one word.
A. It's basically the perimeter squared over area.
(Reporter seeks clarification.)
A. Perimeter squared over area.
(Reporter seeks clarification.)
A. Length of the -- over area.
Q. Okay. So let's step back a little bit. So it says -- the second paragraph says [as read], "Our shape features are based on a combination of heuristic" -- h-e-u-r-i-s-t-i-c -- "shape features area, circularity, eccentricity, major access orientation and a set of algebraic moment invariants," i-n-v-a-r-i-a-n-t-s.
Did I read that correctly?
A. Um-hum.
Q. Okay. For circularity, that is computed

Page 97

as perimeter squared over area; correct?
A. Right.
Q. All right. Area is the number of pixels in the binary image?
A. In the binary blob.
Q. In the binary what?
A. Blob.
Q. Blob. And when you say "blob," what does that refer to?
A. The outline of the object that you're doing a shape query on.
Q. And so are each of these, area, circularity, eccentricity, major axis orientation and algebraic moment invariants, are those each different algorithms or is it all the same algorithm?
A. It's different algorithms.
Q. Okay. And since you worked on the shape feature the most, do you recall any other details that are not disclosed in this paragraph?
A. Well, this evolved and we ended up doing an implementation of something that's not in here.
Q. Okay. And what was that implementation?
A. Something called turning angle per round, turning angle, which is -- it looks at the tangent

			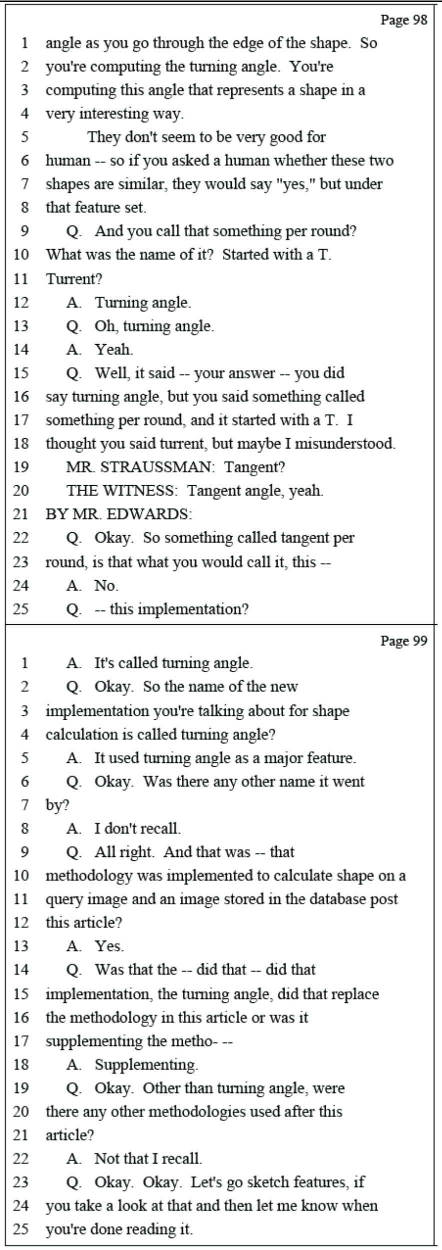 Page 98 angle as you go through the edge of the shape. So you're computing the turning angle. You're computing this angle that represents a shape in a very interesting way. They don't seem to be very good for human -- so if you asked a human whether these two shapes are similar, they would say "yes," but under that feature set. Q. And you call that something per round? What was the name of it? Started with a T. Turrent? A. Turning angle. Q. Oh, turning angle. A. Yeah. Q. Well, it said -- your answer -- you did say turning angle, but you said something called something per round, and it started with a T. I thought you said turrent, but maybe I misunderstood. MR. STRAUSSMAN: Tangent? THE WITNESS: Tangent angle, yeah. BY MR. EDWARDS: Q. Okay. So something called tangent per round, is that what you would call it, this -- A. No. Q. -- this implementation? Page 99 A. It's called turning angle. Q. Okay. So the name of the new implementation you're talking about for shape calculation is called turning angle? A. It used turning angle as a major feature. Q. Okay. Was there any other name it went by? A. I don't recall. Q. All right. And that was -- that methodology was implemented to calculate shape on a query image and an image stored in the database post this article? A. Yes. Q. Was that the -- did that -- did that implementation, the turning angle, did that replace the methodology in this article or was it supplementing the metho- -- A. Supplementing. Q. Okay. Other than turning angle, were there any other methodologies used after this article? A. Not that I recall. Q. Okay. Okay. Let's go sketch features, if you take a look at that and then let me know when you're done reading it.	

Element	'529 Claim Recitation	Exemplary Citations in Reference
		Id. at 119–128: 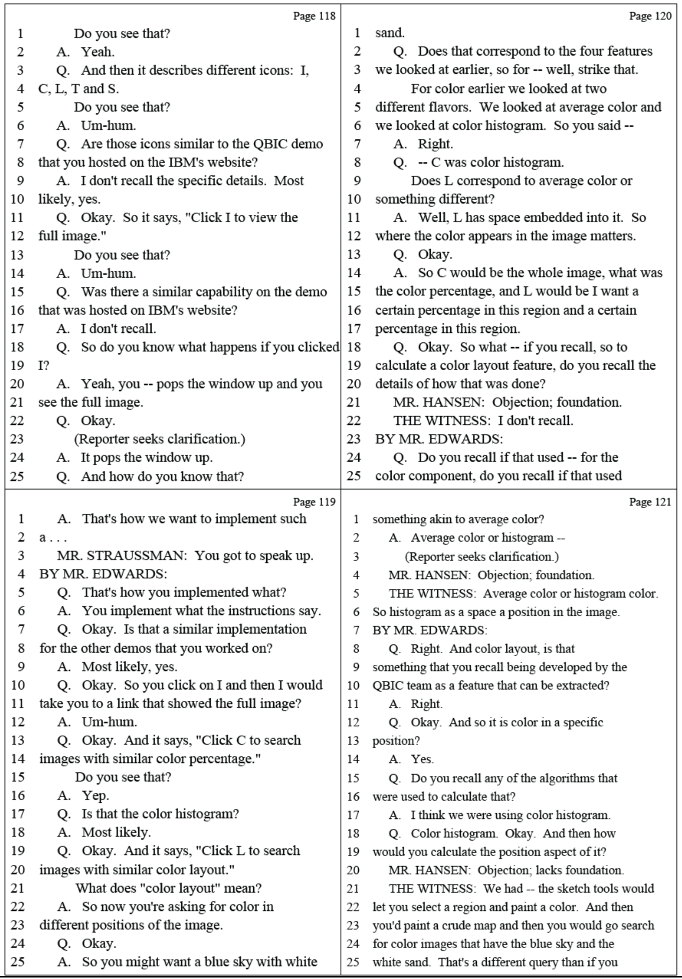 Page 118 Do you see that? A. Yeah. Q. And then it describes different icons: I, C, L, T and S. Do you see that? A. Um-hum. Q. Are those icons similar to the QBIC demo that you hosted on the IBM's website? A. I don't recall the specific details. Most likely, yes. Q. Okay. So it says, "Click I to view the full image." Do you see that? A. Um-hum. Q. Was there a similar capability on the demo that was hosted on IBM's website? A. I don't recall. Q. So do you know what happens if you clicked I? A. Yeah, you -- pops the window up and you see the full image. Q. Okay. (Reporter seeks clarification.) A. It pops the window up. Q. And how do you know that? Page 119 A. That's how we want to implement such a . . . MR. STRAUSSMAN: You got to speak up. BY MR. EDWARDS: Q. That's how you implemented what? A. You implement what the instructions say. Q. Okay. Is that a similar implementation for the other demos that you worked on? A. Most likely, yes. Q. Okay. So you click on I and then I would take you to a link that showed the full image? A. Um-hum. Q. Okay. And it says, "Click C to search images with similar color percentage." Do you see that? A. Yep. Q. Is that the color histogram? A. Most likely. Q. Okay. And it says, "Click L to search images with similar color layout." What does "color layout" mean? A. So now you're asking for color in different positions of the image. Q. Okay. A. So you might want a blue sky with white Page 120 sand. Q. Does that correspond to the four features we looked at earlier, so for -- well, strike that. For color earlier we looked at two different flavors. We looked at average color and we looked at color histogram. So you said -- A. Right. Q. -- C was color histogram. Does L correspond to average color or something different? A. Well, L has space embedded into it. So where the color appears in the image matters. Q. Okay. A. So C would be the whole image, what was the color percentage, and L would be I want a certain percentage in this region and a certain percentage in this region. Q. Okay. So what -- if you recall, so to calculate a color layout feature, do you recall the details of how that was done? MR. HANSEN: Objection; foundation. THE WITNESS: I don't recall. BY MR. EDWARDS: Q. Do you recall if that used -- for the color component, do you recall if that used Page 121 something akin to average color? A. Average color or histogram -- (Reporter seeks clarification.) MR. HANSEN: Objection; foundation. THE WITNESS: Average color or histogram color. So histogram as a space a position in the image. BY MR. EDWARDS: Q. Right. And color layout, is that something that you recall being developed by the QBIC team as a feature that can be extracted? A. Right. Q. Okay. And so it is color in a specific position? A. Yes. Q. Do you recall any of the algorithms that were used to calculate that? A. I think we were using color histogram. Q. Color histogram. Okay. And then how would you calculate the position aspect of it? MR. HANSEN: Objection; lacks foundation. THE WITNESS: We had -- the sketch tools would let you select a region and paint a color. And then you'd paint a crude map and then you would go search for color images that have the blue sky and the white sand. That's a different query than if you

Page 122

had white on top and blue on the bottom. Then you're looking for blue water, white sand.
BY MR. EDWARDS:
Q. Do you recall what the methodologies used -- that were used to calculate the orientation of those colors in the image?
MR. HANSEN: Objection; lacks foundation.
THE WITNESS: There's no orientation per se. It's just position.
BY MR. EDWARDS:
Q. Well, then position. Do you recall what the methodologies are to calculate the position?
MR. HANSEN: Same objection.
THE WITNESS: Position is something the user inputs.
BY MR. EDWARDS:
Q. Something you use what?
A. The user had to input.
Q. Okay. And then let's move on to the next one, T. It says, "Click T to search images with similar global texture."
A. Yeah.
Q. Do you recall, in your work on the QBIC project, a feature called global texture that could be extracted from the image?

Page 123

A. A feature called . . .
Q. Say again?
A. I missed one of your words. A feature called what?
Q. I said do you recall in your work on the QBIC project a feature called global texture that was extracted from the image?
A. Yes.
Q. Okay. Is that -- is that akin to texture that we looked at earlier?
A. Yes.
Q. Okay. And the last one is S. It's, "Search images with special multipass filter."
Do you recall, in your work on the QBIC project, being able to calculate a feature called multipass filter?
A. Probably Gabor functions.
Q. Say again?
A. Gabor functions.
Q. One more time and louder.
A. Gabor functions.
Q. Gabor functions. Okay. What are Gabor functions?
A. Gabor functions are a way of producing multiple resolution closed contours.

Page 124

Q. Okay. And how do you spell Gabor?
A. G-a-b-o-r.
Q. So switch over and look at the second page of this Exhibit 6.
So -- well, first back -- go back to -- go back to the first page. If you see on the left-hand side it says, "Goto: IBM Image Search Demo," and a little finger there is clicking it.
You see it's in red?
A. Yeah.
Q. Okay. So that's clicking on that link. And then the second page is going to the IBM image search demo.
A. Okay.
Q. Okay? On the second page it is showing thumbnails with the different icons that we just walked through, I, C, L, T, S, on each of those icons.
Do you see that?
A. Yep.
Q. Do you recognize this as a similar layout to the QBIC demos you hosted in the '90s?
A. I don't.
Q. You don't?
A. No.

Page 125

Q. This was not -- so the link to the IBM website that was the AIX operating system in Exhibit 3, that did not have a similar thumbnail layout like this?
A. A thumbnail, yes. I don't recall it had the buttons, the I, C, L, T, S.
Q. You don't recall it had all the buttons or some of the buttons?
A. All the buttons.
Q. Okay. So you don't recall that it had any buttons?
A. I don't -- yeah. I don't recall.
Q. It could have, it couldn't have, you just don't remember?
A. I just don't remember, yeah.
Q. Okay. But the one that you hosted on IBM's website, it did have thumbnails in a similar format here?
A. Yeah, it had thumbnails.
Q. And these thumbnails represent the images that are stored in the database?
A. Yes.
Q. Okay. Here, the 470 images would be reflected at least -- one, two, three, four, five, six, seven -- eight of them right here in this

Element	'529 Claim Recitation	Exemplary Citations in Reference
		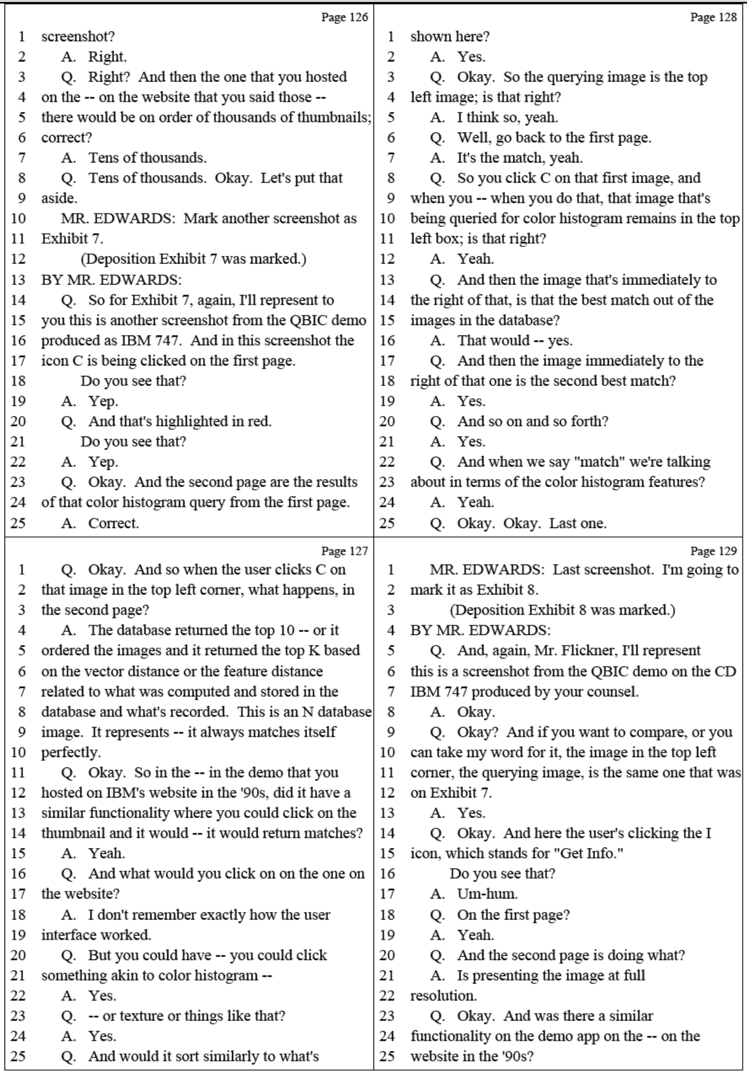 Page 126 screenshot? A. Right. Q. Right? And then the one that you hosted on the -- on the website that you said those -- there would be on order of thousands of thumbnails; correct? A. Tens of thousands. Q. Tens of thousands. Okay. Let's put that aside. MR. EDWARDS: Mark another screenshot as Exhibit 7. (Deposition Exhibit 7 was marked.) BY MR. EDWARDS: Q. So for Exhibit 7, again, I'll represent to you this is another screenshot from the QBIC demo produced as IBM 747. And in this screenshot the icon C is being clicked on the first page. Do you see that? A. Yep. Q. And that's highlighted in red. Do you see that? A. Yep. Q. Okay. And the second page are the results of that color histogram query from the first page. A. Correct. Page 127 Q. Okay. And so when the user clicks C on that image in the top left corner, what happens, in the second page? A. The database returned the top 10 -- or it ordered the images and it returned the top K based on the vector distance or the feature distance related to what was computed and stored in the database and what's recorded. This is an N database image. It represents -- it always matches itself perfectly. Q. Okay. So in the -- in the demo that you hosted on IBM's website in the '90s, did it have a similar functionality where you could click on the thumbnail and it would -- it would return matches? A. Yeah. Q. And what would you click on on the one on the website? A. I don't remember exactly how the user interface worked. Q. But you could have -- you could click something akin to color histogram -- A. Yes. Q. -- or texture or things like that? A. Yes. Q. And would it sort similarly to what's Page 128 shown here? A. Yes. Q. Okay. So the querying image is the top left image; is that right? A. I think so, yeah. Q. Well, go back to the first page. A. It's the match, yeah. Q. So you click C on that first image, and when you -- when you do that, that image that's being queried for color histogram remains in the top left box; is that right? A. Yeah. Q. And then the image that's immediately to the right of that, is that the best match out of the images in the database? A. That would -- yes. Q. And then the image immediately to the right of that one is the second best match? A. Yes. Q. And so on and so forth? A. Yes. Q. And when we say "match" we're talking about in terms of the color histogram features? A. Yeah. Q. Okay. Okay. Last one. Page 129 MR. EDWARDS: Last screenshot. I'm going to mark it as Exhibit 8. (Deposition Exhibit 8 was marked.) BY MR. EDWARDS: Q. And, again, Mr. Flickner, I'll represent this is a screenshot from the QBIC demo on the CD IBM 747 produced by your counsel. A. Okay. Q. Okay? And if you want to compare, or you can take my word for it, the image in the top left corner, the querying image, is the same one that was on Exhibit 7. A. Yes. Q. Okay. And here the user's clicking the I icon, which stands for "Get Info." Do you see that? A. Um-hum. Q. On the first page? A. Yeah. Q. And the second page is doing what? A. Is presenting the image at full resolution. Q. Okay. And was there a similar functionality on the demo app on the -- on the website in the '90s? *Id*. at 135 – 136:

Element	'529 Claim Recitation	Exemplary Citations in Reference
		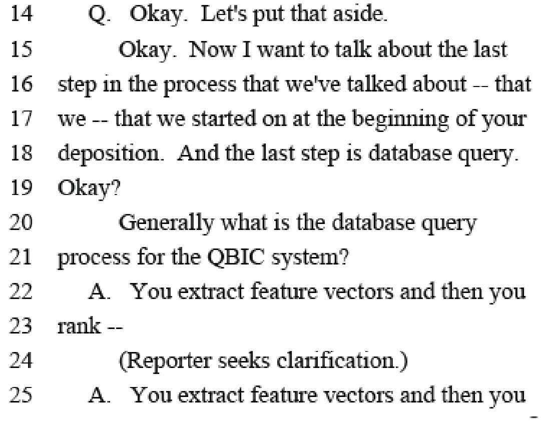 Q. Okay. Let's put that aside. Okay. Now I want to talk about the last step in the process that we've talked about -- that we -- that we started on at the beginning of your deposition. And the last step is database query. Okay? Generally what is the database query process for the QBIC system? A. You extract feature vectors and then you rank -- (Reporter seeks clarification.) A. You extract feature vectors and then you 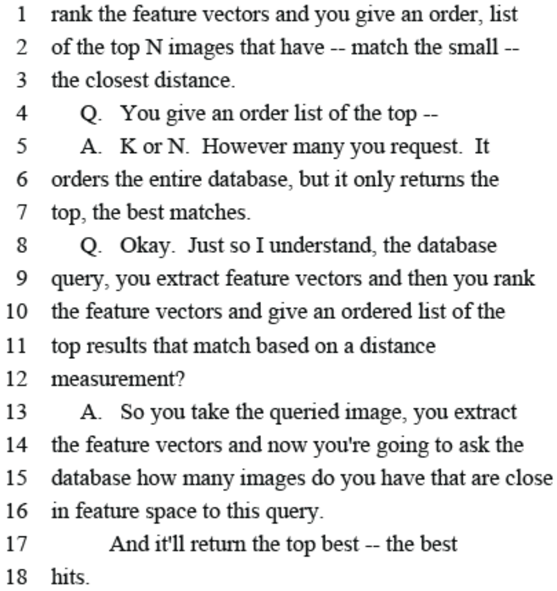 rank the feature vectors and you give an order, list of the top N images that have -- match the small -- the closest distance. Q. You give an order list of the top -- A. K or N. However many you request. It orders the entire database, but it only returns the top, the best matches. Q. Okay. Just so I understand, the database query, you extract feature vectors and then you rank the feature vectors and give an ordered list of the top results that match based on a distance measurement? A. So you take the queried image, you extract the feature vectors and now you're going to ask the database how many images do you have that are close in feature space to this query. And it'll return the top best -- the best hits. *Id.* at 139–165:

Element	'529 Claim Recitation	Exemplary Citations in Reference
		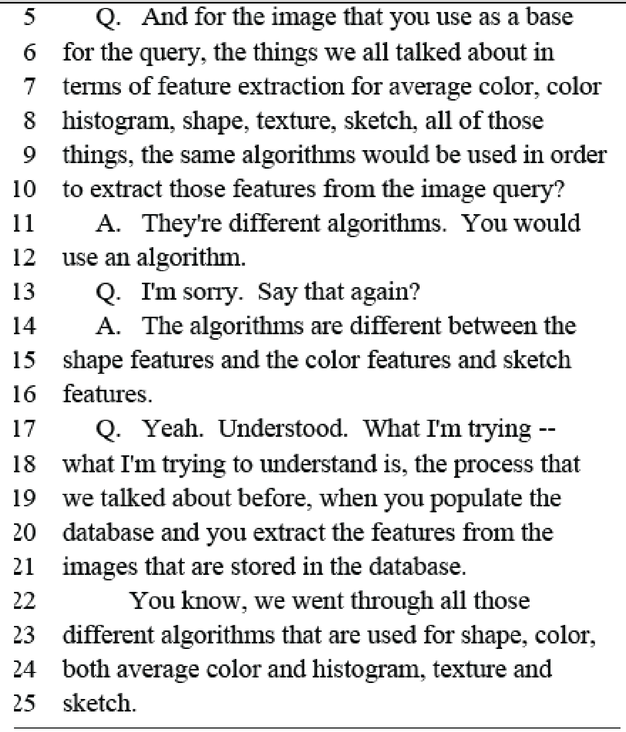 Q. And for the image that you use as a base for the query, the things we all talked about in terms of feature extraction for average color, color histogram, shape, texture, sketch, all of those things, the same algorithms would be used in order to extract those features from the image query? A. They're different algorithms. You would use an algorithm. Q. I'm sorry. Say that again? A. The algorithms are different between the shape features and the color features and sketch features. Q. Yeah. Understood. What I'm trying -- what I'm trying to understand is, the process that we talked about before, when you populate the database and you extract the features from the images that are stored in the database. You know, we went through all those different algorithms that are used for shape, color, both average color and histogram, texture and sketch.

Element	’529 Claim Recitation	Exemplary Citations in Reference
		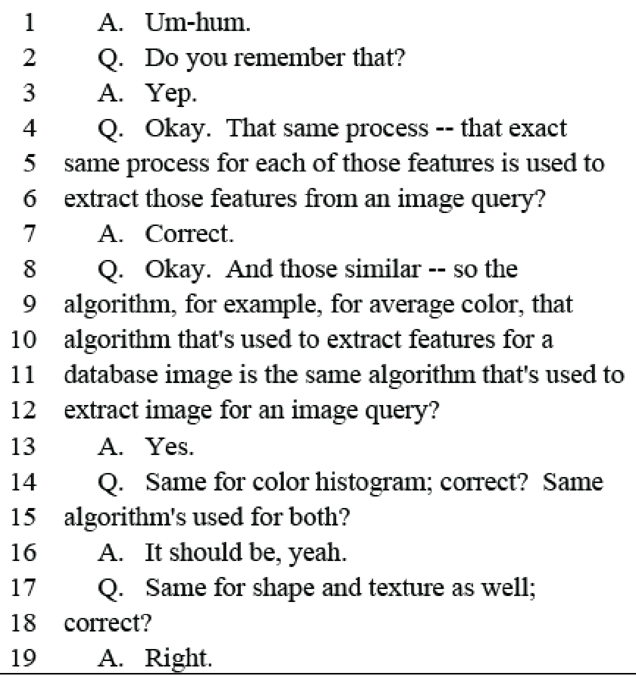 A. Um-hum. Q. Do you remember that? A. Yep. Q. Okay. That same process -- that exact same process for each of those features is used to extract those features from an image query? A. Correct. Q. Okay. And those similar -- so the algorithm, for example, for average color, that algorithm that's used to extract features for a database image is the same algorithm that's used to extract image for an image query? A. Yes. Q. Same for color histogram; correct? Same algorithm's used for both? A. It should be, yeah. Q. Same for shape and texture as well; correct? A. Right.

		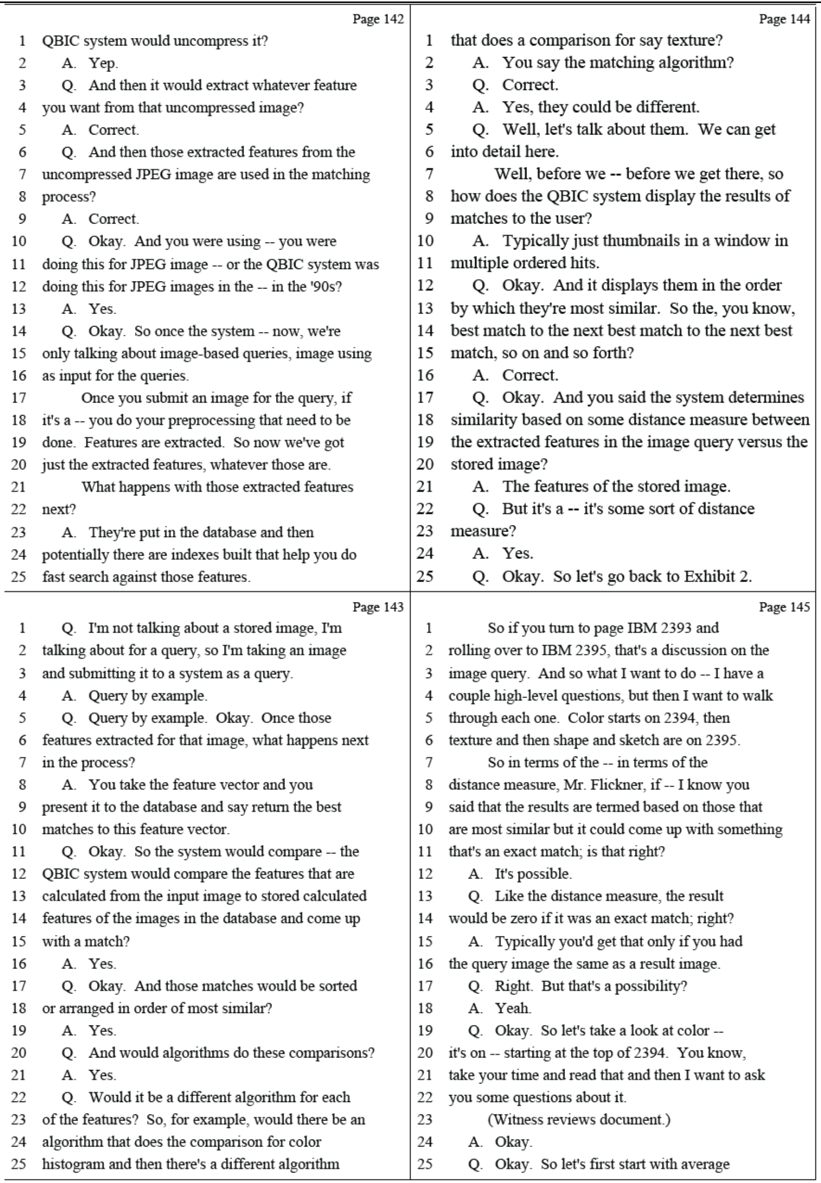

Page 142

QBIC system would uncompress it?

A. Yep.

Q. And then it would extract whatever feature you want from that uncompressed image?

A. Correct.

Q. And then those extracted features from the uncompressed JPEG image are used in the matching process?

A. Correct.

Q. Okay. And you were using -- you were doing this for JPEG image -- or the QBIC system was doing this for JPEG images in the -- in the '90s?

A. Yes.

Q. Okay. So once the system -- now, we're only talking about image-based queries, image using as input for the queries.

Once you submit an image for the query, if it's a -- you do your preprocessing that need to be done. Features are extracted. So now we've got just the extracted features, whatever those are.

What happens with those extracted features next?

A. They're put in the database and then potentially there are indexes built that help you do fast search against those features.

Page 143

Q. I'm not talking about a stored image, I'm talking about for a query, so I'm taking an image and submitting it to a system as a query.

A. Query by example.

Q. Query by example. Okay. Once those features extracted for that image, what happens next in the process?

A. You take the feature vector and you present it to the database and say return the best matches to this feature vector.

Q. Okay. So the system would compare -- the QBIC system would compare the features that are calculated from the input image to stored calculated features of the images in the database and come up with a match?

A. Yes.

Q. Okay. And those matches would be sorted or arranged in order of most similar?

A. Yes.

Q. And would algorithms do these comparisons?

A. Yes.

Q. Would it be a different algorithm for each of the features? So, for example, would there be an algorithm that does the comparison for color histogram and then there's a different algorithm

Page 144

that does a comparison for say texture?

A. You say the matching algorithm?

Q. Correct.

A. Yes, they could be different.

Q. Well, let's talk about them. We can get into detail here.

Well, before we -- before we get there, so how does the QBIC system display the results of matches to the user?

A. Typically just thumbnails in a window in multiple ordered hits.

Q. Okay. And it displays them in the order by which they're most similar. So the, you know, best match to the next best match to the next best match, so on and so forth?

A. Correct.

Q. Okay. And you said the system determines similarity based on some distance measure between the extracted features in the image query versus the stored image?

A. The features of the stored image.

Q. But it's a -- it's some sort of distance measure?

A. Yes.

Q. Okay. So let's go back to Exhibit 2.

Page 145

So if you turn to page IBM 2393 and rolling over to IBM 2395, that's a discussion on the image query. And so what I want to do -- I have a couple high-level questions, but then I want to walk through each one. Color starts on 2394, then texture and then shape and sketch are on 2395.

So in terms of the -- in terms of the distance measure, Mr. Flickner, if -- I know you said that the results are termed based on those that are most similar but it could come up with something that's an exact match; is that right?

A. It's possible.

Q. Like the distance measure, the result would be zero if it was an exact match; right?

A. Typically you'd get that only if you had the query image the same as a result image.

Q. Right. But that's a possibility?

A. Yeah.

Q. Okay. So let's take a look at color -- it's on -- starting at the top of 2394. You know, take your time and read that and then I want to ask you some questions about it.

(Witness reviews document.)

A. Okay.

Q. Okay. So let's first start with average

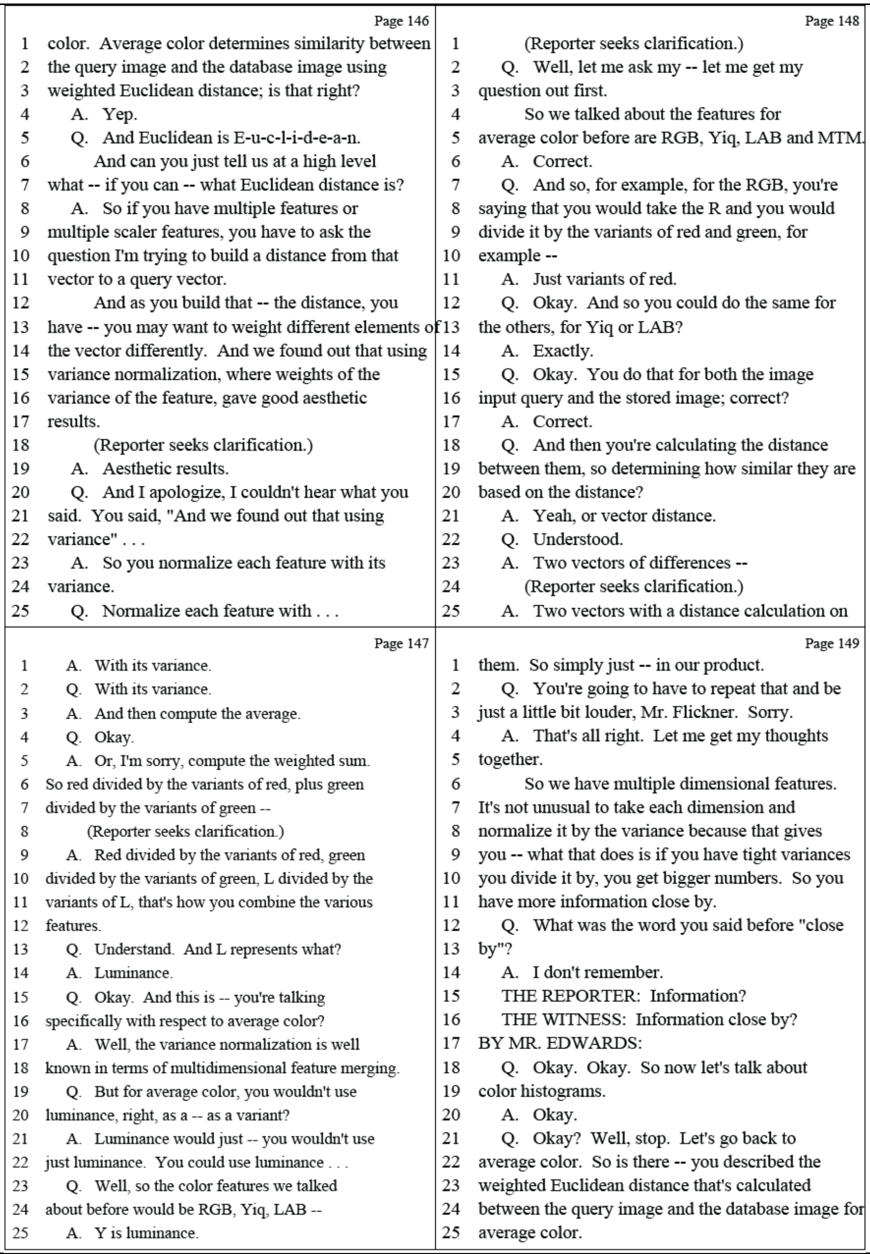

Page 146

color. Average color determines similarity between the query image and the database image using weighted Euclidean distance; is that right?

A. Yep.

Q. And Euclidean is E-u-c-l-i-d-e-a-n.

And can you just tell us at a high level what -- if you can -- what Euclidean distance is?

A. So if you have multiple features or multiple scaler features, you have to ask the question I'm trying to build a distance from that vector to a query vector.

And as you build that -- the distance, you have -- you may want to weight different elements of the vector differently. And we found out that using variance normalization, where weights of the variance of the feature, gave good aesthetic results.

(Reporter seeks clarification.)

A. Aesthetic results.

Q. And I apologize, I couldn't hear what you said. You said, "And we found out that using variance" . . .

A. So you normalize each feature with its variance.

Q. Normalize each feature with . . .

Page 147

A. With its variance.

Q. With its variance.

A. And then compute the average.

Q. Okay.

A. Or, I'm sorry, compute the weighted sum. So red divided by the variants of red, plus green divided by the variants of green --

(Reporter seeks clarification.)

A. Red divided by the variants of red, green divided by the variants of green, L divided by the variants of L, that's how you combine the various features.

Q. Understand. And L represents what?

A. Luminance.

Q. Okay. And this is -- you're talking specifically with respect to average color?

A. Well, the variance normalization is well known in terms of multidimensional feature merging.

Q. But for average color, you wouldn't use luminance, right, as a -- as a variant?

A. Luminance would just -- you wouldn't use just luminance. You could use luminance . . .

Q. Well, so the color features we talked about before would be RGB, Yiq, LAB --

A. Y is luminance.

Page 148

(Reporter seeks clarification.)

Q. Well, let me ask my -- let me get my question out first.

So we talked about the features for average color before are RGB, Yiq, LAB and MTM.

A. Correct.

Q. And so, for example, for the RGB, you're saying that you would take the R and you would divide it by the variants of red and green, for example --

A. Just variants of red.

Q. Okay. And so you could do the same for the others, for Yiq or LAB?

A. Exactly.

Q. Okay. You do that for both the image input query and the stored image; correct?

A. Correct.

Q. And then you're calculating the distance between them, so determining how similar they are based on the distance?

A. Yeah, or vector distance.

Q. Understood.

A. Two vectors of differences --

(Reporter seeks clarification.)

A. Two vectors with a distance calculation on

Page 149

them. So simply just -- in our product.

Q. You're going to have to repeat that and be just a little bit louder, Mr. Flickner. Sorry.

A. That's all right. Let me get my thoughts together.

So we have multiple dimensional features. It's not unusual to take each dimension and normalize it by the variance because that gives you -- what that does is if you have tight variances you divide it by, you get bigger numbers. So you have more information close by.

Q. What was the word you said before "close by"?

A. I don't remember.

THE REPORTER: Information?

THE WITNESS: Information close by?

BY MR. EDWARDS:

Q. Okay. Okay. So now let's talk about color histograms.

A. Okay.

Q. Okay? Well, stop. Let's go back to average color. So is there -- you described the weighted Euclidean distance that's calculated between the query image and the database image for average color.

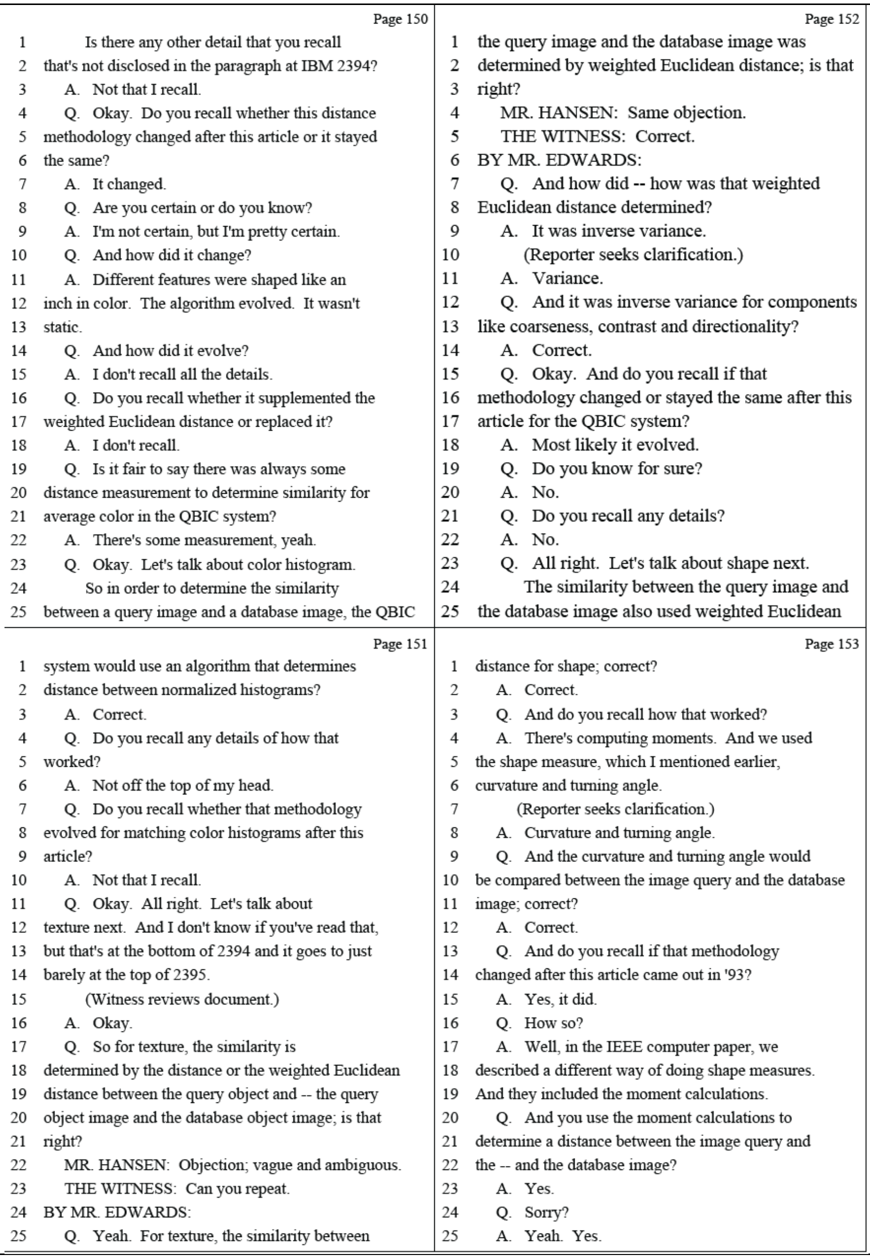

Page 150

Is there any other detail that you recall that's not disclosed in the paragraph at IBM 2394?
A. Not that I recall.
Q. Okay. Do you recall whether this distance methodology changed after this article or it stayed the same?
A. It changed.
Q. Are you certain or do you know?
A. I'm not certain, but I'm pretty certain.
Q. And how did it change?
A. Different features were shaped like an inch in color. The algorithm evolved. It wasn't static.
Q. And how did it evolve?
A. I don't recall all the details.
Q. Do you recall whether it supplemented the weighted Euclidean distance or replaced it?
A. I don't recall.
Q. Is it fair to say there was always some distance measurement to determine similarity for average color in the QBIC system?
A. There's some measurement, yeah.
Q. Okay. Let's talk about color histogram.
So in order to determine the similarity between a query image and a database image, the QBIC

Page 151

system would use an algorithm that determines distance between normalized histograms?
A. Correct.
Q. Do you recall any details of how that worked?
A. Not off the top of my head.
Q. Do you recall whether that methodology evolved for matching color histograms after this article?
A. Not that I recall.
Q. Okay. All right. Let's talk about texture next. And I don't know if you've read that, but that's at the bottom of 2394 and it goes to just barely at the top of 2395.
(Witness reviews document.)
A. Okay.
Q. So for texture, the similarity is determined by the distance or the weighted Euclidean distance between the query object and -- the query object image and the database object image; is that right?
MR. HANSEN: Objection; vague and ambiguous.
THE WITNESS: Can you repeat.
BY MR. EDWARDS:
Q. Yeah. For texture, the similarity between

Page 152

the query image and the database image was determined by weighted Euclidean distance; is that right?
MR. HANSEN: Same objection.
THE WITNESS: Correct.
BY MR. EDWARDS:
Q. And how did -- how was that weighted Euclidean distance determined?
A. It was inverse variance.
(Reporter seeks clarification.)
A. Variance.
Q. And it was inverse variance for components like coarseness, contrast and directionality?
A. Correct.
Q. Okay. And do you recall if that methodology changed or stayed the same after this article for the QBIC system?
A. Most likely it evolved.
Q. Do you know for sure?
A. No.
Q. Do you recall any details?
A. No.
Q. All right. Let's talk about shape next.
The similarity between the query image and the database image also used weighted Euclidean

Page 153

distance for shape; correct?
A. Correct.
Q. And do you recall how that worked?
A. There's computing moments. And we used the shape measure, which I mentioned earlier, curvature and turning angle.
(Reporter seeks clarification.)
A. Curvature and turning angle.
Q. And the curvature and turning angle would be compared between the image query and the database image; correct?
A. Correct.
Q. And do you recall if that methodology changed after this article came out in '93?
A. Yes, it did.
Q. How so?
A. Well, in the IEEE computer paper, we described a different way of doing shape measures. And they included the moment calculations.
Q. And you use the moment calculations to determine a distance between the image query and the -- and the database image?
A. Yes.
Q. Sorry?
A. Yeah. Yes.

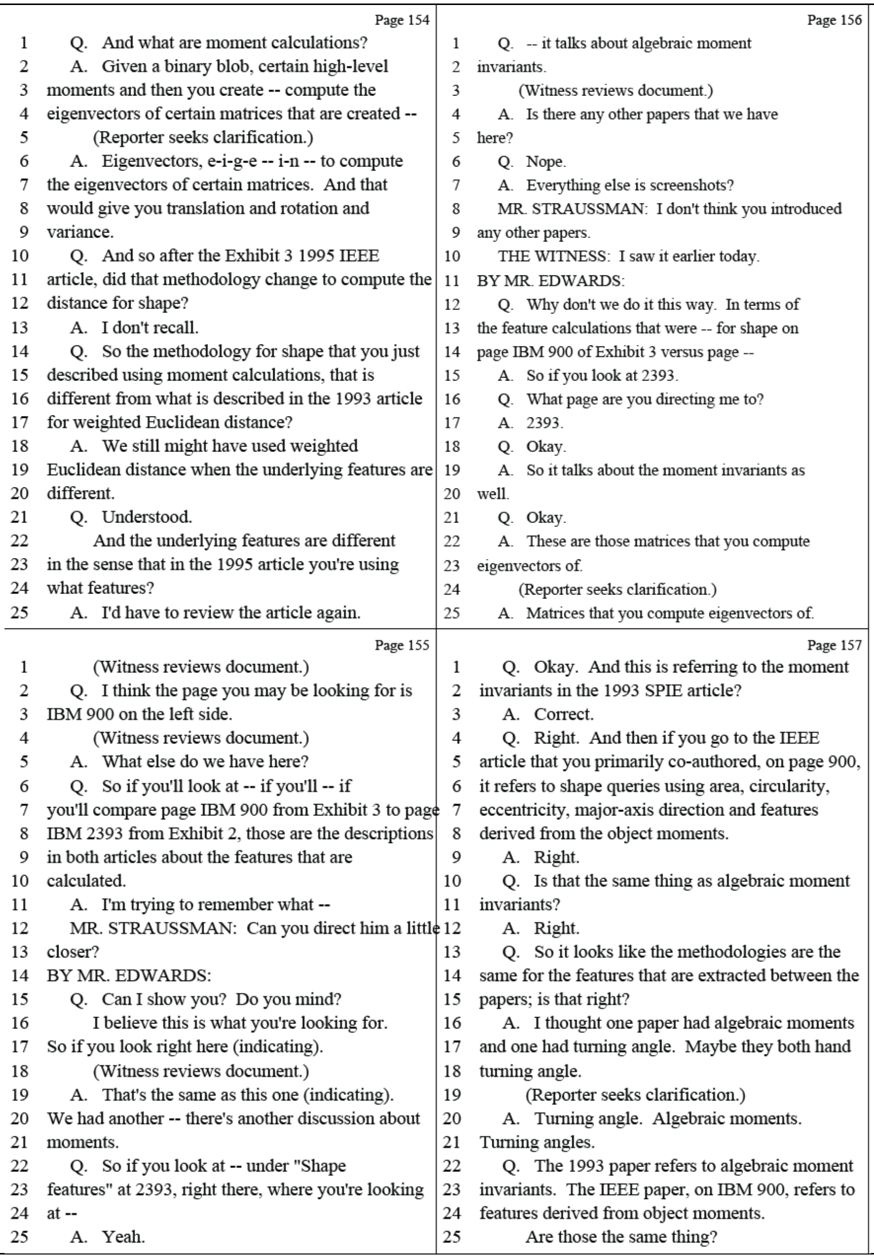

Page 154

Q. And what are moment calculations?
A. Given a binary blob, certain high-level moments and then you create -- compute the eigenvectors of certain matrices that are created --
(Reporter seeks clarification.)
A. Eigenvectors, e-i-g-e -- i-n -- to compute the eigenvectors of certain matrices. And that would give you translation and rotation and variance.
Q. And so after the Exhibit 3 1995 IEEE article, did that methodology change to compute the distance for shape?
A. I don't recall.
Q. So the methodology for shape that you just described using moment calculations, that is different from what is described in the 1993 article for weighted Euclidean distance?
A. We still might have used weighted Euclidean distance when the underlying features are different.
Q. Understood.
And the underlying features are different in the sense that in the 1995 article you're using what features?
A. I'd have to review the article again.

Page 155

(Witness reviews document.)
Q. I think the page you may be looking for is IBM 900 on the left side.
(Witness reviews document.)
A. What else do we have here?
Q. So if you'll look at -- if you'll -- if you'll compare page IBM 900 from Exhibit 3 to page IBM 2393 from Exhibit 2, those are the descriptions in both articles about the features that are calculated.
A. I'm trying to remember what --
MR. STRAUSSMAN: Can you direct him a little closer?
BY MR. EDWARDS:
Q. Can I show you? Do you mind?
I believe this is what you're looking for. So if you look right here (indicating).
(Witness reviews document.)
A. That's the same as this one (indicating). We had another -- there's another discussion about moments.
Q. So if you look at -- under "Shape features" at 2393, right there, where you're looking at --
A. Yeah.

Page 156

Q. -- it talks about algebraic moment invariants.
(Witness reviews document.)
A. Is there any other papers that we have here?
Q. Nope.
A. Everything else is screenshots?
MR. STRAUSSMAN: I don't think you introduced any other papers.
THE WITNESS: I saw it earlier today.
BY MR. EDWARDS:
Q. Why don't we do it this way. In terms of the feature calculations that were -- for shape on page IBM 900 of Exhibit 3 versus page --
A. So if you look at 2393.
Q. What page are you directing me to?
A. 2393.
Q. Okay.
A. So it talks about the moment invariants as well.
Q. Okay.
A. These are those matrices that you compute eigenvectors of.
(Reporter seeks clarification.)
A. Matrices that you compute eigenvectors of.

Page 157

Q. Okay. And this is referring to the moment invariants in the 1993 SPIE article?
A. Correct.
Q. Right. And then if you go to the IEEE article that you primarily co-authored, on page 900, it refers to shape queries using area, circularity, eccentricity, major-axis direction and features derived from the object moments.
A. Right.
Q. Is that the same thing as algebraic moment invariants?
A. Right.
Q. So it looks like the methodologies are the same for the features that are extracted between the papers; is that right?
A. I thought one paper had algebraic moments and one had turning angle. Maybe they both hand turning angle.
(Reporter seeks clarification.)
A. Turning angle. Algebraic moments. Turning angles.
Q. The 1993 paper refers to algebraic moment invariants. The IEEE paper, on IBM 900, refers to features derived from object moments.
Are those the same thing?

			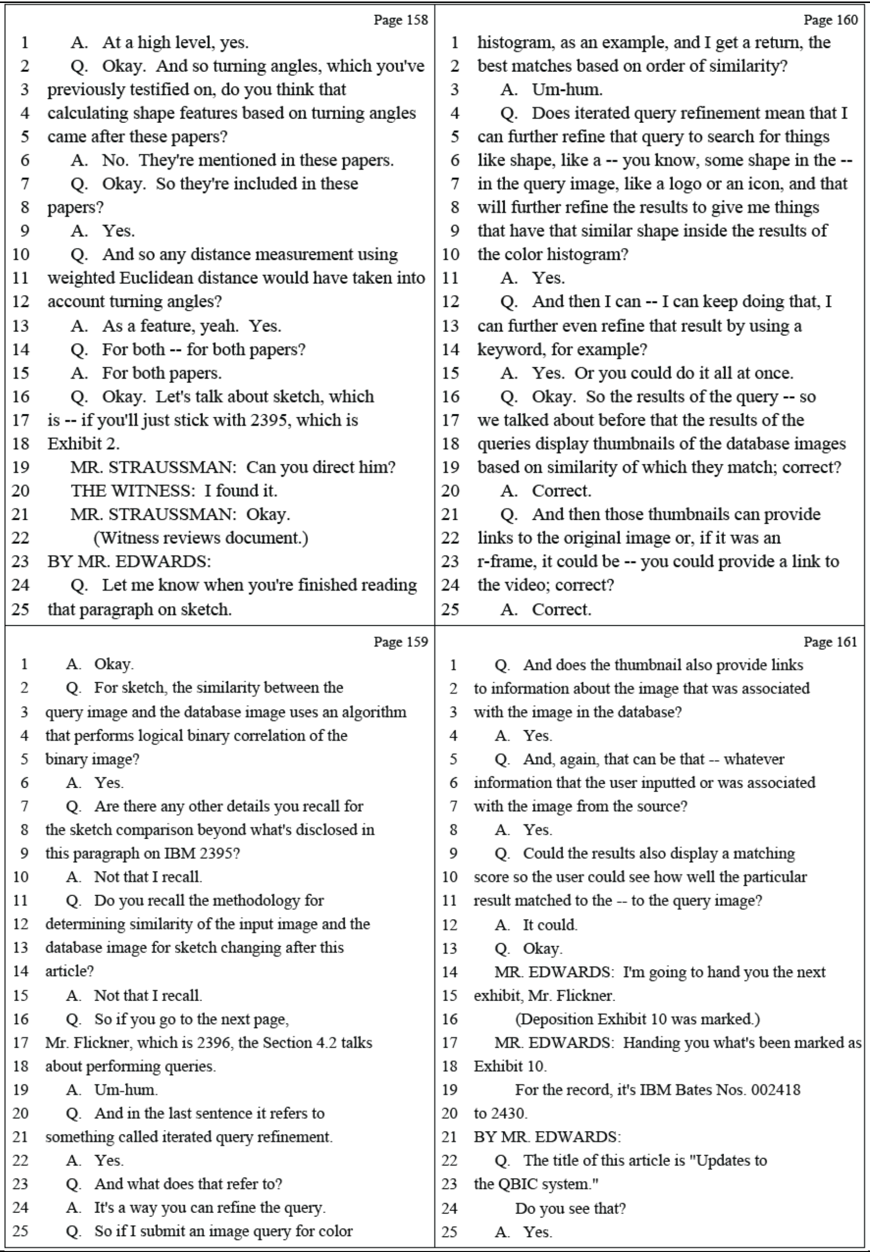 Page 158 A. At a high level, yes. Q. Okay. And so turning angles, which you've previously testified on, do you think that calculating shape features based on turning angles came after these papers? A. No. They're mentioned in these papers. Q. Okay. So they're included in these papers? A. Yes. Q. And so any distance measurement using weighted Euclidean distance would have taken into account turning angles? A. As a feature, yeah. Yes. Q. For both -- for both papers? A. For both papers. Q. Okay. Let's talk about sketch, which is -- if you'll just stick with 2395, which is Exhibit 2. MR. STRAUSSMAN: Can you direct him? THE WITNESS: I found it. MR. STRAUSSMAN: Okay. (Witness reviews document.) BY MR. EDWARDS: Q. Let me know when you're finished reading that paragraph on sketch. Page 159 A. Okay. Q. For sketch, the similarity between the query image and the database image uses an algorithm that performs logical binary correlation of the binary image? A. Yes. Q. Are there any other details you recall for the sketch comparison beyond what's disclosed in this paragraph on IBM 2395? A. Not that I recall. Q. Do you recall the methodology for determining similarity of the input image and the database image for sketch changing after this article? A. Not that I recall. Q. So if you go to the next page, Mr. Flickner, which is 2396, the Section 4.2 talks about performing queries. A. Um-hum. Q. And in the last sentence it refers to something called iterated query refinement. A. Yes. Q. And what does that refer to? A. It's a way you can refine the query. Q. So if I submit an image query for color histogram, as an example, and I get a return, the best matches based on order of similarity? Page 160 A. Um-hum. Q. Does iterated query refinement mean that I can further refine that query to search for things like shape, like a -- you know, some shape in the -- in the query image, like a logo or an icon, and that will further refine the results to give me things that have that similar shape inside the results of the color histogram? A. Yes. Q. And then I can -- I can keep doing that, I can further even refine that result by using a keyword, for example? A. Yes. Or you could do it all at once. Q. Okay. So the results of the query -- so we talked about before that the results of the queries display thumbnails of the database images based on similarity of which they match; correct? A. Correct. Q. And then those thumbnails can provide links to the original image or, if it was an r-frame, it could be -- you could provide a link to the video; correct? A. Correct. Page 161 Q. And does the thumbnail also provide links to information about the image that was associated with the image in the database? A. Yes. Q. And, again, that can be that -- whatever information that the user inputted or was associated with the image from the source? A. Yes. Q. Could the results also display a matching score so the user could see how well the particular result matched to the -- to the query image? A. It could. Q. Okay. MR. EDWARDS: I'm going to hand you the next exhibit, Mr. Flickner. (Deposition Exhibit 10 was marked.) MR. EDWARDS: Handing you what's been marked as Exhibit 10. For the record, it's IBM Bates Nos. 002418 to 2430. BY MR. EDWARDS: Q. The title of this article is "Updates to the QBIC system." Do you see that? A. Yes.

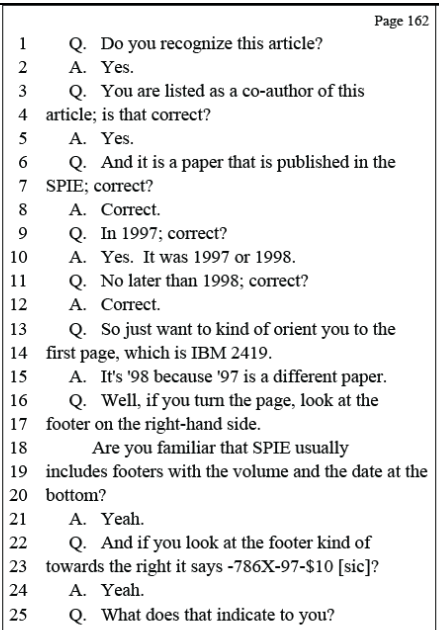

Page 162

Q. Do you recognize this article?
A. Yes.
Q. You are listed as a co-author of this article; is that correct?
A. Yes.
Q. And it is a paper that is published in the SPIE; correct?
A. Correct.
Q. In 1997; correct?
A. Yes. It was 1997 or 1998.
Q. No later than 1998; correct?
A. Correct.
Q. So just want to kind of orient you to the first page, which is IBM 2419.
A. It's '98 because '97 is a different paper.
Q. Well, if you turn the page, look at the footer on the right-hand side.
Are you familiar that SPIE usually includes footers with the volume and the date at the bottom?
A. Yeah.
Q. And if you look at the footer kind of towards the right it says -786X-97-$10 [sic]?
A. Yeah.
Q. What does that indicate to you?

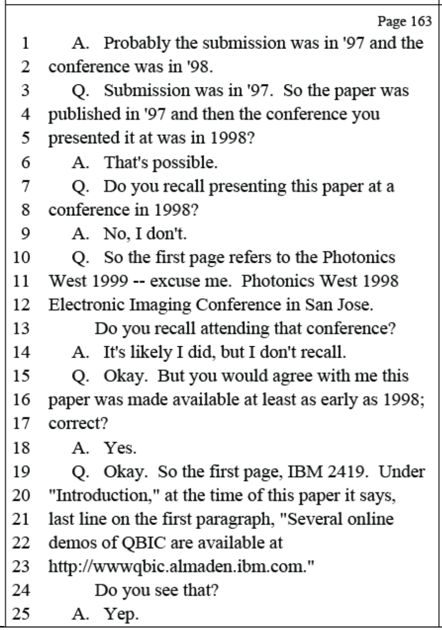

Page 163

A. Probably the submission was in '97 and the conference was in '98.
Q. Submission was in '97. So the paper was published in '97 and then the conference you presented it at was in 1998?
A. That's possible.
Q. Do you recall presenting this paper at a conference in 1998?
A. No, I don't.
Q. So the first page refers to the Photonics West 1999 -- excuse me. Photonics West 1998 Electronic Imaging Conference in San Jose.
Do you recall attending that conference?
A. It's likely I did, but I don't recall.
Q. Okay. But you would agree with me this paper was made available at least as early as 1998; correct?
A. Yes.
Q. Okay. So the first page, IBM 2419. Under "Introduction," at the time of this paper it says, last line on the first paragraph, "Several online demos of QBIC are available at http://wwwqbic.almaden.ibm.com."
Do you see that?
A. Yep.

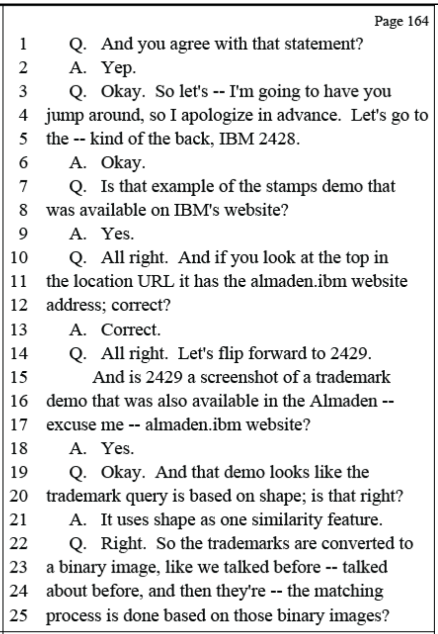

Page 164

Q. And you agree with that statement?
A. Yep.
Q. Okay. So let's -- I'm going to have you jump around, so I apologize in advance. Let's go to the -- kind of the back, IBM 2428.
A. Okay.
Q. Is that example of the stamps demo that was available on IBM's website?
A. Yes.
Q. All right. And if you look at the top in the location URL it has the almaden.ibm website address; correct?
A. Correct.
Q. All right. Let's flip forward to 2429.
And is 2429 a screenshot of a trademark demo that was also available in the Almaden -- excuse me -- almaden.ibm website?
A. Yes.
Q. Okay. And that demo looks like the trademark query is based on shape; is that right?
A. It uses shape as one similarity feature.
Q. Right. So the trademarks are converted to a binary image, like we talked before -- talked about before, and then they're -- the matching process is done based on those binary images?

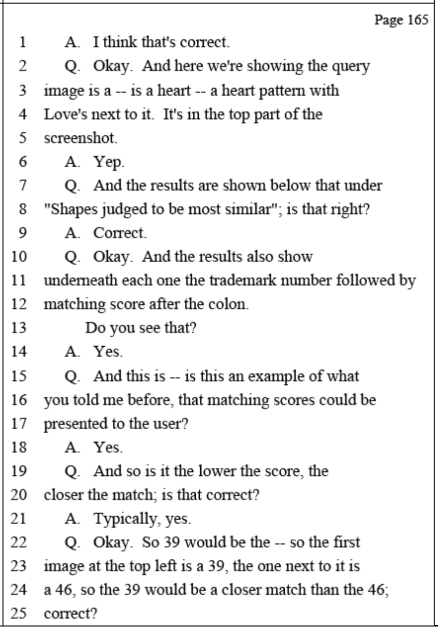

Page 165

A. I think that's correct.
Q. Okay. And here we're showing the query image is a -- is a heart -- a heart pattern with Love's next to it. It's in the top part of the screenshot.
A. Yep.
Q. And the results are shown below that under "Shapes judged to be most similar"; is that right?
A. Correct.
Q. Okay. And the results also show underneath each one the trademark number followed by matching score after the colon.
Do you see that?
A. Yes.
Q. And this is -- is this an example of what you told me before, that matching scores could be presented to the user?
A. Yes.
Q. And so is it the lower the score, the closer the match; is that correct?
A. Typically, yes.
Q. Okay. So 39 would be the -- so the first image at the top left is a 39, the one next to it is a 46, so the 39 would be a closer match than the 46; correct?

Element	'529 Claim Recitation	Exemplary Citations in Reference
		Id. at 174–182: 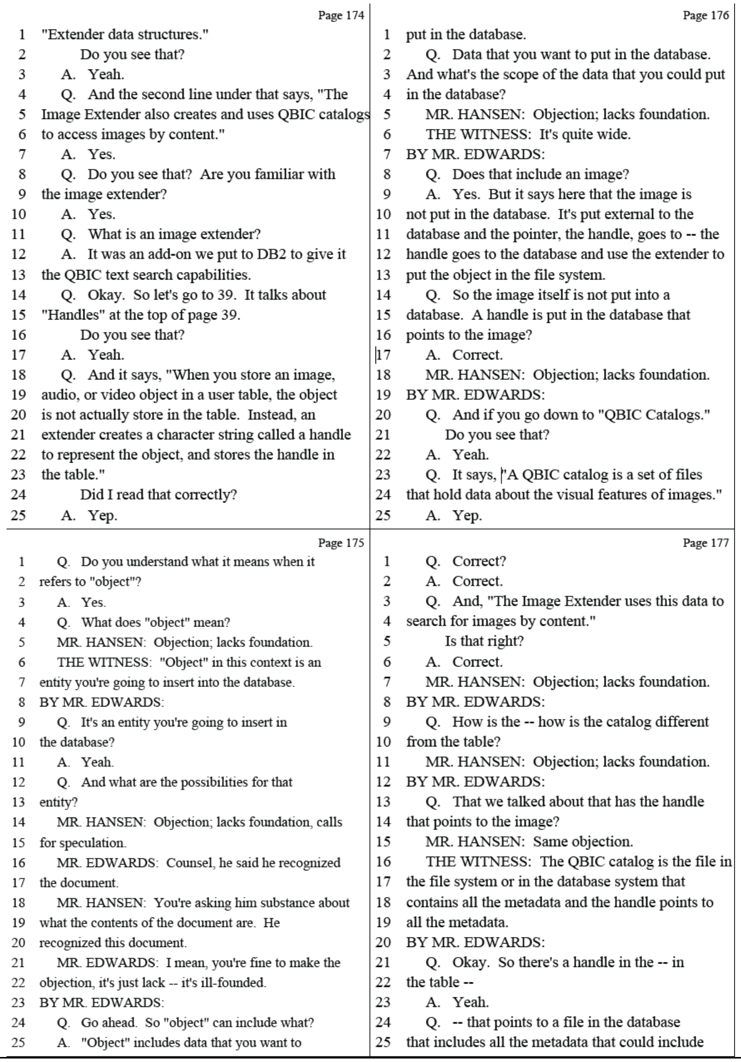 Page 174 "Extender data structures." Do you see that? A. Yeah. Q. And the second line under that says, "The Image Extender also creates and uses QBIC catalogs to access images by content." A. Yes. Q. Do you see that? Are you familiar with the image extender? A. Yes. Q. What is an image extender? A. It was an add-on we put to DB2 to give it the QBIC text search capabilities. Q. Okay. So let's go to 39. It talks about "Handles" at the top of page 39. Do you see that? A. Yeah. Q. And it says, "When you store an image, audio, or video object in a user table, the object is not actually store in the table. Instead, an extender creates a character string called a handle to represent the object, and stores the handle in the table." Did I read that correctly? A. Yep. Page 175 Q. Do you understand what it means when it refers to "object"? A. Yes. Q. What does "object" mean? MR. HANSEN: Objection; lacks foundation. THE WITNESS: "Object" in this context is an entity you're going to insert into the database. BY MR. EDWARDS: Q. It's an entity you're going to insert in the database? A. Yeah. Q. And what are the possibilities for that entity? MR. HANSEN: Objection; lacks foundation, calls for speculation. MR. EDWARDS: Counsel, he said he recognized the document. MR. HANSEN: You're asking him substance about what the contents of the document are. He recognized this document. MR. EDWARDS: I mean, you're fine to make the objection, it's just lack -- it's ill-founded. BY MR. EDWARDS: Q. Go ahead. So "object" can include what? A. "Object" includes data that you want to Page 176 put in the database. Q. Data that you want to put in the database. And what's the scope of the data that you could put in the database? MR. HANSEN: Objection; lacks foundation. THE WITNESS: It's quite wide. BY MR. EDWARDS: Q. Does that include an image? A. Yes. But it says here that the image is not put in the database. It's put external to the database and the pointer, the handle, goes to -- the handle goes to the database and use the extender to put the object in the file system. Q. So the image itself is not put into a database. A handle is put in the database that points to the image? A. Correct. MR. HANSEN: Objection; lacks foundation. BY MR. EDWARDS: Q. And if you go down to "QBIC Catalogs." Do you see that? A. Yeah. Q. It says, "A QBIC catalog is a set of files that hold data about the visual features of images." A. Yep. Page 177 Q. Correct? A. Correct. Q. And, "The Image Extender uses this data to search for images by content." Is that right? A. Correct. MR. HANSEN: Objection; lacks foundation. BY MR. EDWARDS: Q. How is the -- how is the catalog different from the table? MR. HANSEN: Objection; lacks foundation. BY MR. EDWARDS: Q. That we talked about that has the handle that points to the image? MR. HANSEN: Same objection. THE WITNESS: The QBIC catalog is the file in the file system or in the database system that contains all the metadata and the handle points to all the metadata. BY MR. EDWARDS: Q. Okay. So there's a handle in the -- in the table -- A. Yeah. Q. -- that points to a file in the database that includes all the metadata that could include

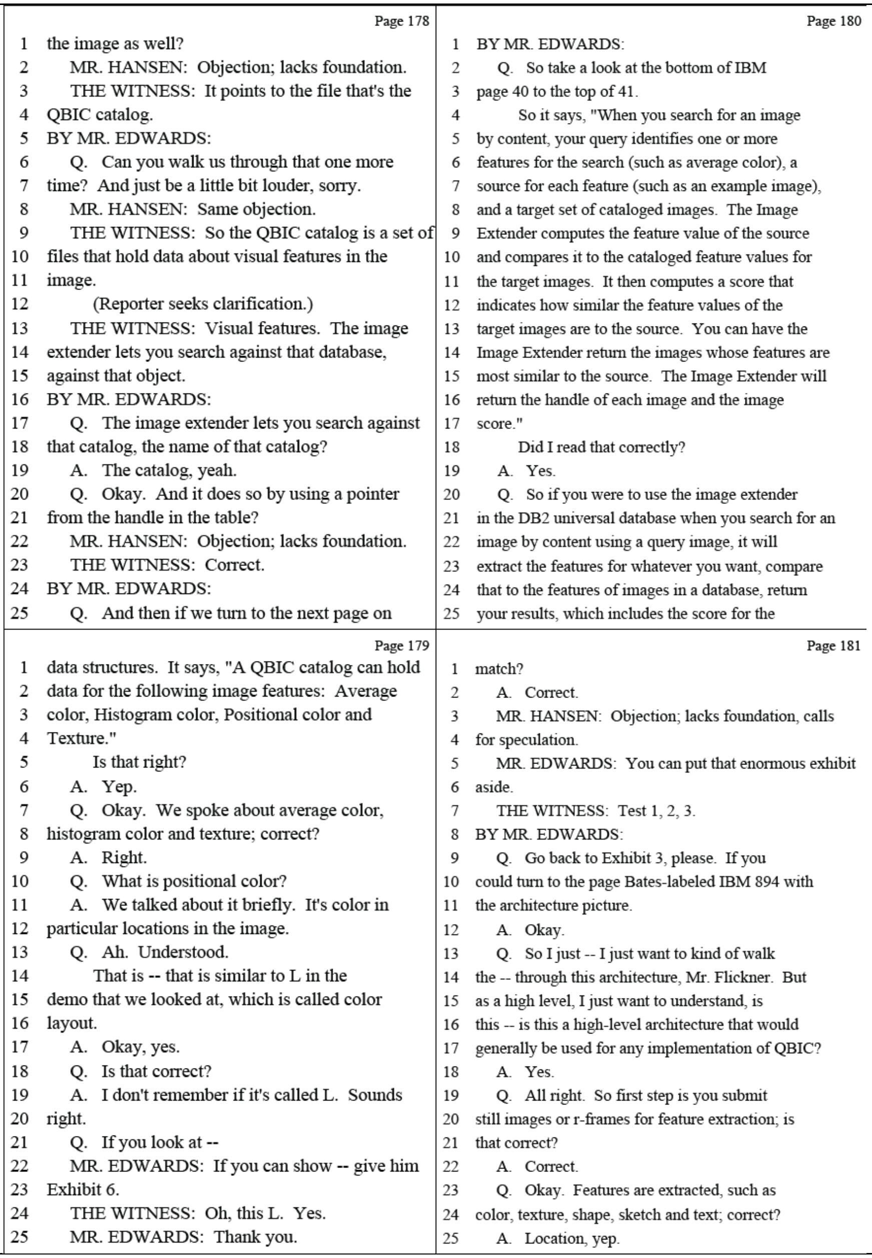

Page 178

the image as well?
MR. HANSEN: Objection; lacks foundation.
THE WITNESS: It points to the file that's the QBIC catalog.
BY MR. EDWARDS:
Q. Can you walk us through that one more time? And just be a little bit louder, sorry.
MR. HANSEN: Same objection.
THE WITNESS: So the QBIC catalog is a set of files that hold data about visual features in the image.
(Reporter seeks clarification.)
THE WITNESS: Visual features. The image extender lets you search against that database, against that object.
BY MR. EDWARDS:
Q. The image extender lets you search against that catalog, the name of that catalog?
A. The catalog, yeah.
Q. Okay. And it does so by using a pointer from the handle in the table?
MR. HANSEN: Objection; lacks foundation.
THE WITNESS: Correct.
BY MR. EDWARDS:
Q. And then if we turn to the next page on

Page 179

data structures. It says, "A QBIC catalog can hold data for the following image features: Average color, Histogram color, Positional color and Texture."
Is that right?
A. Yep.
Q. Okay. We spoke about average color, histogram color and texture; correct?
A. Right.
Q. What is positional color?
A. We talked about it briefly. It's color in particular locations in the image.
Q. Ah. Understood.
That is -- that is similar to L in the demo that we looked at, which is called color layout.
A. Okay, yes.
Q. Is that correct?
A. I don't remember if it's called L. Sounds right.
Q. If you look at --
MR. EDWARDS: If you can show -- give him Exhibit 6.
THE WITNESS: Oh, this L. Yes.
MR. EDWARDS: Thank you.

Page 180

BY MR. EDWARDS:
Q. So take a look at the bottom of IBM page 40 to the top of 41.
So it says, "When you search for an image by content, your query identifies one or more features for the search (such as average color), a source for each feature (such as an example image), and a target set of cataloged images. The Image Extender computes the feature value of the source and compares it to the cataloged feature values for the target images. It then computes a score that indicates how similar the feature values of the target images are to the source. You can have the Image Extender return the images whose features are most similar to the source. The Image Extender will return the handle of each image and the image score."
Did I read that correctly?
A. Yes.
Q. So if you were to use the image extender in the DB2 universal database when you search for an image by content using a query image, it will extract the features for whatever you want, compare that to the features of images in a database, return your results, which includes the score for the

Page 181

match?
A. Correct.
MR. HANSEN: Objection; lacks foundation, calls for speculation.
MR. EDWARDS: You can put that enormous exhibit aside.
THE WITNESS: Test 1, 2, 3.
BY MR. EDWARDS:
Q. Go back to Exhibit 3, please. If you could turn to the page Bates-labeled IBM 894 with the architecture picture.
A. Okay.
Q. So I just -- I just want to kind of walk the -- through this architecture, Mr. Flickner. But as a high level, I just want to understand, is this -- is this a high-level architecture that would generally be used for any implementation of QBIC?
A. Yes.
Q. All right. So first step is you submit still images or r-frames for feature extraction; is that correct?
A. Correct.
Q. Okay. Features are extracted, such as color, texture, shape, sketch and text; correct?
A. Location, yep.

Element	'529 Claim Recitation	Exemplary Citations in Reference
		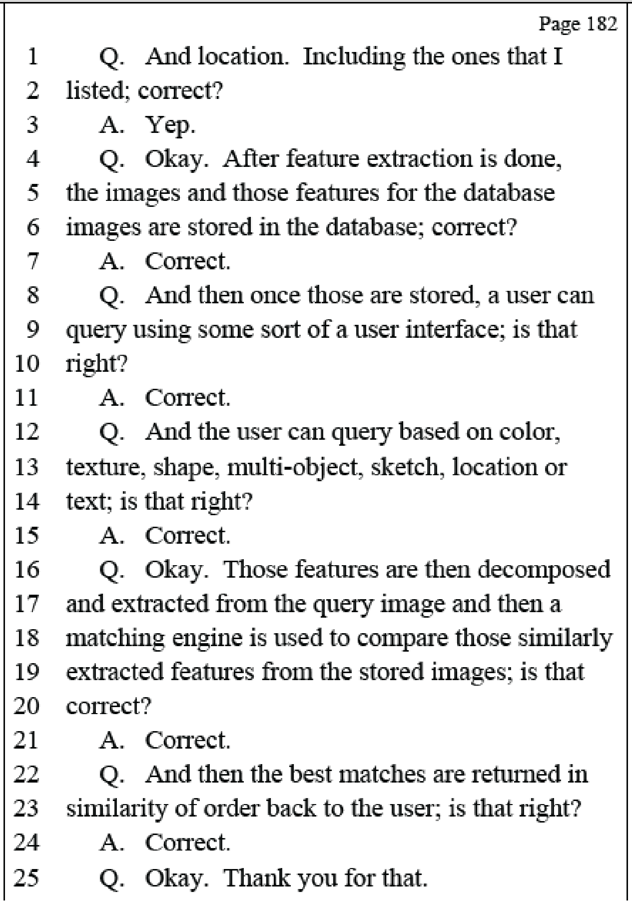 Page 182 Q. And location. Including the ones that I listed; correct? A. Yep. Q. Okay. After feature extraction is done, the images and those features for the database images are stored in the database; correct? A. Correct. Q. And then once those are stored, a user can query using some sort of a user interface; is that right? A. Correct. Q. And the user can query based on color, texture, shape, multi-object, sketch, location or text; is that right? A. Correct. Q. Okay. Those features are then decomposed and extracted from the query image and then a matching engine is used to compare those similarly extracted features from the stored images; is that correct? A. Correct. Q. And then the best matches are returned in similarity of order back to the user; is that right? A. Correct. Q. Okay. Thank you for that. *Id.* at 190–191: 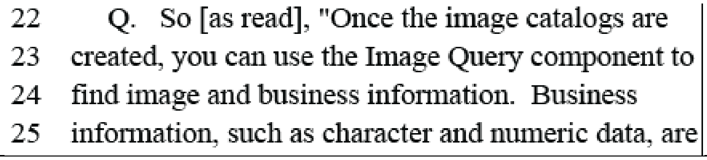 Q. So [as read], "Once the image catalogs are created, you can use the Image Query component to find image and business information. Business information, such as character and numeric data, are

		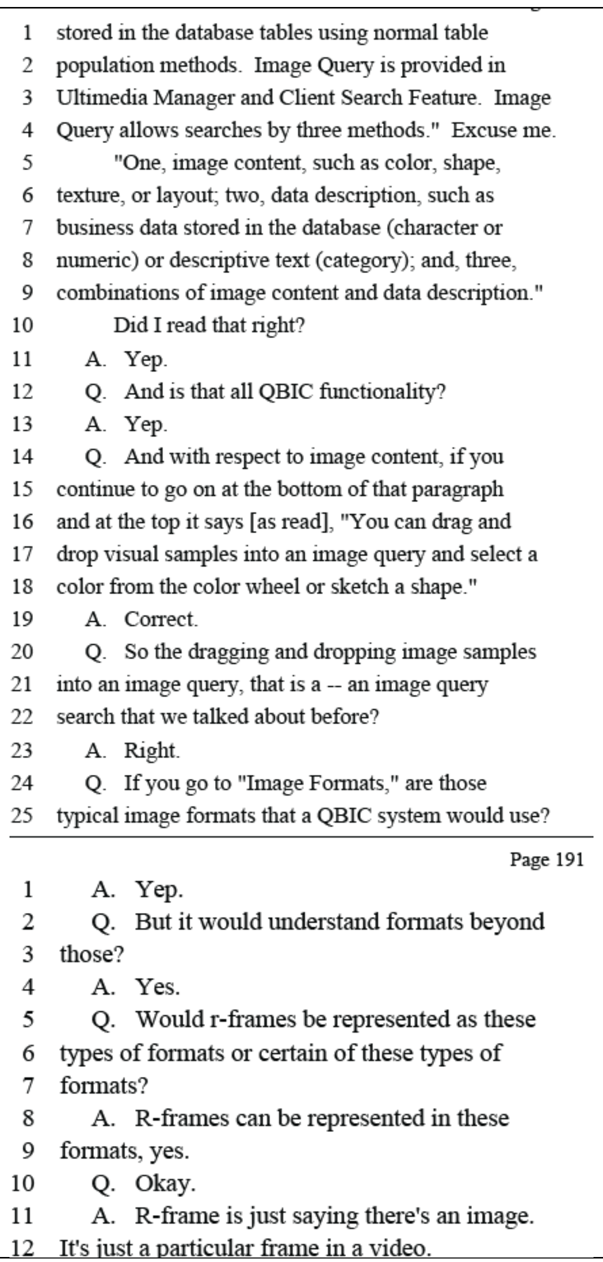 stored in the database tables using normal table population methods. Image Query is provided in Ultimedia Manager and Client Search Feature. Image Query allows searches by three methods." Excuse me. "One, image content, such as color, shape, texture, or layout; two, data description, such as business data stored in the database (character or numeric) or descriptive text (category); and, three, combinations of image content and data description." Did I read that right? A. Yep. Q. And is that all QBIC functionality? A. Yep. Q. And with respect to image content, if you continue to go on at the bottom of that paragraph and at the top it says [as read], "You can drag and drop visual samples into an image query and select a color from the color wheel or sketch a shape." A. Correct. Q. So the dragging and dropping image samples into an image query, that is a -- an image query search that we talked about before? A. Right. Q. If you go to "Image Formats," are those typical image formats that a QBIC system would use? Page 191 A. Yep. Q. But it would understand formats beyond those? A. Yes. Q. Would r-frames be represented as these types of formats or certain of these types of formats? A. R-frames can be represented in these formats, yes. Q. Okay. A. R-frame is just saying there's an image. It's just a particular frame in a video.

Element	'529 Claim Recitation	Exemplary Citations in Reference
		Id. at 198 – 199: 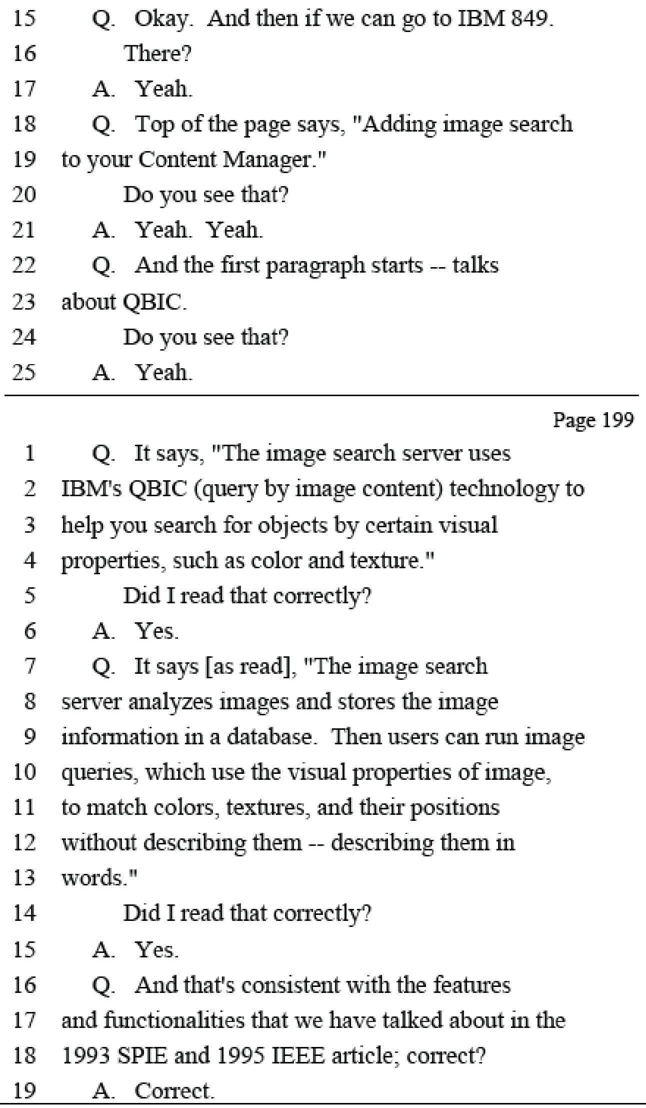 Q. Okay. And then if we can go to IBM 849. There? A. Yeah. Q. Top of the page says, "Adding image search to your Content Manager." Do you see that? A. Yeah. Yeah. Q. And the first paragraph starts -- talks about QBIC. Do you see that? A. Yeah. Page 199 Q. It says, "The image search server uses IBM's QBIC (query by image content) technology to help you search for objects by certain visual properties, such as color and texture." Did I read that correctly? A. Yes. Q. It says [as read], "The image search server analyzes images and stores the image information in a database. Then users can run image queries, which use the visual properties of image, to match colors, textures, and their positions without describing them -- describing them in words." Did I read that correctly? A. Yes. Q. And that's consistent with the features and functionalities that we have talked about in the 1993 SPIE and 1995 IEEE article; correct? A. Correct.

Element	'529 Claim Recitation	Exemplary Citations in Reference
		Id. at 305: 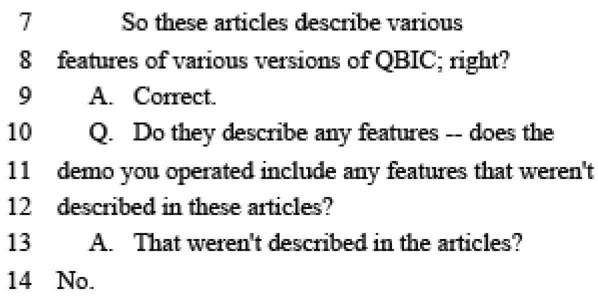 So these articles describe various features of various versions of QBIC; right? A. Correct. Q. Do they describe any features -- does the demo you operated include any features that weren't described in these articles? A. That weren't described in the articles? No. *Id*. at 35: 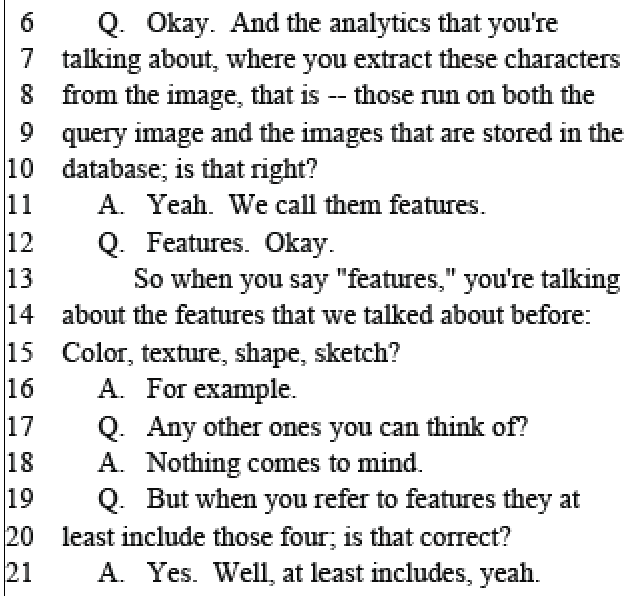 Q. Okay. And the analytics that you're talking about, where you extract these characters from the image, that is -- those run on both the query image and the images that are stored in the database; is that right? A. Yeah. We call them features. Q. Features. Okay. So when you say "features," you're talking about the features that we talked about before: Color, texture, shape, sketch? A. For example. Q. Any other ones you can think of? A. Nothing comes to mind. Q. But when you refer to features they at least include those four; is that correct? A. Yes. Well, at least includes, yeah.

Element	'529 Claim Recitation	Exemplary Citations in Reference
		Id. at 109-110: 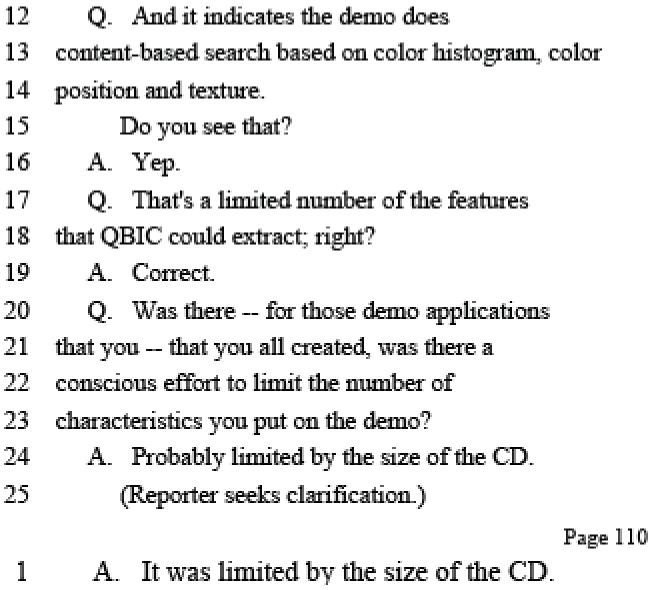 Q. And it indicates the demo does content-based search based on color histogram, color position and texture. Do you see that? A. Yep. Q. That's a limited number of the features that QBIC could extract; right? A. Correct. Q. Was there -- for those demo applications that you -- that you all created, was there a conscious effort to limit the number of characteristics you put on the demo? A. Probably limited by the size of the CD. (Reporter seeks clarification.) Page 110 A. It was limited by the size of the CD. *Id.* at 311: 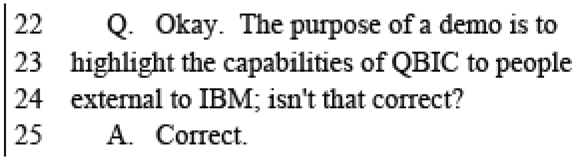 Q. Okay. The purpose of a demo is to highlight the capabilities of QBIC to people external to IBM; isn't that correct? A. Correct. *Id.* at 312 – 313:

Element	'529 Claim Recitation	Exemplary Citations in Reference
		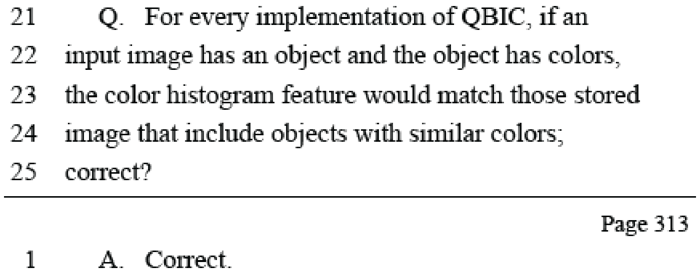 Q. For every implementation of QBIC, if an input image has an object and the object has colors, the color histogram feature would match those stored image that include objects with similar colors; correct? Page 313 A. Correct. *Id.* at 313-317:

Element	'529 Claim Recitation	Exemplary Citations in Reference
		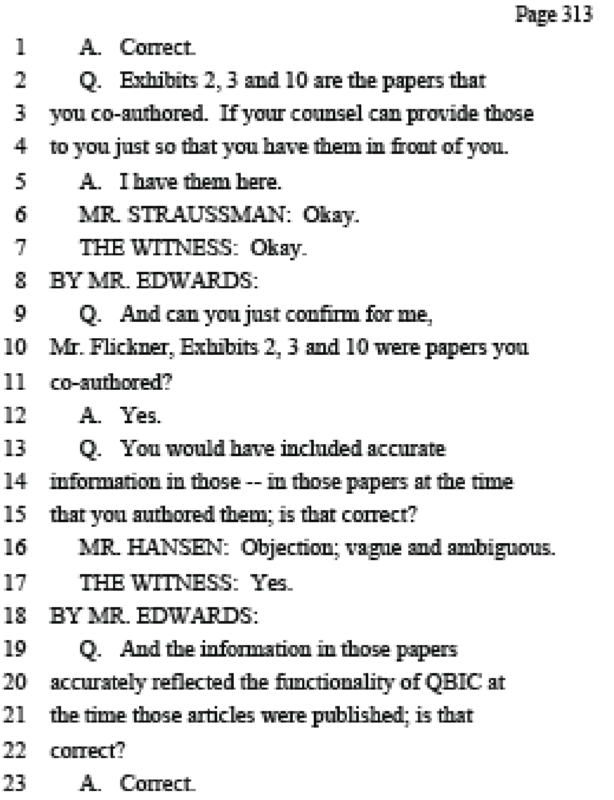 Page 313 A. Correct. Q. Exhibits 2, 3 and 10 are the papers that you co-authored. If your counsel can provide those to you just so that you have them in front of you. A. I have them here. MR. STRAUSSMAN: Okay. THE WITNESS: Okay. BY MR. EDWARDS: Q. And can you just confirm for me, Mr. Flickner, Exhibits 2, 3 and 10 were papers you co-authored? A. Yes. Q. You would have included accurate information in those -- in those papers at the time that you authored them; is that correct? MR. HANSEN: Objection; vague and ambiguous. THE WITNESS: Yes. BY MR. EDWARDS: Q. And the information in those papers accurately reflected the functionality of QBIC at the time those articles were published; is that correct? A. Correct.

		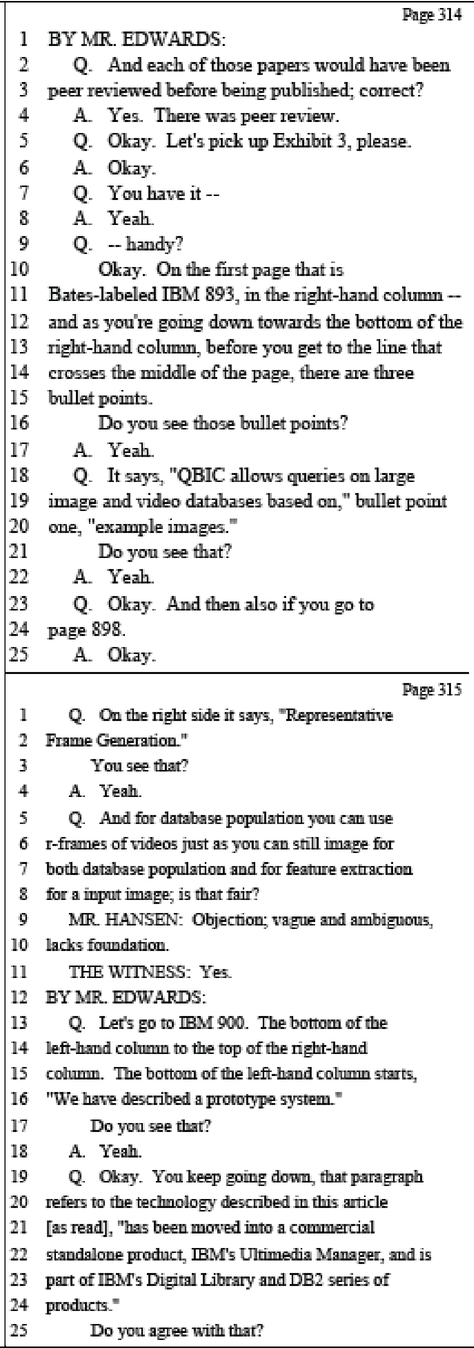Page 314 BY MR. EDWARDS: Q. And each of those papers would have been peer reviewed before being published; correct? A. Yes. There was peer review. Q. Okay. Let's pick up Exhibit 3, please. A. Okay. Q. You have it -- A. Yeah. Q. -- handy? Okay. On the first page that is Bates-labeled IBM 893, in the right-hand column -- and as you're going down towards the bottom of the right-hand column, before you get to the line that crosses the middle of the page, there are three bullet points. Do you see those bullet points? A. Yeah. Q. It says, "QBIC allows queries on large image and video databases based on," bullet point one, "example images." Do you see that? A. Yeah. Q. Okay. And then also if you go to page 898. A. Okay. Page 315 Q. On the right side it says, "Representative Frame Generation." You see that? A. Yeah. Q. And for database population you can use r-frames of videos just as you can still image for both database population and for feature extraction for a input image; is that fair? MR. HANSEN: Objection; vague and ambiguous, lacks foundation. THE WITNESS: Yes. BY MR. EDWARDS: Q. Let's go to IBM 900. The bottom of the left-hand column to the top of the right-hand column. The bottom of the left-hand column starts, "We have described a prototype system." Do you see that? A. Yeah. Q. Okay. You keep going down, that paragraph refers to the technology described in this article [as read], "has been moved into a commercial standalone product, IBM's Ultimedia Manager, and is part of IBM's Digital Library and DB2 series of products." Do you agree with that?

Element	'529 Claim Recitation	Exemplary Citations in Reference
		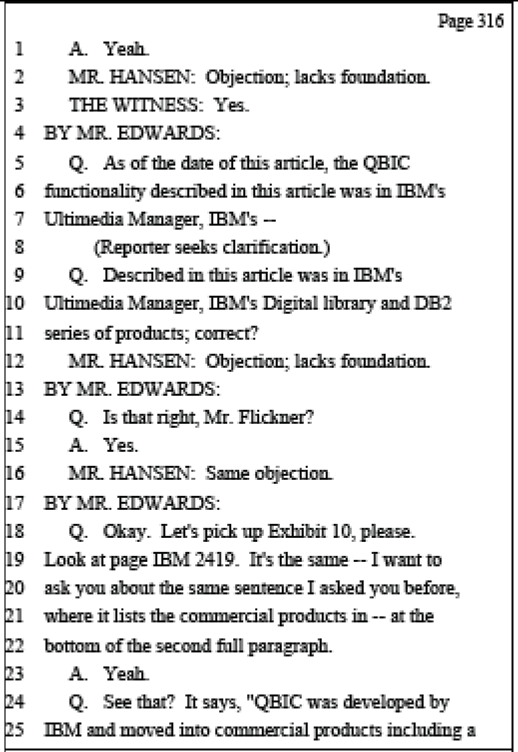 Page 316 A. Yeah. MR. HANSEN: Objection; lacks foundation. THE WITNESS: Yes. BY MR. EDWARDS: Q. As of the date of this article, the QBIC functionality described in this article was in IBM's Ultimedia Manager, IBM's -- (Reporter seeks clarification.) Q. Described in this article was in IBM's Ultimedia Manager, IBM's Digital library and DB2 series of products; correct? MR. HANSEN: Objection; lacks foundation. BY MR. EDWARDS: Q. Is that right, Mr. Flickner? A. Yes. MR. HANSEN: Same objection. BY MR. EDWARDS: Q. Okay. Let's pick up Exhibit 10, please. Look at page IBM 2419. It's the same -- I want to ask you about the same sentence I asked you before, where it lists the commercial products in -- at the bottom of the second full paragraph. A. Yeah. Q. See that? It says, "QBIC was developed by IBM and moved into commercial products including a 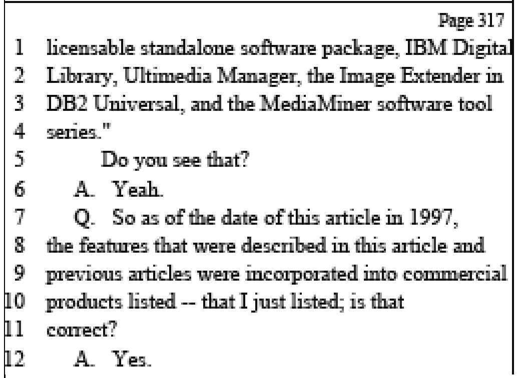 Page 317 licensable standalone software package, IBM Digital Library, Ultimedia Manager, the Image Extender in DB2 Universal, and the MediaMiner software tool series." Do you see that? A. Yeah. Q. So as of the date of this article in 1997, the features that were described in this article and previous articles were incorporated into commercial products listed -- that I just listed; is that correct? A. Yes.

Element	'529 Claim Recitation	Exemplary Citations in Reference
		Id. at 317–322: 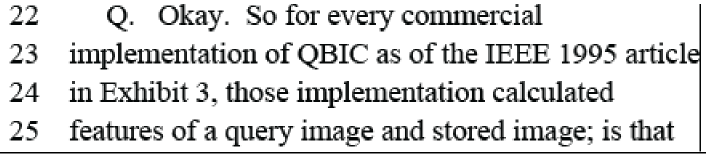 Q. Okay. So for every commercial implementation of QBIC as of the IEEE 1995 article in Exhibit 3, those implementation calculated features of a query image and stored image; is that

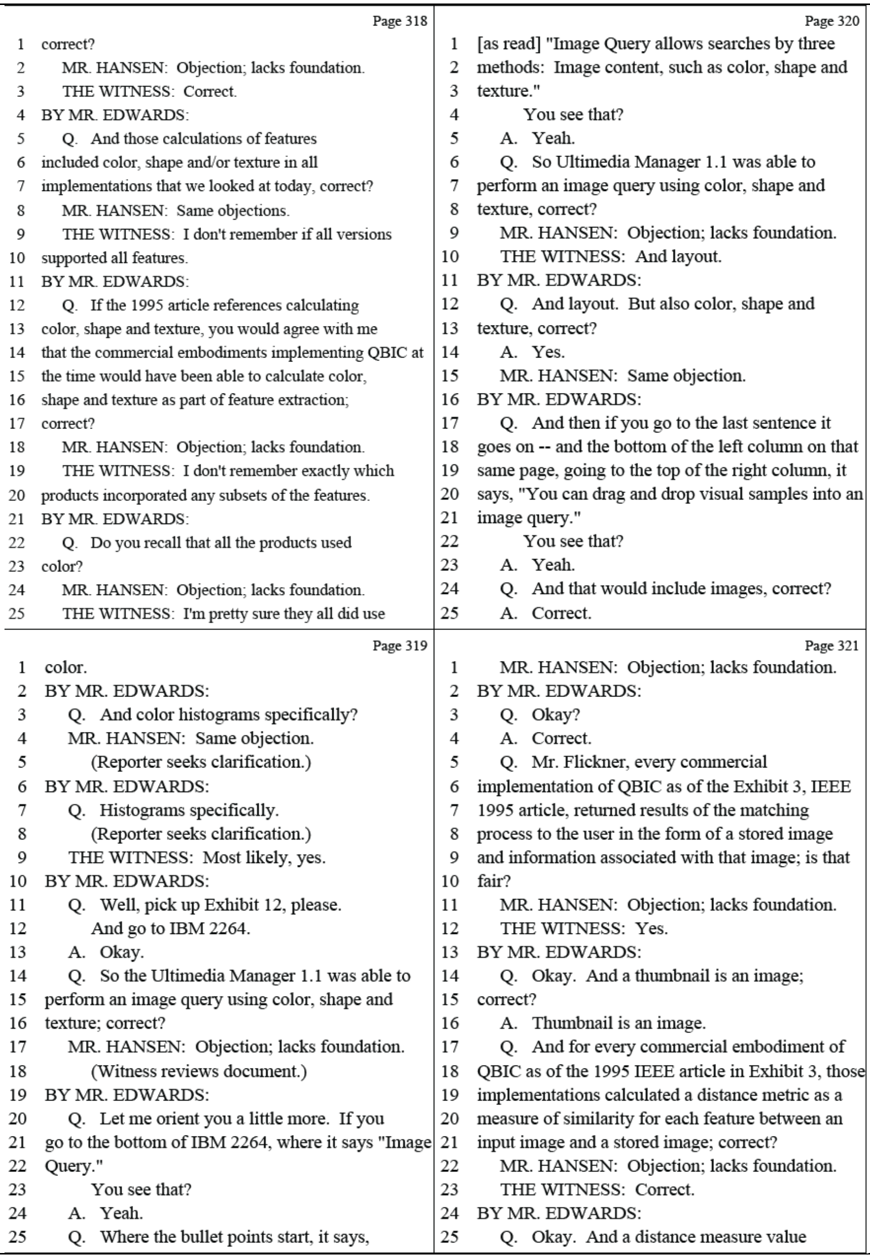

Page 318

correct?
MR. HANSEN: Objection; lacks foundation.
THE WITNESS: Correct.
BY MR. EDWARDS:
Q. And those calculations of features included color, shape and/or texture in all implementations that we looked at today, correct?
MR. HANSEN: Same objections.
THE WITNESS: I don't remember if all versions supported all features.
BY MR. EDWARDS:
Q. If the 1995 article references calculating color, shape and texture, you would agree with me that the commercial embodiments implementing QBIC at the time would have been able to calculate color, shape and texture as part of feature extraction; correct?
MR. HANSEN: Objection; lacks foundation.
THE WITNESS: I don't remember exactly which products incorporated any subsets of the features.
BY MR. EDWARDS:
Q. Do you recall that all the products used color?
MR. HANSEN: Objection; lacks foundation.
THE WITNESS: I'm pretty sure they all did use

Page 319

color.
BY MR. EDWARDS:
Q. And color histograms specifically?
MR. HANSEN: Same objection.
(Reporter seeks clarification.)
BY MR. EDWARDS:
Q. Histograms specifically.
(Reporter seeks clarification.)
THE WITNESS: Most likely, yes.
BY MR. EDWARDS:
Q. Well, pick up Exhibit 12, please.
And go to IBM 2264.
A. Okay.
Q. So the Ultimedia Manager 1.1 was able to perform an image query using color, shape and texture; correct?
MR. HANSEN: Objection; lacks foundation.
(Witness reviews document.)
BY MR. EDWARDS:
Q. Let me orient you a little more. If you go to the bottom of IBM 2264, where it says "Image Query."
You see that?
A. Yeah.
Q. Where the bullet points start, it says,

Page 320

[as read] "Image Query allows searches by three methods: Image content, such as color, shape and texture."
You see that?
A. Yeah.
Q. So Ultimedia Manager 1.1 was able to perform an image query using color, shape and texture, correct?
MR. HANSEN: Objection; lacks foundation.
THE WITNESS: And layout.
BY MR. EDWARDS:
Q. And layout. But also color, shape and texture, correct?
A. Yes.
MR. HANSEN: Same objection.
BY MR. EDWARDS:
Q. And then if you go to the last sentence it goes on -- and the bottom of the left column on that same page, going to the top of the right column, it says, "You can drag and drop visual samples into an image query."
You see that?
A. Yeah.
Q. And that would include images, correct?
A. Correct.

Page 321

MR. HANSEN: Objection; lacks foundation.
BY MR. EDWARDS:
Q. Okay?
A. Correct.
Q. Mr. Flickner, every commercial implementation of QBIC as of the Exhibit 3, IEEE 1995 article, returned results of the matching process to the user in the form of a stored image and information associated with that image; is that fair?
MR. HANSEN: Objection; lacks foundation.
THE WITNESS: Yes.
BY MR. EDWARDS:
Q. Okay. And a thumbnail is an image; correct?
A. Thumbnail is an image.
Q. And for every commercial embodiment of QBIC as of the 1995 IEEE article in Exhibit 3, those implementations calculated a distance metric as a measure of similarity for each feature between an input image and a stored image; correct?
MR. HANSEN: Objection; lacks foundation.
THE WITNESS: Correct.
BY MR. EDWARDS:
Q. Okay. And a distance measure value

Element	'529 Claim Recitation	Exemplary Citations in Reference
		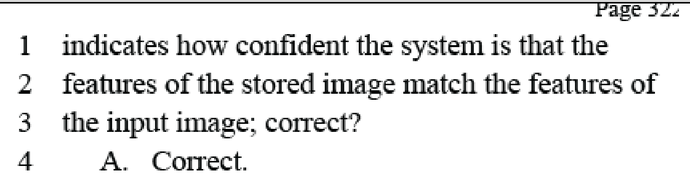 Page 322 indicates how confident the system is that the features of the stored image match the features of the input image; correct? A. Correct. IBM000001 at 39 – 41

Element	'529 Claim Recitation	Exemplary Citations in Reference
		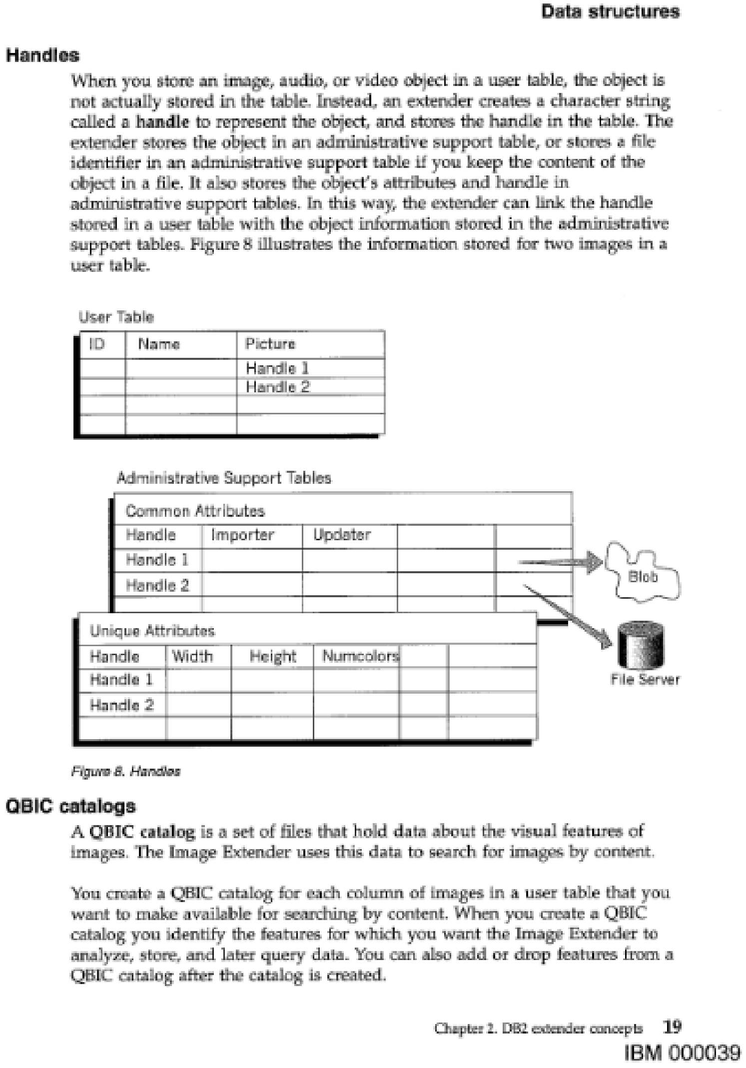 Data structures **Handles** When you store an image, audio, or video object in a user table, the object is not actually stored in the table. Instead, an extender creates a character string called a **handle** to represent the object, and stores the handle in the table. The extender stores the object in an administrative support table, or stores a file identifier in an administrative support table if you keep the content of the object in a file. It also stores the object's attributes and handle in administrative support tables. In this way, the extender can link the handle stored in a user table with the object information stored in the administrative support tables. Figure 8 illustrates the information stored for two images in a user table. *Figure 8. Handles* **QBIC catalogs** A **QBIC catalog** is a set of files that hold data about the visual features of images. The Image Extender uses this data to search for images by content. You create a QBIC catalog for each column of images in a user table that you want to make available for searching by content. When you create a QBIC catalog you identify the features for which you want the Image Extender to analyze, store, and later query data. You can also add or drop features from a QBIC catalog after the catalog is created. Chapter 2. DB2 extender concepts 19 IBM 000039

Element	'529 Claim Recitation	Exemplary Citations in Reference
		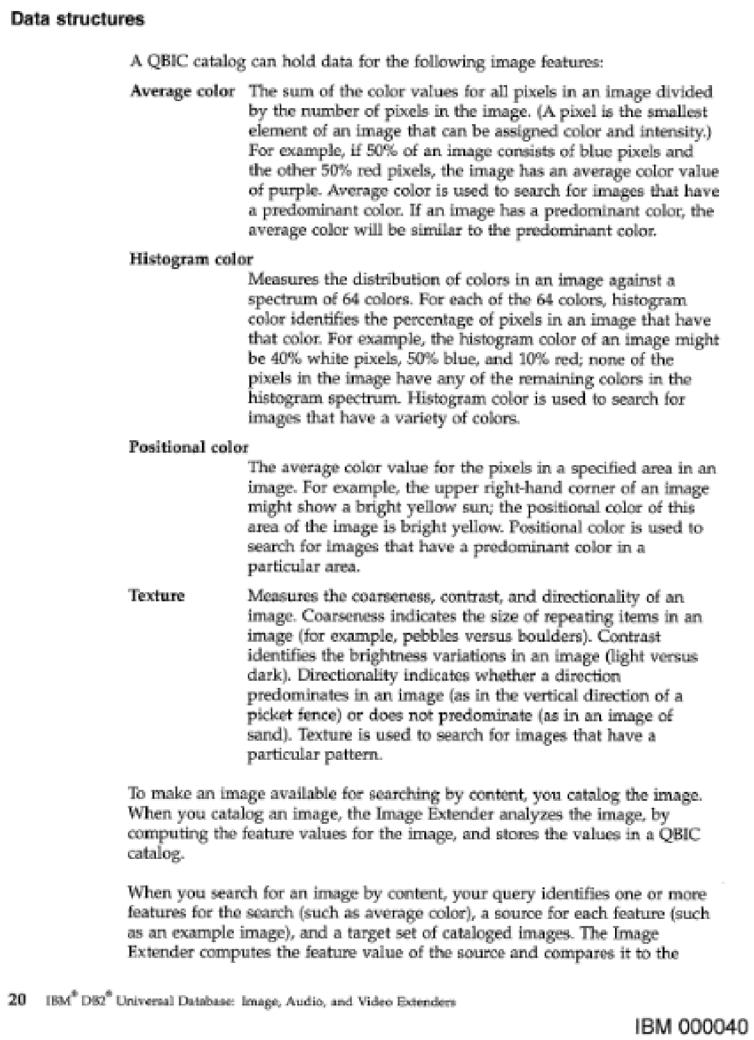 **Data structures** A QBIC catalog can hold data for the following image features: **Average color** The sum of the color values for all pixels in an image divided by the number of pixels in the image. (A pixel is the smallest element of an image that can be assigned color and intensity.) For example, if 50% of an image consists of blue pixels and the other 50% red pixels, the image has an average color value of purple. Average color is used to search for images that have a predominant color. If an image has a predominant color, the average color will be similar to the predominant color. **Histogram color** Measures the distribution of colors in an image against a spectrum of 64 colors. For each of the 64 colors, histogram color identifies the percentage of pixels in an image that have that color. For example, the histogram color of an image might be 40% white pixels, 50% blue, and 10% red; none of the pixels in the image have any of the remaining colors in the histogram spectrum. Histogram color is used to search for images that have a variety of colors. **Positional color** The average color value for the pixels in a specified area in an image. For example, the upper right-hand corner of an image might show a bright yellow sun; the positional color of this area of the image is bright yellow. Positional color is used to search for images that have a predominant color in a particular area. **Texture** Measures the coarseness, contrast, and directionality of an image. Coarseness indicates the size of repeating items in an image (for example, pebbles versus boulders). Contrast identifies the brightness variations in an image (light versus dark). Directionality indicates whether a direction predominates in an image (as in the vertical direction of a picket fence) or does not predominate (as in an image of sand). Texture is used to search for images that have a particular pattern. To make an image available for searching by content, you catalog the image. When you catalog an image, the Image Extender analyzes the image, by computing the feature values for the image, and stores the values in a QBIC catalog. When you search for an image by content, your query identifies one or more features for the search (such as average color), a source for each feature (such as an example image), and a target set of cataloged images. The Image Extender computes the feature value of the source and compares it to the 20 IBM® DB2® Universal Database: Image, Audio, and Video Extenders IBM 000040

Element	'529 Claim Recitation	Exemplary Citations in Reference
		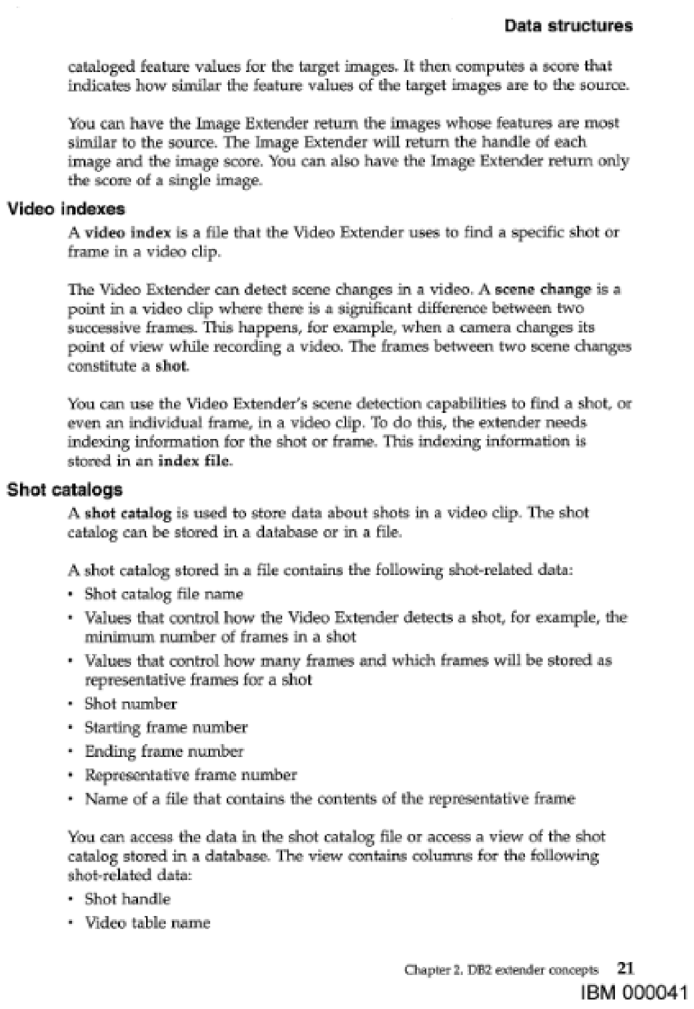 Data structures cataloged feature values for the target images. It then computes a score that indicates how similar the feature values of the target images are to the source. You can have the Image Extender return the images whose features are most similar to the source. The Image Extender will return the handle of each image and the image score. You can also have the Image Extender return only the score of a single image. **Video indexes** A **video index** is a file that the Video Extender uses to find a specific shot or frame in a video clip. The Video Extender can detect scene changes in a video. A **scene change** is a point in a video clip where there is a significant difference between two successive frames. This happens, for example, when a camera changes its point of view while recording a video. The frames between two scene changes constitute a **shot**. You can use the Video Extender's scene detection capabilities to find a shot, or even an individual frame, in a video clip. To do this, the extender needs indexing information for the shot or frame. This indexing information is stored in an **index file**. **Shot catalogs** A **shot catalog** is used to store data about shots in a video clip. The shot catalog can be stored in a database or in a file. A shot catalog stored in a file contains the following shot-related data: • Shot catalog file name • Values that control how the Video Extender detects a shot, for example, the minimum number of frames in a shot • Values that control how many frames and which frames will be stored as representative frames for a shot • Shot number • Starting frame number • Ending frame number • Representative frame number • Name of a file that contains the contents of the representative frame You can access the data in the shot catalog file or access a view of the shot catalog stored in a database. The view contains columns for the following shot-related data: • Shot handle • Video table name Chapter 2. DB2 extender concepts 21 IBM 000041 IBM0002263 at 2264

Element	'529 Claim Recitation	Exemplary Citations in Reference
		

Description

Ultimedia Manager 1.1 opens new dimensions in visual queries, helping you manage your image database more efficiently. With its easy-to-use interface, even entry-level users can perform sophisticated searches. You can search by image content, data description, or a combination of both. Once the images are found, you can view the images in full resolution, use the images in another application, or read the associated textual comments.

You can use Ultimedia Manager 1.1 either as a stand-alone product or in a LAN environment. As a stand-alone product, Ultimedia Manager 1.1 employs full image classification and image query capabilities. Within the LAN environment, the Ultimedia Manager image catalogs can be sorted on a DB2 for OS/2 or DB2 for AIX database server. The image catalogs can be accessed by Windows or OS/2 clients using the Client Search Feature or the Visualizer Ultimedia Query.

Ultimedia Manager runs as an OS/2 application to a DB2 for OS/2 or DB2 for AIX database server in a LAN, or as a stand-alone application with a local DB2 for OS/2 database. The Client Search Feature runs on DOS or Windows, and accesses files in dBASE flat file format. A database manager is not required to access catalogs in dBASE file format.

Image Classification

With the Image Classifier, you can create visual indexes based on color, shape, category, and layout elements, which are areas of the image containing the patterns you want to use for searches. You can store these elements in catalogs in DB2 for OS/2 databases. The catalogs can be used locally or, through the use of utility programs, be transferred to DB2 for OS/2 or AIX servers. In addition, you can convert the catalog data into dBASE flat file format for access by Windows client applications.

The Image Classifier and Analyzer allows automatic or interactive classification of images for later image queries, flexibility in storage of catalog information, and access to multiple users. You can organize large image catalogs so that later searches for images can be performed easily. You can classify automatically using a batch program, or interactively by specifying areas of interest in an image and indicating the visual cue (color, shape, or texture) to be saved for later queries.

Image Query

Once the image catalogs are created, you can use the Image Query component to find image and business information. Business information, such as character and numeric data, are stored in the database tables using normal table population methods. Image Query is provided in Ultimedia Manager and Client Search Feature.

Image Query allows searches by three methods:

- Image content, such as color, shape, texture, or layout
- Data description, such as business data stored in the database (character or numeric) or descriptive text (category)
- Combinations of image content and data description.

Image Content: To search by image content, specify the attributes (color, shape, texture, or layout). You can drag and drop visual samples into an image query, select a color from the color wheel, or sketch a shape.

Data Description: To search by data description, enter character and numeric data values in the database tables using the simple fill-in-the-blanks query panel. You can enter data such as name, number, date, or a range of dates. You can integrate character or numeric data as part of the image query, enabling a more effective search for better results. The simple fill-in-the-blanks method for specifying desired conditions is so easy even non-programmers can use database table information as part of their searches.

Combinations: Image queries using either image content or data description, or combinations of both, allow you to search through thousands of images more quickly and effectively than the traditional text and keyword search methods.

Image Formats

Ultimedia Manager 1.1 understands these image formats:

- Audio Visual Connection (AVC) images
- OS/2 bitmap
- Windows 3.x bitmap
- PCX
- Intel and Motorola TIFF
- TARGA TGA
- TIFF (FAX)
- GIF
- Kodak Photo-CD

Product Positioning

The advanced image search technology in Ultimedia Manager, query by image content, was developed in conjunction with the IBM Almaden Research Center. It provides function not available with conventional, text-only image management systems.

Ultimedia Manager fulfills image search and retrieval needs of business users and creative professionals through use of online image libraries or collections. It is intended for users who have large image libraries and need fast, flexible access to their image data, users who may also want to add character or numeric data to their image queries, and users who feel limited by text-based systems.

Ultimedia Manager 1.1 works with the DB2 Family of products and the Visualizer Ultimedia Query product. Catalogs created by Ultimedia Manager can be stored in DB2 for OS/2 database tables. Image catalogs also can be stored in DB2 for OS/2 or DB2 for AIX database servers in client/server environments.

OS/2 and Windows clients in a LAN environment using the Client Search Feature and the Visualizer Ultimedia Query can access catalogs created with Ultimedia Manager 1.1. Visualizer Ultimedia Query provides users with the capability to:

- Add multimedia columns (such as image, video, audio, and formatted text) to DB2 for OS/2 database tables
- Query and display images
- View or listen to video and audio information
- Display formatted text

Users of Ultimedia Manager 1.1 and Client Search Feature are focused on image management and query applications. Visualizer Ultimedia Query users are

298-042 -2-

IBM 0002264

Element	'529 Claim Recitation	Exemplary Citations in Reference
		IBM000777 at 794 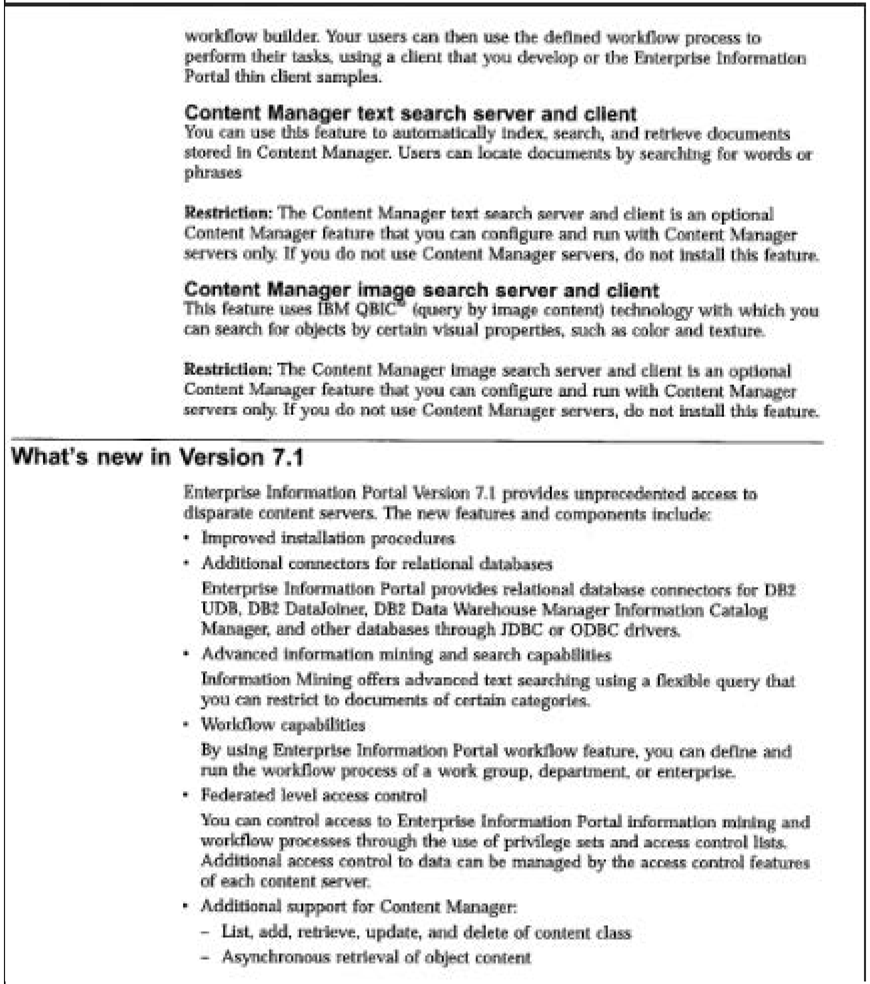 workflow builder. Your users can then use the defined workflow process to perform their tasks, using a client that you develop or the Enterprise Information Portal thin client samples. **Content Manager text search server and client** You can use this feature to automatically index, search, and retrieve documents stored in Content Manager. Users can locate documents by searching for words or phrases **Restriction:** The Content Manager text search server and client is an optional Content Manager feature that you can configure and run with Content Manager servers only. If you do not use Content Manager servers, do not install this feature. **Content Manager image search server and client** This feature uses IBM QBIC™ (query by image content) technology with which you can search for objects by certain visual properties, such as color and texture. **Restriction:** The Content Manager image search server and client is an optional Content Manager feature that you can configure and run with Content Manager servers only. If you do not use Content Manager servers, do not install this feature. **What's new in Version 7.1** Enterprise Information Portal Version 7.1 provides unprecedented access to disparate content servers. The new features and components include: • Improved installation procedures • Additional connectors for relational databases Enterprise Information Portal provides relational database connectors for DB2 UDB, DB2 DataJoiner, DB2 Data Warehouse Manager Information Catalog Manager, and other databases through JDBC or ODBC drivers. • Advanced information mining and search capabilities Information Mining offers advanced text searching using a flexible query that you can restrict to documents of certain categories. • Workflow capabilities By using Enterprise Information Portal workflow feature, you can define and run the workflow process of a work group, department, or enterprise. • Federated level access control You can control access to Enterprise Information Portal information mining and workflow processes through the use of privilege sets and access control lists. Additional access control to data can be managed by the access control features of each content server. • Additional support for Content Manager: – List, add, retrieve, update, and delete of content class – Asynchronous retrieval of object content To the extent it is found that QBIC does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that QBIC does not anticipate this claim, QBIC renders it obvious, either alone, in combination with the knowledge of a

Element	'529 Claim Recitation	Exemplary Citations in Reference
		person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine QBIC with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1E]	associate the object with information; and	QBIC discloses associate the object with information. The QBIC Project: Querying Images by Content Using Color, Texture, and Shape (IBM 0002390–404) at 2396-97 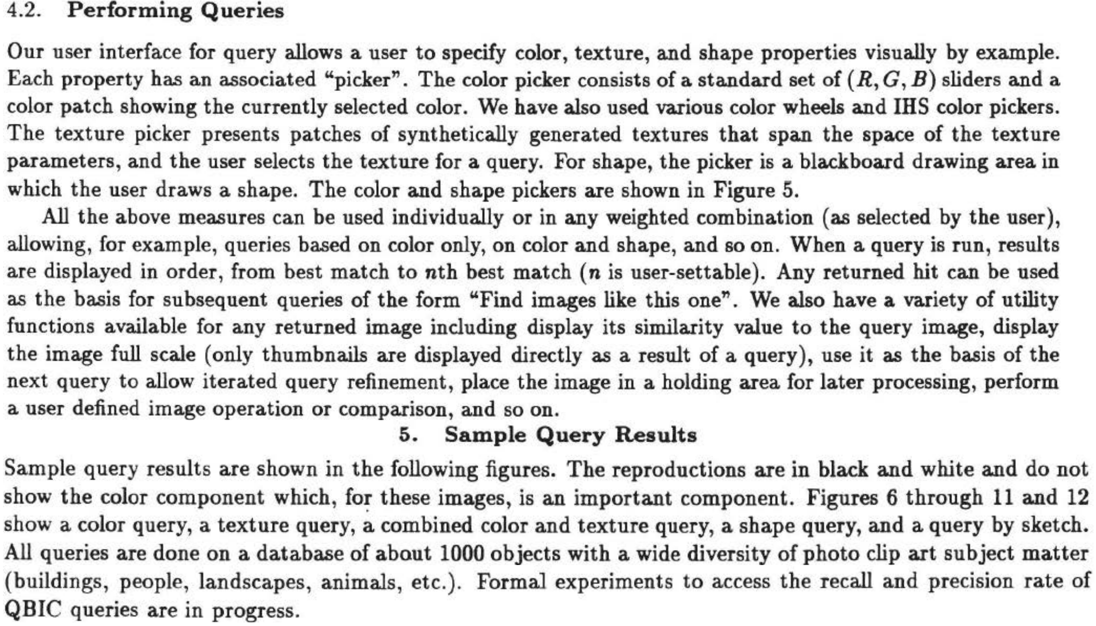 4.2. **Performing Queries** Our user interface for query allows a user to specify color, texture, and shape properties visually by example. Each property has an associated "picker". The color picker consists of a standard set of (R, G, B) sliders and a color patch showing the currently selected color. We have also used various color wheels and IHS color pickers. The texture picker presents patches of synthetically generated textures that span the space of the texture parameters, and the user selects the texture for a query. For shape, the picker is a blackboard drawing area in which the user draws a shape. The color and shape pickers are shown in Figure 5. All the above measures can be used individually or in any weighted combination (as selected by the user), allowing, for example, queries based on color only, on color and shape, and so on. When a query is run, results are displayed in order, from best match to nth best match (n is user-settable). Any returned hit can be used as the basis for subsequent queries of the form "Find images like this one". We also have a variety of utility functions available for any returned image including display its similarity value to the query image, display the image full scale (only thumbnails are displayed directly as a result of a query), use it as the basis of the next query to allow iterated query refinement, place the image in a holding area for later processing, perform a user defined image operation or comparison, and so on. 5. **Sample Query Results** Sample query results are shown in the following figures. The reproductions are in black and white and do not show the color component which, for these images, is an important component. Figures 6 through 11 and 12 show a color query, a texture query, a combined color and texture query, a shape query, and a query by sketch. All queries are done on a database of about 1000 objects with a wide diversity of photo clip art subject matter (buildings, people, landscapes, animals, etc.). Formal experiments to access the recall and precision rate of QBIC queries are in progress. 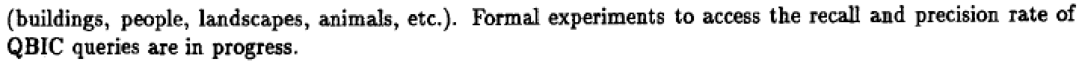 (buildings, people, landscapes, animals, etc.). Formal experiments to access the recall and precision rate of QBIC queries are in progress.

Element	'529 Claim Recitation	Exemplary Citations in Reference
		IBM0002390 at 2391-92 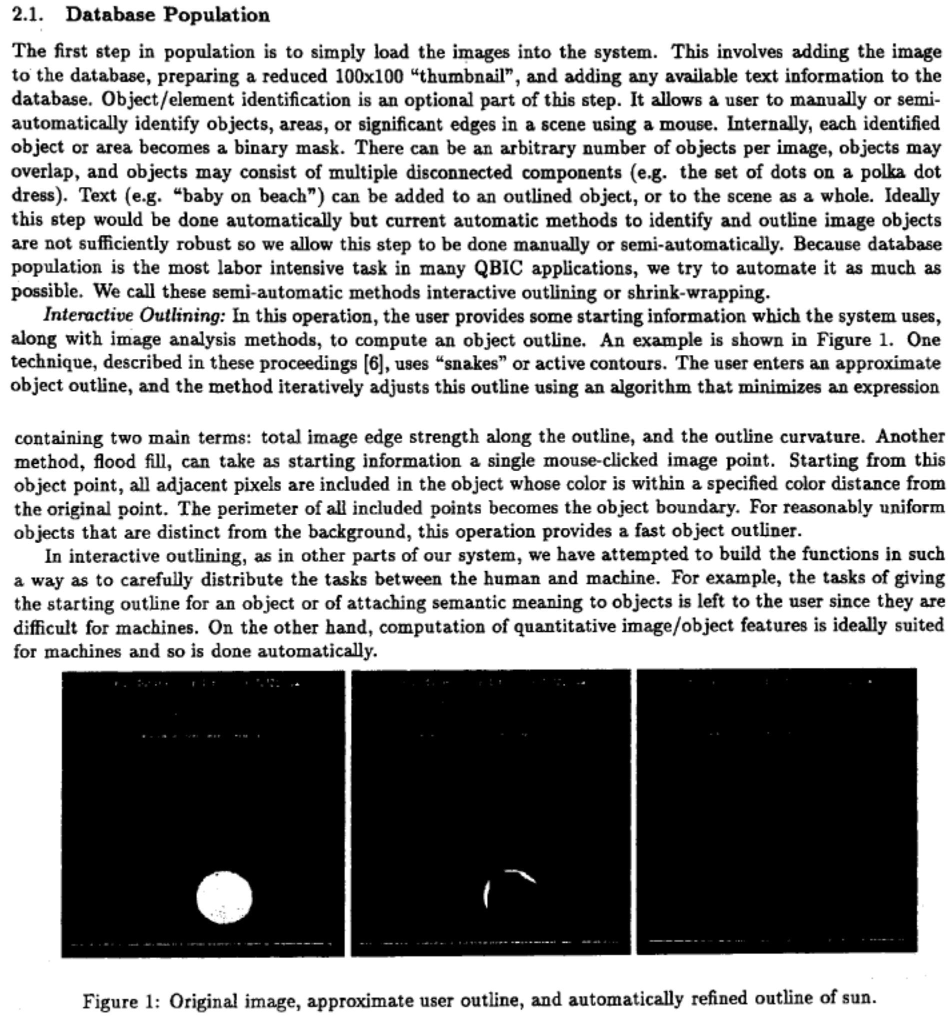 2.1. **Database Population** The first step in population is to simply load the images into the system. This involves adding the image to the database, preparing a reduced 100x100 "thumbnail", and adding any available text information to the database. Object/element identification is an optional part of this step. It allows a user to manually or semi-automatically identify objects, areas, or significant edges in a scene using a mouse. Internally, each identified object or area becomes a binary mask. There can be an arbitrary number of objects per image, objects may overlap, and objects may consist of multiple disconnected components (e.g. the set of dots on a polka dot dress). Text (e.g. "baby on beach") can be added to an outlined object, or to the scene as a whole. Ideally this step would be done automatically but current automatic methods to identify and outline image objects are not sufficiently robust so we allow this step to be done manually or semi-automatically. Because database population is the most labor intensive task in many QBIC applications, we try to automate it as much as possible. We call these semi-automatic methods interactive outlining or shrink-wrapping. *Interactive Outlining:* In this operation, the user provides some starting information which the system uses, along with image analysis methods, to compute an object outline. An example is shown in Figure 1. One technique, described in these proceedings [6], uses "snakes" or active contours. The user enters an approximate object outline, and the method iteratively adjusts this outline using an algorithm that minimizes an expression containing two main terms: total image edge strength along the outline, and the outline curvature. Another method, flood fill, can take as starting information a single mouse-clicked image point. Starting from this object point, all adjacent pixels are included in the object whose color is within a specified color distance from the original point. The perimeter of all included points becomes the object boundary. For reasonably uniform objects that are distinct from the background, this operation provides a fast object outliner. In interactive outlining, as in other parts of our system, we have attempted to build the functions in such a way as to carefully distribute the tasks between the human and machine. For example, the tasks of giving the starting outline for an object or of attaching semantic meaning to objects is left to the user since they are difficult for machines. On the other hand, computation of quantitative image/object features is ideally suited for machines and so is done automatically. Figure 1: Original image, approximate user outline, and automatically refined outline of sun.

Element	'529 Claim Recitation	Exemplary Citations in Reference
		The QBIC Project: Querying Images by Content Using Color, Texture, and Shape (IBM 0002390–404) at Figs. 7–9 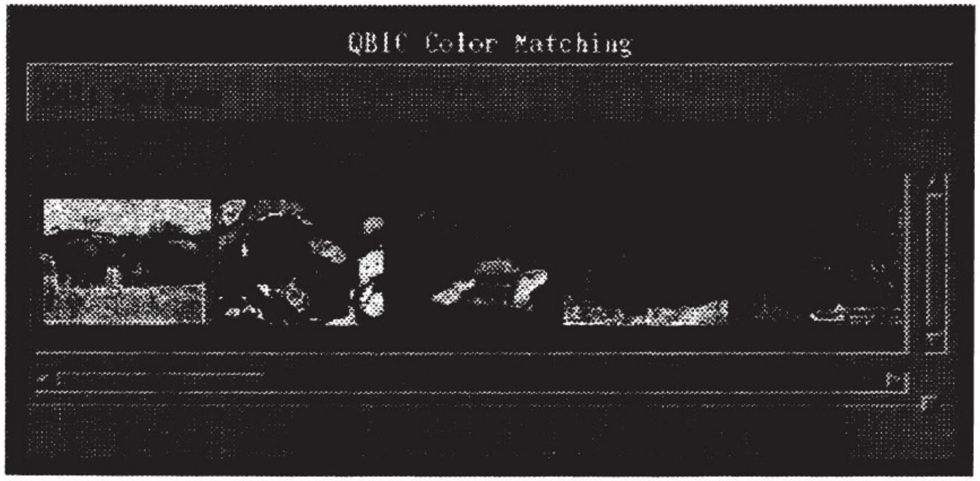 Figure 7: Result of query by example using color only.

Element	'529 Claim Recitation	Exemplary Citations in Reference
		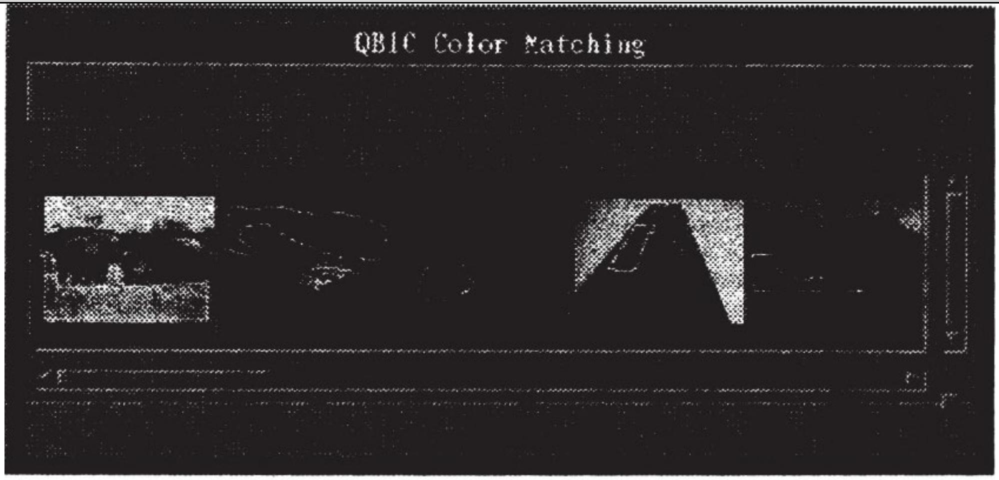 Figure 8: Result of query by example using texture only.

Element	'529 Claim Recitation	Exemplary Citations in Reference
		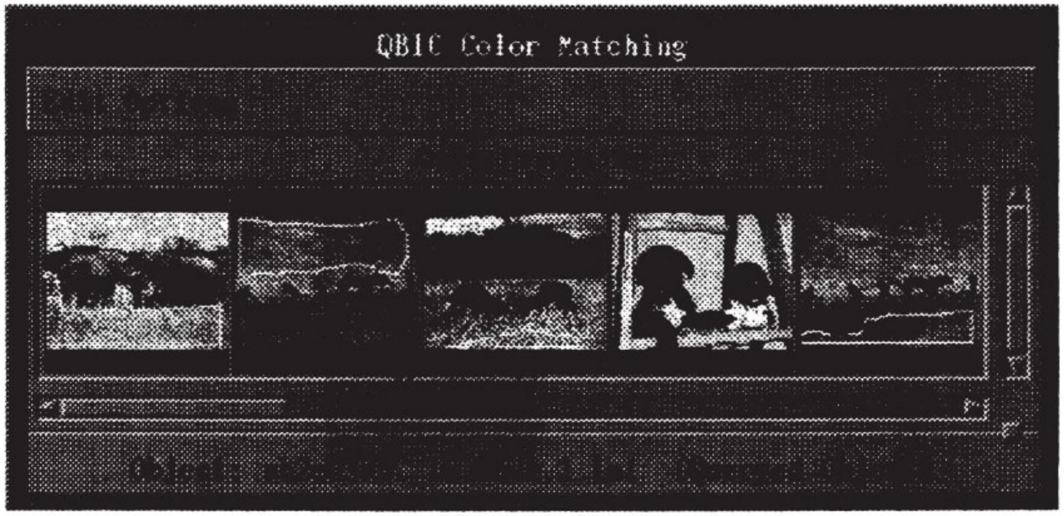 Figure 9: Result of query by example using color and texture. QBIC (Query by Image Content) (BOFA00001391–392) at 2 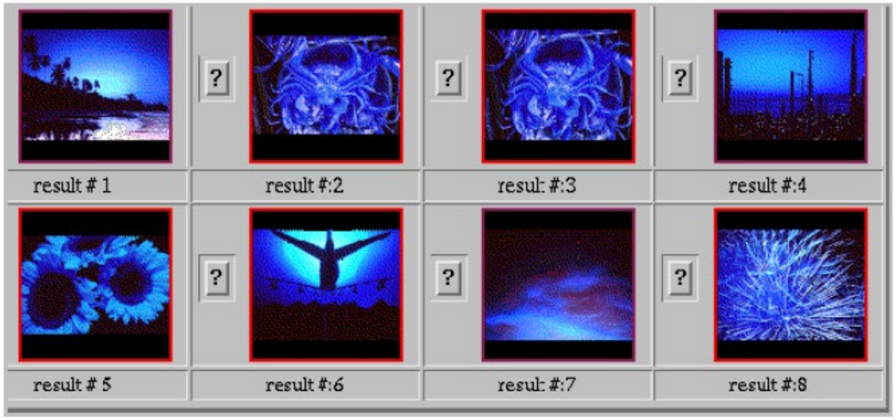 **Result of QBIC Search**

Element	'529 Claim Recitation	Exemplary Citations in Reference
		Query by Image and Video Content: The QBIC System (IBM 0000893–902): "To achieve this functionality, QBIC has two main components: database population (the process of creating an image database) and database query. During the population, images and videos are processed to extract features describing their content-colors, textures, shapes, and camera and object motion-and the features are stored in a database. During the query, the user composes a query graphically. Features are generated from the graphical query and then input to a matching engine that finds images or videos from the database with similar features. Figure 2 shows the system architecture." *Id.* at p. 3. "For both population and query, the QBIC data model has • still images or scenes (full images) that contain objects (subsets of an image), and • video shots that consist of sets of contiguous frames and contain motion objects. For still images, the QBIC data model distinguishes between 'scenes' (or images) and 'objects.' A scene is an image or single representative frame of video. An object is a part of a scene-for example, the fox in Figure 3-or a moving entity in a video. For still image database population, features are extracted from images and objects and stored in a database as shown in the top left part of Figure 2. Videos are broken into clips called shots. Representative frames, or r-frames, are generated for each extracted shot. R-frames are treated as still images, and features are extracted and stored in the database. Further processing of shots generates motion objects-for example, a car moving across the screen. Queries are allowed on objects ('Find images with a red, round object'), scenes ('Find images that have approximately 30-percent red and IS-percent blue colors'), shots ('Find all shots panning from left to right'), or any combination ('Find images that have 30 percent red and contain a blue textured object')." *Id.* at p. 3. "INTEGRATION WITH TEXT AND PARAMETRIC ANNOTATION. Query by visual content complements and extends existing query methods. Systems must be able to integrate queries combining date, subject matter, price, and availability with content properties such as color, texture, and shape." *Id.* at p. 8.

Element	'529 Claim Recitation	Exemplary Citations in Reference
		Image, Audio, and Video Extenders Administration and Programming (IBM 000001–617): "Your applications can even search for images by content. Imagine an application that uses visual examples to search for images. With such an application, users could select an example image and have the application find other images that have colors or textures similar to those in the example. With DB2 extenders' Query by Image Content (QBIC) capability, you can create applications that search for images in this visual way." *Id*. at p. 4. "To make an image available for searching by content, you catalog the image. When you catalog an image, the Image Extender analyzes the image, by computing the feature values for the image, and stores the values in a QBIC catalog. When you search for an image by content, your query identifies one or more features for the search (such as average color), a source for each feature (such as an example image), and a target set of cataloged images. The Image Extender computes the feature value of the source and compares it to the cataloged feature values for the target images. It then computes a score that indicates how similar the feature values of the target images are to the source. You can have the Image Extender return the images whose features are most similar to the source. The Image Extender will return the handle of each image and the image score. You can also have the Image Extender return only the score of a single image." *Id*. at pp. 20–21.

Element	'529 Claim Recitation	Exemplary Citations in Reference
		QBIC Demo (IBM000747): 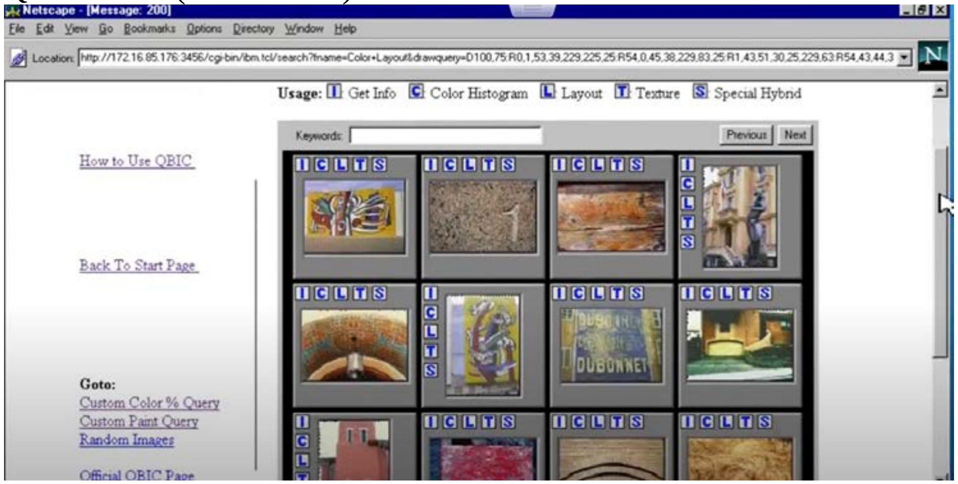 Netscape - [Message: 200] Location: http://172.16.85.176:3456/cgi-bin/ibm.tcl/search?fname=Color+Layout&drawquery=D100,75:R0,1,53,39,229,225,25:R54,0,45,38,229,83,25:R1,43,51,30,25,229,63:R54,43,44,3 Keywords: QBIC Demo (IBM000747):  Location: http://172.16.85.176:3456/cgi-bin/ibm.tcl/info?france50%2Ejpg Image france50.jpg

Element	'529 Claim Recitation	Exemplary Citations in Reference
		Flickner Depo at 201-202: 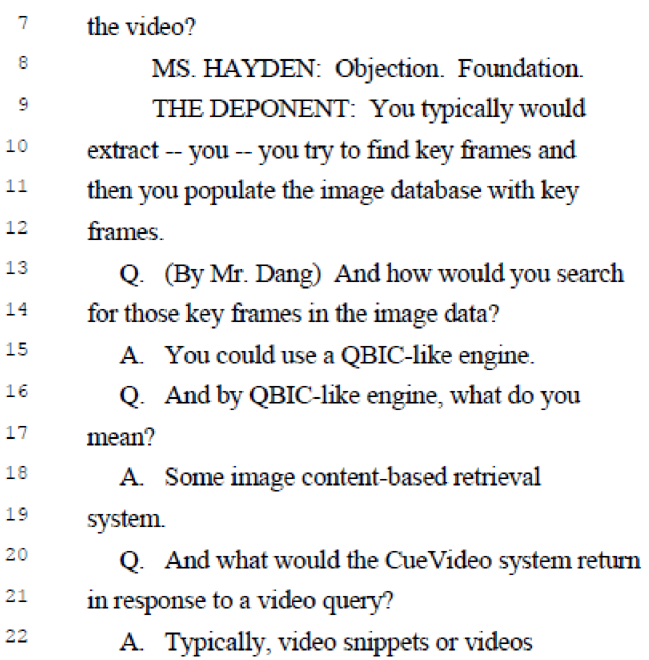 Q. Now, to search a video by content using CueVideo, how would the features be extracted from the video? MS. HAYDEN: Objection. Foundation. THE DEPONENT: You typically would extract -- you -- you try to find key frames and then you populate the image database with key frames. Q. (By Mr. Dang) And how would you search for those key frames in the image data? A. You could use a QBIC-like engine. Q. And by QBIC-like engine, what do you mean? A. Some image content-based retrieval system. Q. And what would the CueVideo system return in response to a video query? A. Typically, video snippets or videos

Element	'529 Claim Recitation	Exemplary Citations in Reference
		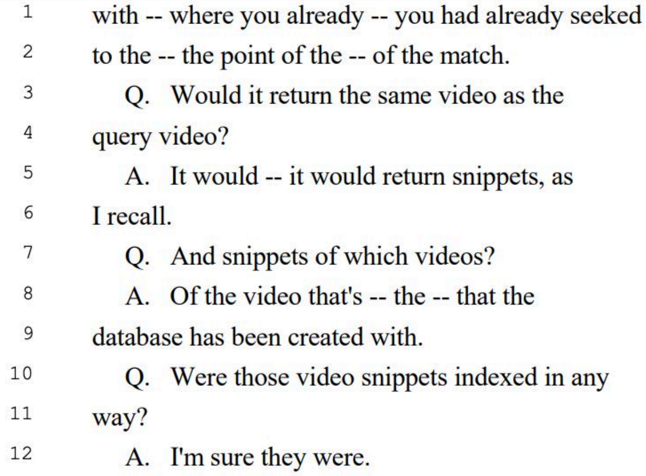 with -- where you already -- you had already seeked to the -- the point of the -- of the match. Q. Would it return the same video as the query video? A. It would -- it would return snippets, as I recall. Q. And snippets of which videos? A. Of the video that's -- the -- that the database has been created with. Q. Were those video snippets indexed in any way? A. I'm sure they were. Flickner Depo at 131-132: 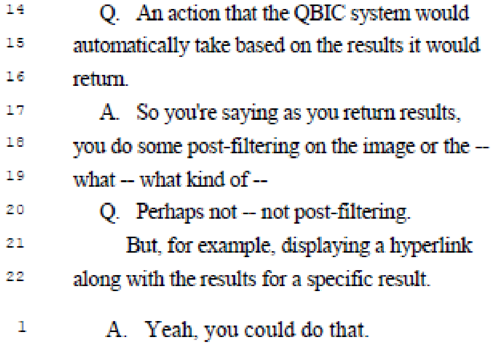 Q. An action that the QBIC system would automatically take based on the results it would return. A. So you're saying as you return results, you do some post-filtering on the image or the -- what -- what kind of -- Q. Perhaps not -- not post-filtering. But, for example, displaying a hyperlink along with the results for a specific result. A. Yeah, you could do that. *See also* Flickner Depo at 54-55, 75-76, 78-79, 88, 89-90, 159, 209-210. Flickner 8/29/23 Depo at 53–56:

Element	'529 Claim Recitation	Exemplary Citations in Reference
		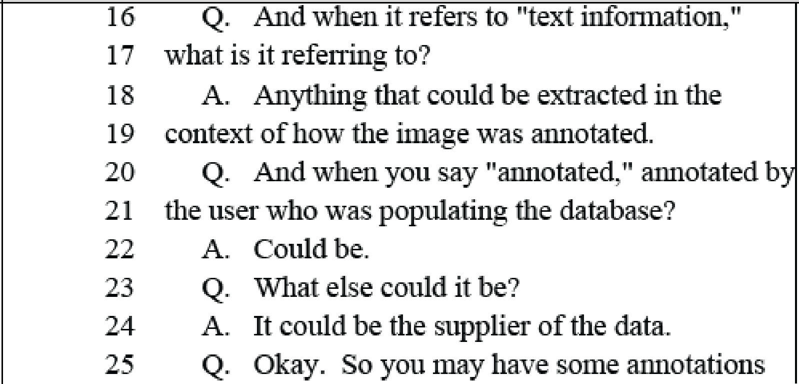 Q. And when it refers to "text information," what is it referring to? A. Anything that could be extracted in the context of how the image was annotated. Q. And when you say "annotated," annotated by the user who was populating the database? A. Could be. Q. What else could it be? A. It could be the supplier of the data. Q. Okay. So you may have some annotations

Element	'529 Claim Recitation	Exemplary Citations in Reference
		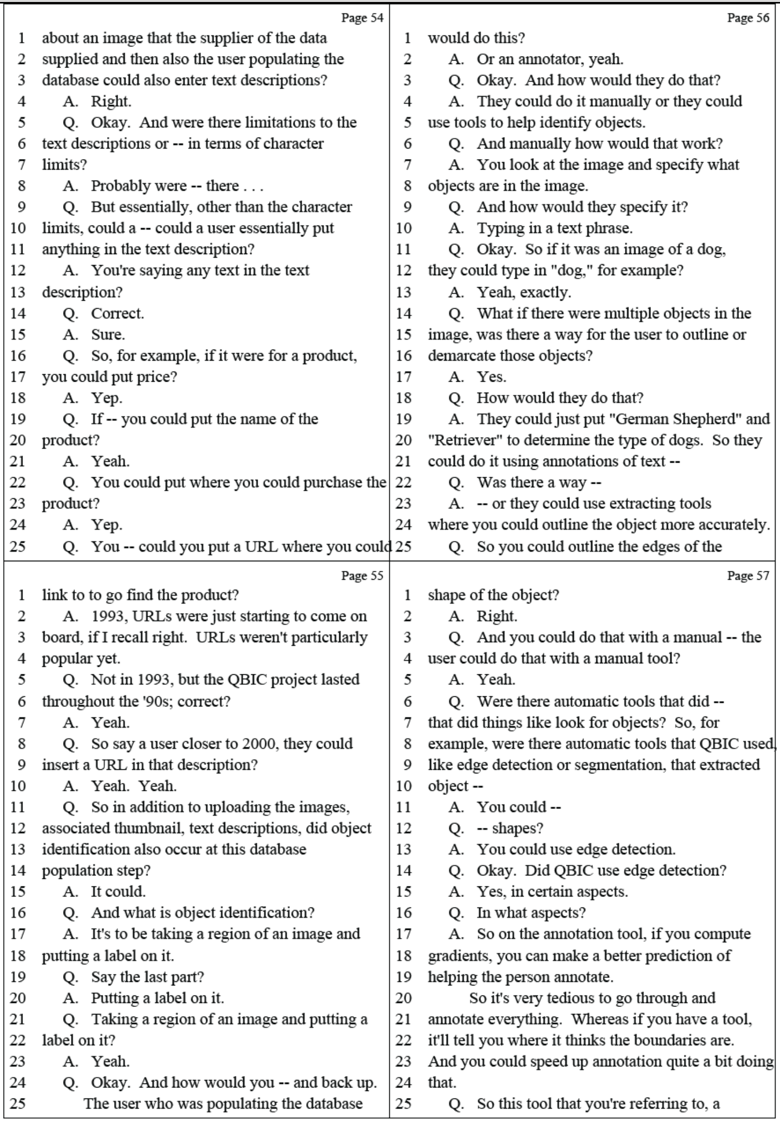

Page 54

about an image that the supplier of the data supplied and then also the user populating the database could also enter text descriptions?
A. Right.
Q. Okay. And were there limitations to the text descriptions or -- in terms of character limits?
A. Probably were -- there . . .
Q. But essentially, other than the character limits, could a -- could a user essentially put anything in the text description?
A. You're saying any text in the text description?
Q. Correct.
A. Sure.
Q. So, for example, if it were for a product, you could put price?
A. Yep.
Q. If -- you could put the name of the product?
A. Yeah.
Q. You could put where you could purchase the product?
A. Yep.
Q. You -- could you put a URL where you could

Page 55

link to to go find the product?
A. 1993, URLs were just starting to come on board, if I recall right. URLs weren't particularly popular yet.
Q. Not in 1993, but the QBIC project lasted throughout the '90s; correct?
A. Yeah.
Q. So say a user closer to 2000, they could insert a URL in that description?
A. Yeah. Yeah.
Q. So in addition to uploading the images, associated thumbnail, text descriptions, did object identification also occur at this database population step?
A. It could.
Q. And what is object identification?
A. It's to be taking a region of an image and putting a label on it.
Q. Say the last part?
A. Putting a label on it.
Q. Taking a region of an image and putting a label on it?
A. Yeah.
Q. Okay. And how would you -- and back up.
The user who was populating the database

Page 56

would do this?
A. Or an annotator, yeah.
Q. Okay. And how would they do that?
A. They could do it manually or they could use tools to help identify objects.
Q. And manually how would that work?
A. You look at the image and specify what objects are in the image.
Q. And how would they specify it?
A. Typing in a text phrase.
Q. Okay. So if it was an image of a dog, they could type in "dog," for example?
A. Yeah, exactly.
Q. What if there were multiple objects in the image, was there a way for the user to outline or demarcate those objects?
A. Yes.
Q. How would they do that?
A. They could just put "German Shepherd" and "Retriever" to determine the type of dogs. So they could do it using annotations of text --
Q. Was there a way --
A. -- or they could use extracting tools where you could outline the object more accurately.
Q. So you could outline the edges of the

Page 57

shape of the object?
A. Right.
Q. And you could do that with a manual -- the user could do that with a manual tool?
A. Yeah.
Q. Were there automatic tools that did -- that did things like look for objects? So, for example, were there automatic tools that QBIC used, like edge detection or segmentation, that extracted object --
A. You could --
Q. -- shapes?
A. You could use edge detection.
Q. Okay. Did QBIC use edge detection?
A. Yes, in certain aspects.
Q. In what aspects?
A. So on the annotation tool, if you compute gradients, you can make a better prediction of helping the person annotate.
So it's very tedious to go through and annotate everything. Whereas if you have a tool, it'll tell you where it thinks the boundaries are. And you could speed up annotation quite a bit doing that.
Q. So this tool that you're referring to, a

Element	'529 Claim Recitation	Exemplary Citations in Reference
		Id. at 60 – 62: 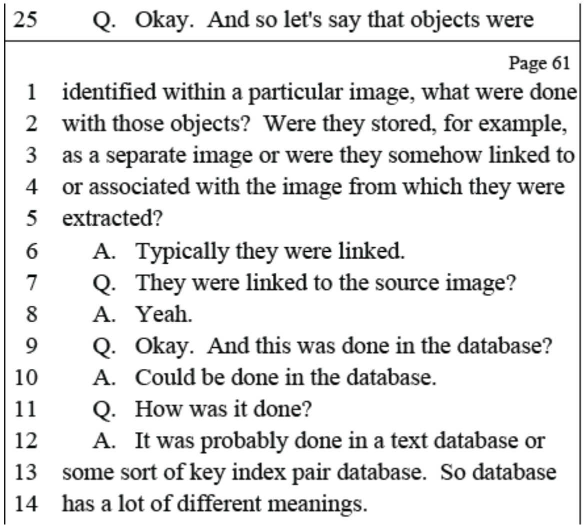 Q. Okay. And so let's say that objects were Page 61 identified within a particular image, what were done with those objects? Were they stored, for example, as a separate image or were they somehow linked to or associated with the image from which they were extracted? A. Typically they were linked. Q. They were linked to the source image? A. Yeah. Q. Okay. And this was done in the database? A. Could be done in the database. Q. How was it done? A. It was probably done in a text database or some sort of key index pair database. So database has a lot of different meanings. 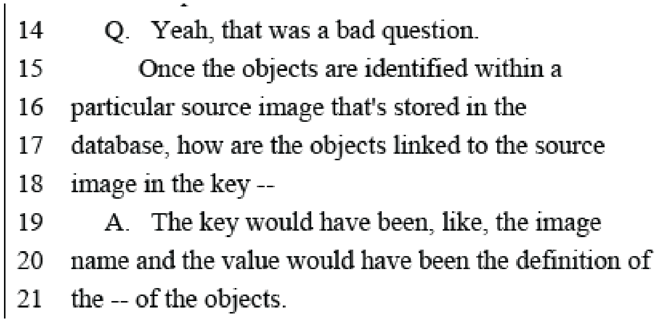 Q. Yeah, that was a bad question. Once the objects are identified within a particular source image that's stored in the database, how are the objects linked to the source image in the key -- A. The key would have been, like, the image name and the value would have been the definition of the -- of the objects. *Id.* at 103:

Element	'529 Claim Recitation	Exemplary Citations in Reference
		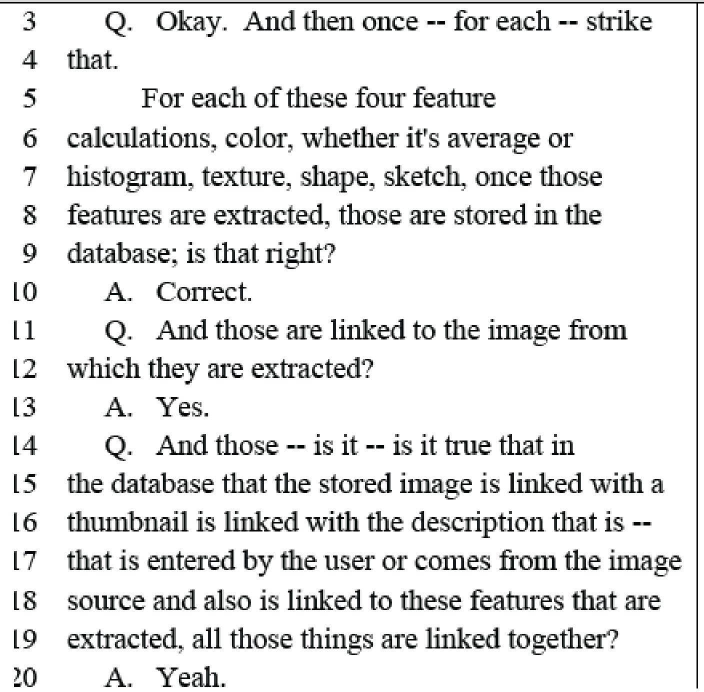 Q. Okay. And then once -- for each -- strike that. For each of these four feature calculations, color, whether it's average or histogram, texture, shape, sketch, once those features are extracted, those are stored in the database; is that right? A. Correct. Q. And those are linked to the image from which they are extracted? A. Yes. Q. And those -- is it -- is it true that in the database that the stored image is linked with a thumbnail is linked with the description that is -- that is entered by the user or comes from the image source and also is linked to these features that are extracted, all those things are linked together? A. Yeah.

Element	'529 Claim Recitation	Exemplary Citations in Reference
		Id. at 129 – 131: 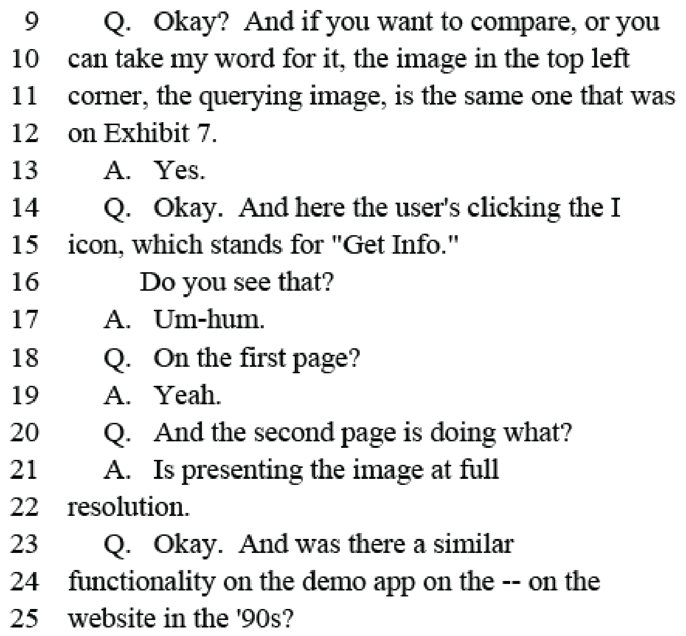 Q. Okay? And if you want to compare, or you can take my word for it, the image in the top left corner, the querying image, is the same one that was on Exhibit 7. A. Yes. Q. Okay. And here the user's clicking the I icon, which stands for "Get Info." Do you see that? A. Um-hum. Q. On the first page? A. Yeah. Q. And the second page is doing what? A. Is presenting the image at full resolution. Q. Okay. And was there a similar functionality on the demo app on the -- on the website in the '90s?

			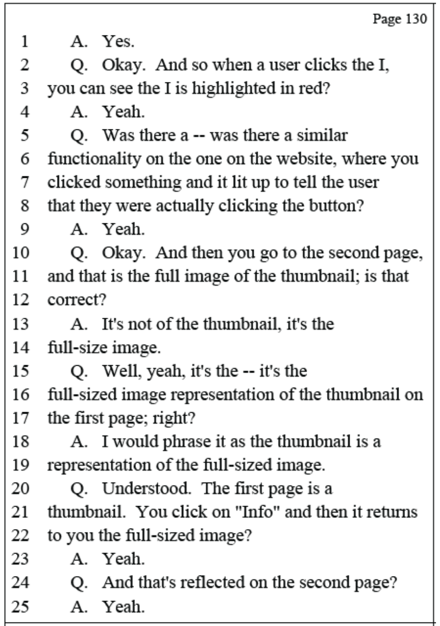 Page 130 A. Yes. Q. Okay. And so when a user clicks the I, you can see the I is highlighted in red? A. Yeah. Q. Was there a -- was there a similar functionality on the one on the website, where you clicked something and it lit up to tell the user that they were actually clicking the button? A. Yeah. Q. Okay. And then you go to the second page, and that is the full image of the thumbnail; is that correct? A. It's not of the thumbnail, it's the full-size image. Q. Well, yeah, it's the -- it's the full-sized image representation of the thumbnail on the first page; right? A. I would phrase it as the thumbnail is a representation of the full-sized image. Q. Understood. The first page is a thumbnail. You click on "Info" and then it returns to you the full-sized image? A. Yeah. Q. And that's reflected on the second page? A. Yeah. 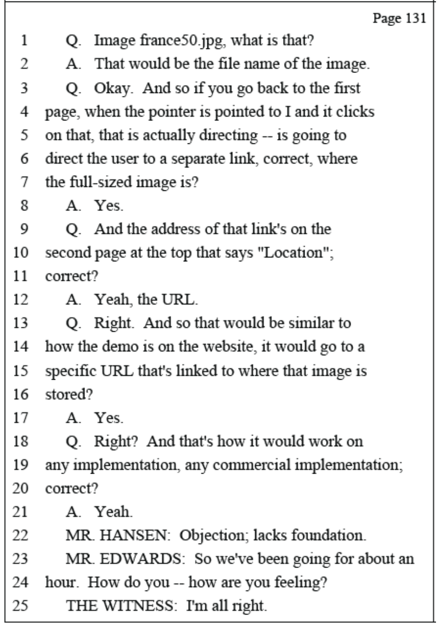 Page 131 Q. Image france50.jpg, what is that? A. That would be the file name of the image. Q. Okay. And so if you go back to the first page, when the pointer is pointed to I and it clicks on that, that is actually directing -- is going to direct the user to a separate link, correct, where the full-sized image is? A. Yes. Q. And the address of that link's on the second page at the top that says "Location"; correct? A. Yeah, the URL. Q. Right. And so that would be similar to how the demo is on the website, it would go to a specific URL that's linked to where that image is stored? A. Yes. Q. Right? And that's how it would work on any implementation, any commercial implementation; correct? A. Yeah. MR. HANSEN: Objection; lacks foundation. MR. EDWARDS: So we've been going for about an hour. How do you -- how are you feeling? THE WITNESS: I'm all right.

Element	'529 Claim Recitation	Exemplary Citations in Reference
		Id. at 161: 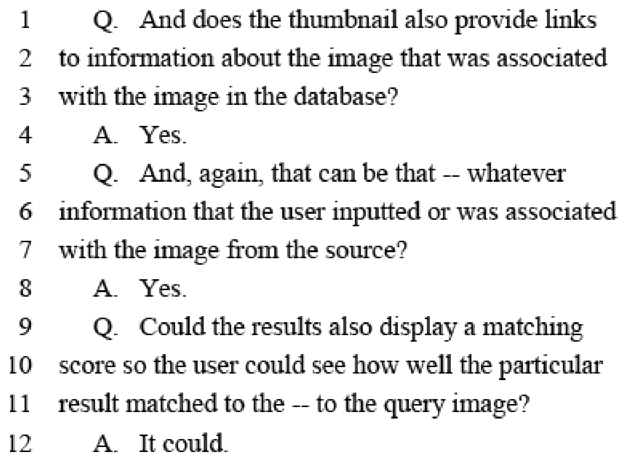 Q. And does the thumbnail also provide links to information about the image that was associated with the image in the database? A. Yes. Q. And, again, that can be that -- whatever information that the user inputted or was associated with the image from the source? A. Yes. Q. Could the results also display a matching score so the user could see how well the particular result matched to the -- to the query image? A. It could. *Id.* at 165:

Element	'529 Claim Recitation	Exemplary Citations in Reference
		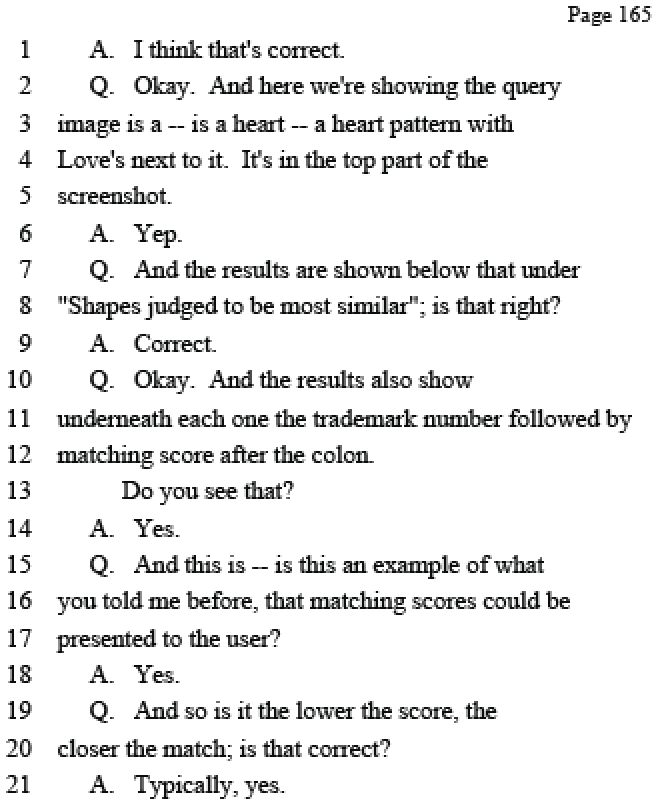 Page 165 A. I think that's correct. Q. Okay. And here we're showing the query image is a -- is a heart -- a heart pattern with Love's next to it. It's in the top part of the screenshot. A. Yep. Q. And the results are shown below that under "Shapes judged to be most similar"; is that right? A. Correct. Q. Okay. And the results also show underneath each one the trademark number followed by matching score after the colon. Do you see that? A. Yes. Q. And this is -- is this an example of what you told me before, that matching scores could be presented to the user? A. Yes. Q. And so is it the lower the score, the closer the match; is that correct? A. Typically, yes. *Id.* at 180-181:

Element	'529 Claim Recitation	Exemplary Citations in Reference
		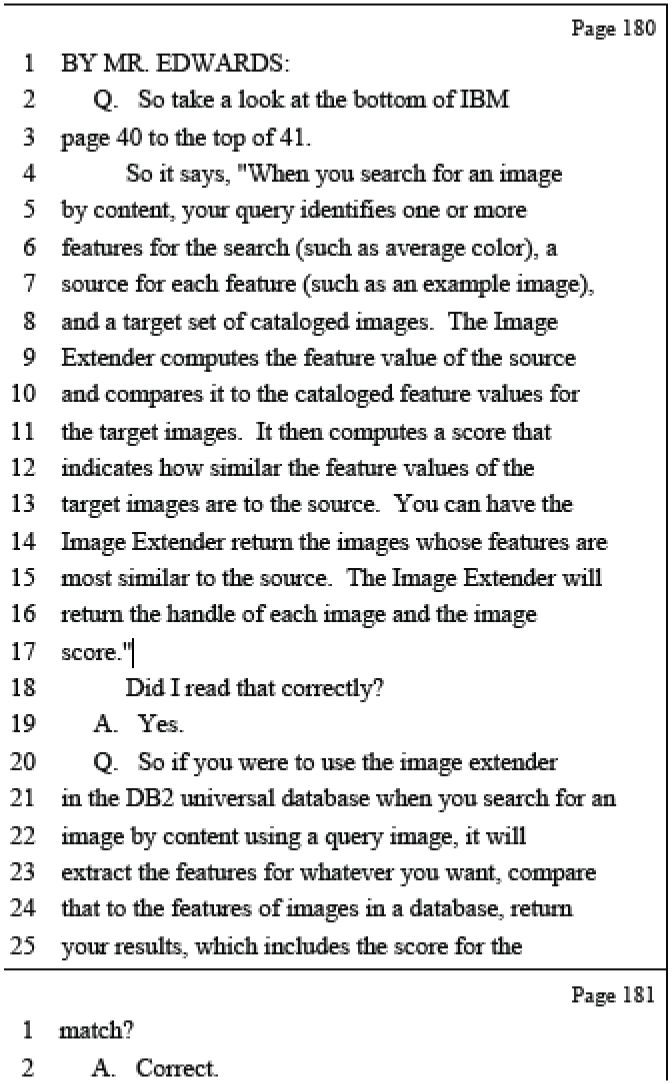 Page 180 BY MR. EDWARDS: Q. So take a look at the bottom of IBM page 40 to the top of 41. So it says, "When you search for an image by content, your query identifies one or more features for the search (such as average color), a source for each feature (such as an example image), and a target set of cataloged images. The Image Extender computes the feature value of the source and compares it to the cataloged feature values for the target images. It then computes a score that indicates how similar the feature values of the target images are to the source. You can have the Image Extender return the images whose features are most similar to the source. The Image Extender will return the handle of each image and the image score." Did I read that correctly? A. Yes. Q. So if you were to use the image extender in the DB2 universal database when you search for an image by content using a query image, it will extract the features for whatever you want, compare that to the features of images in a database, return your results, which includes the score for the Page 181 match? A. Correct.

Element	'529 Claim Recitation	Exemplary Citations in Reference
		Id. at 305: 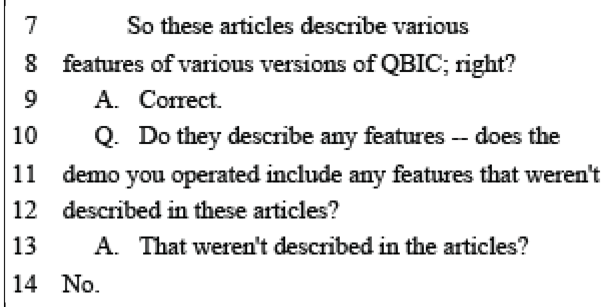 So these articles describe various features of various versions of QBIC; right? A. Correct. Q. Do they describe any features -- does the demo you operated include any features that weren't described in these articles? A. That weren't described in the articles? No. *Id.* at 35: 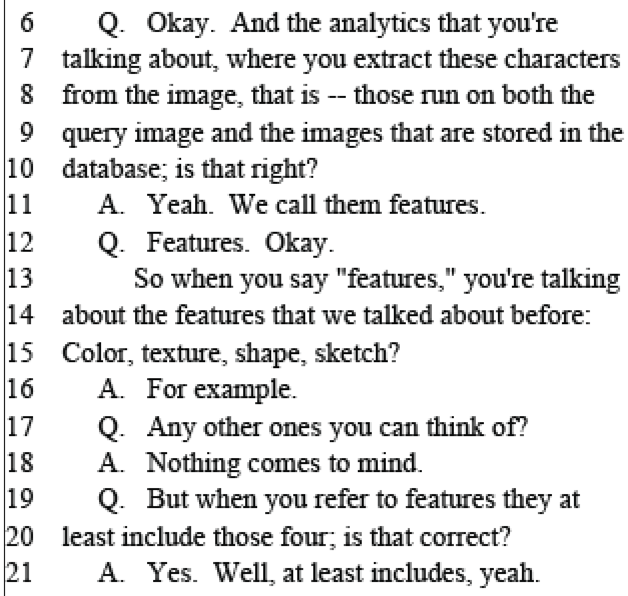 Q. Okay. And the analytics that you're talking about, where you extract these characters from the image, that is -- those run on both the query image and the images that are stored in the database; is that right? A. Yeah. We call them features. Q. Features. Okay. So when you say "features," you're talking about the features that we talked about before: Color, texture, shape, sketch? A. For example. Q. Any other ones you can think of? A. Nothing comes to mind. Q. But when you refer to features they at least include those four; is that correct? A. Yes. Well, at least includes, yeah.

Element	'529 Claim Recitation	Exemplary Citations in Reference
		Id. at 109-110: 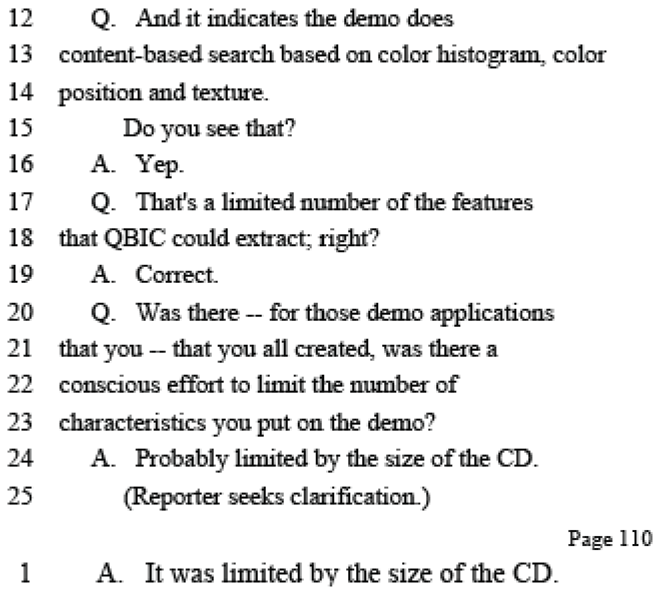 Q. And it indicates the demo does content-based search based on color histogram, color position and texture. Do you see that? A. Yep. Q. That's a limited number of the features that QBIC could extract; right? A. Correct. Q. Was there -- for those demo applications that you -- that you all created, was there a conscious effort to limit the number of characteristics you put on the demo? A. Probably limited by the size of the CD. (Reporter seeks clarification.) Page 110 A. It was limited by the size of the CD. *Id.* at 311: 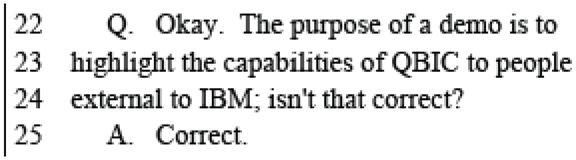 Q. Okay. The purpose of a demo is to highlight the capabilities of QBIC to people external to IBM; isn't that correct? A. Correct. *Id.* at 313-317:

Element	'529 Claim Recitation	Exemplary Citations in Reference
		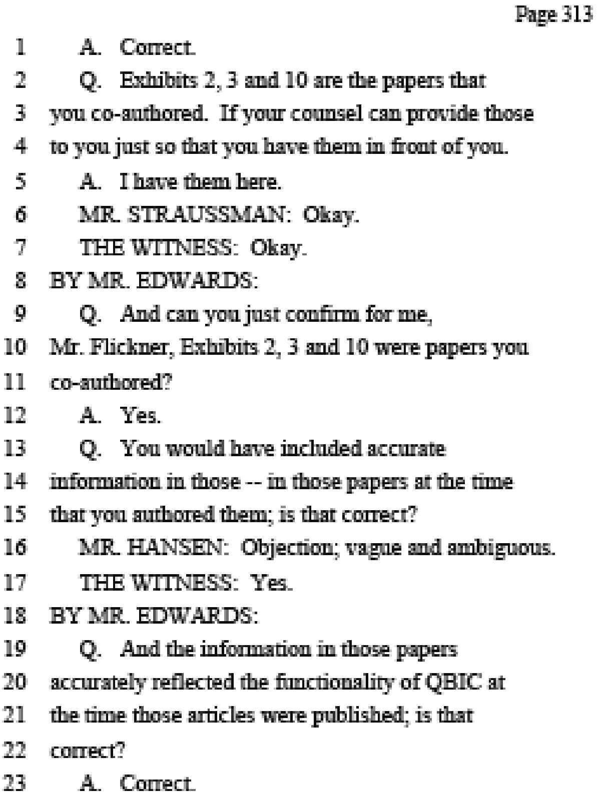 Page 313 A. Correct. Q. Exhibits 2, 3 and 10 are the papers that you co-authored. If your counsel can provide those to you just so that you have them in front of you. A. I have them here. MR. STRAUSSMAN: Okay. THE WITNESS: Okay. BY MR. EDWARDS: Q. And can you just confirm for me, Mr. Flickner, Exhibits 2, 3 and 10 were papers you co-authored? A. Yes. Q. You would have included accurate information in those -- in those papers at the time that you authored them; is that correct? MR. HANSEN: Objection; vague and ambiguous. THE WITNESS: Yes. BY MR. EDWARDS: Q. And the information in those papers accurately reflected the functionality of QBIC at the time those articles were published; is that correct? A. Correct.

		Page 314 BY MR. EDWARDS: Q. And each of those papers would have been peer reviewed before being published; correct? A. Yes. There was peer review. Q. Okay. Let's pick up Exhibit 3, please. A. Okay. Q. You have it -- A. Yeah. Q. -- handy? Okay. On the first page that is Bates-labeled IBM 893, in the right-hand column -- and as you're going down towards the bottom of the right-hand column, before you get to the line that crosses the middle of the page, there are three bullet points. Do you see those bullet points? A. Yeah. Q. It says, "QBIC allows queries on large image and video databases based on," bullet point one, "example images." Do you see that? A. Yeah. Q. Okay. And then also if you go to page 898. A. Okay. Page 315 Q. On the right side it says, "Representative Frame Generation." You see that? A. Yeah. Q. And for database population you can use r-frames of videos just as you can still image for both database population and for feature extraction for a input image; is that fair? MR. HANSEN: Objection; vague and ambiguous, lacks foundation. THE WITNESS: Yes. BY MR. EDWARDS: Q. Let's go to IBM 900. The bottom of the left-hand column to the top of the right-hand column. The bottom of the left-hand column starts, "We have described a prototype system." Do you see that? A. Yeah. Q. Okay. You keep going down, that paragraph refers to the technology described in this article [as read], "has been moved into a commercial standalone product, IBM's Ultimedia Manager, and is part of IBM's Digital Library and DB2 series of products." Do you agree with that?

Element	'529 Claim Recitation	Exemplary Citations in Reference
		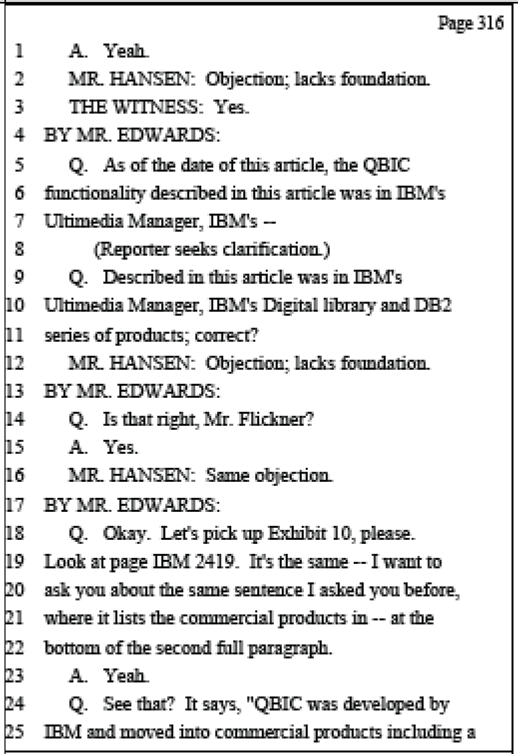 Page 316 A. Yeah. MR. HANSEN: Objection; lacks foundation. THE WITNESS: Yes. BY MR. EDWARDS: Q. As of the date of this article, the QBIC functionality described in this article was in IBM's Ultimedia Manager, IBM's -- (Reporter seeks clarification.) Q. Described in this article was in IBM's Ultimedia Manager, IBM's Digital library and DB2 series of products; correct? MR. HANSEN: Objection; lacks foundation. BY MR. EDWARDS: Q. Is that right, Mr. Flickner? A. Yes. MR. HANSEN: Same objection. BY MR. EDWARDS: Q. Okay. Let's pick up Exhibit 10, please. Look at page IBM 2419. It's the same -- I want to ask you about the same sentence I asked you before, where it lists the commercial products in -- at the bottom of the second full paragraph. A. Yeah. Q. See that? It says, "QBIC was developed by IBM and moved into commercial products including a 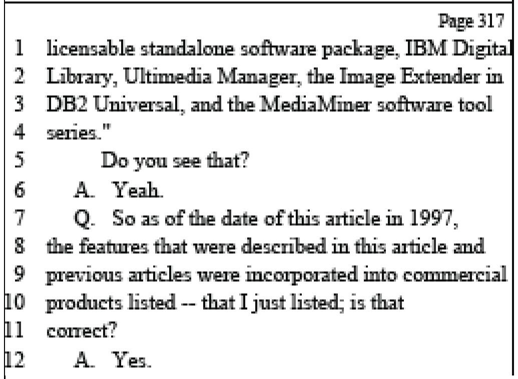 Page 317 licensable standalone software package, IBM Digital Library, Ultimedia Manager, the Image Extender in DB2 Universal, and the MediaMiner software tool series." Do you see that? A. Yeah. Q. So as of the date of this article in 1997, the features that were described in this article and previous articles were incorporated into commercial products listed -- that I just listed; is that correct? A. Yes.

Element	'529 Claim Recitation	Exemplary Citations in Reference
		Id. at 321: 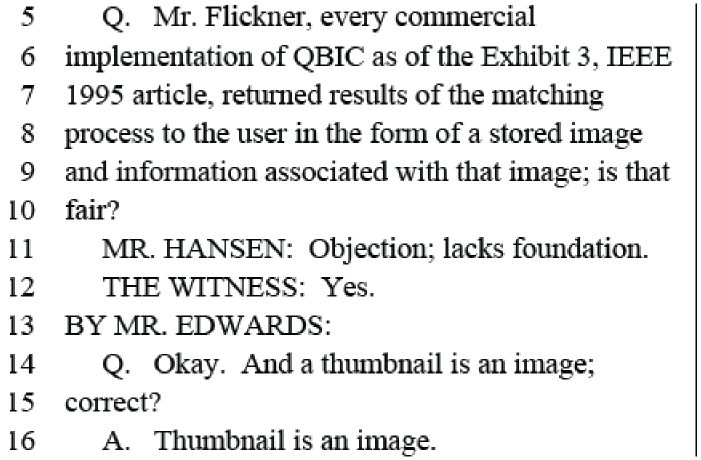 Q. Mr. Flickner, every commercial implementation of QBIC as of the Exhibit 3, IEEE 1995 article, returned results of the matching process to the user in the form of a stored image and information associated with that image; is that fair? MR. HANSEN: Objection; lacks foundation. THE WITNESS: Yes. BY MR. EDWARDS: Q. Okay. And a thumbnail is an image; correct? A. Thumbnail is an image. To the extent it is found that QBIC does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that QBIC does not anticipate this claim, QBIC renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine QBIC with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1F]	return the information to the network-enabled device; and	QBIC discloses return the information to the network-enabled device. The QBIC Project: Querying Images by Content Using Color, Texture, and Shape

Element	'529 Claim Recitation	Exemplary Citations in Reference
		(IBM 0002390–404) at 2395-96. 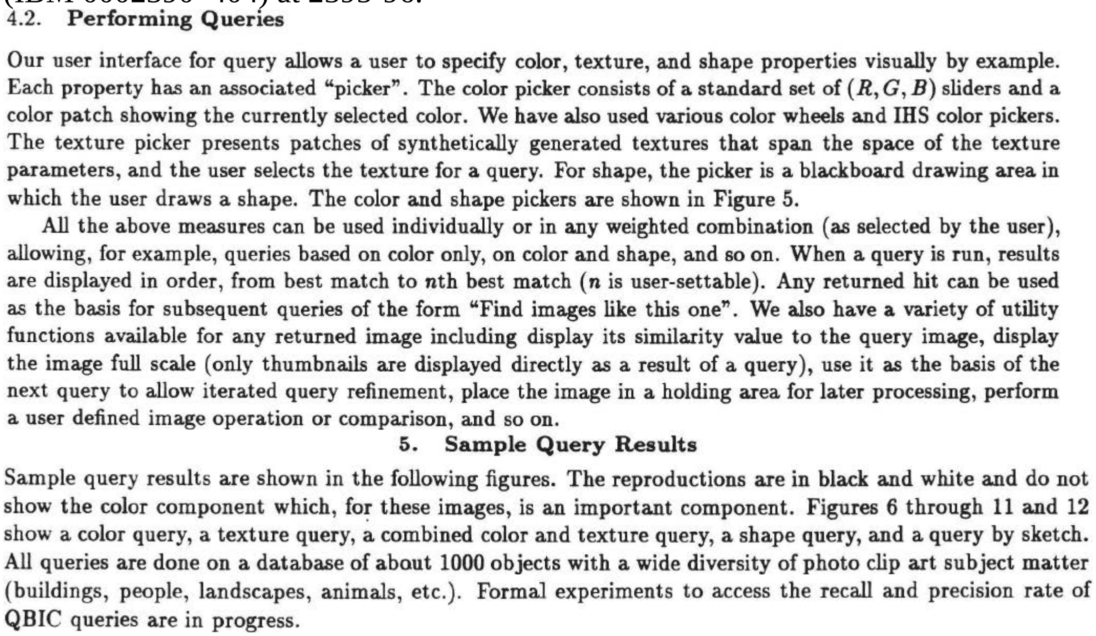 4.2. **Performing Queries** Our user interface for query allows a user to specify color, texture, and shape properties visually by example. Each property has an associated "picker". The color picker consists of a standard set of (R, G, B) sliders and a color patch showing the currently selected color. We have also used various color wheels and IHS color pickers. The texture picker presents patches of synthetically generated textures that span the space of the texture parameters, and the user selects the texture for a query. For shape, the picker is a blackboard drawing area in which the user draws a shape. The color and shape pickers are shown in Figure 5. All the above measures can be used individually or in any weighted combination (as selected by the user), allowing, for example, queries based on color only, on color and shape, and so on. When a query is run, results are displayed in order, from best match to nth best match (n is user-settable). Any returned hit can be used as the basis for subsequent queries of the form "Find images like this one". We also have a variety of utility functions available for any returned image including display its similarity value to the query image, display the image full scale (only thumbnails are displayed directly as a result of a query), use it as the basis of the next query to allow iterated query refinement, place the image in a holding area for later processing, perform a user defined image operation or comparison, and so on. 5. **Sample Query Results** Sample query results are shown in the following figures. The reproductions are in black and white and do not show the color component which, for these images, is an important component. Figures 6 through 11 and 12 show a color query, a texture query, a combined color and texture query, a shape query, and a query by sketch. All queries are done on a database of about 1000 objects with a wide diversity of photo clip art subject matter (buildings, people, landscapes, animals, etc.). Formal experiments to access the recall and precision rate of QBIC queries are in progress. 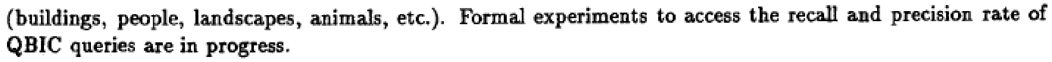 (buildings, people, landscapes, animals, etc.). Formal experiments to access the recall and precision rate of QBIC queries are in progress. The QBIC Project: Querying Images by Content Using Color, Texture, and Shape (IBM 0002390–404) at Figs. 7–9

Element	'529 Claim Recitation	Exemplary Citations in Reference
		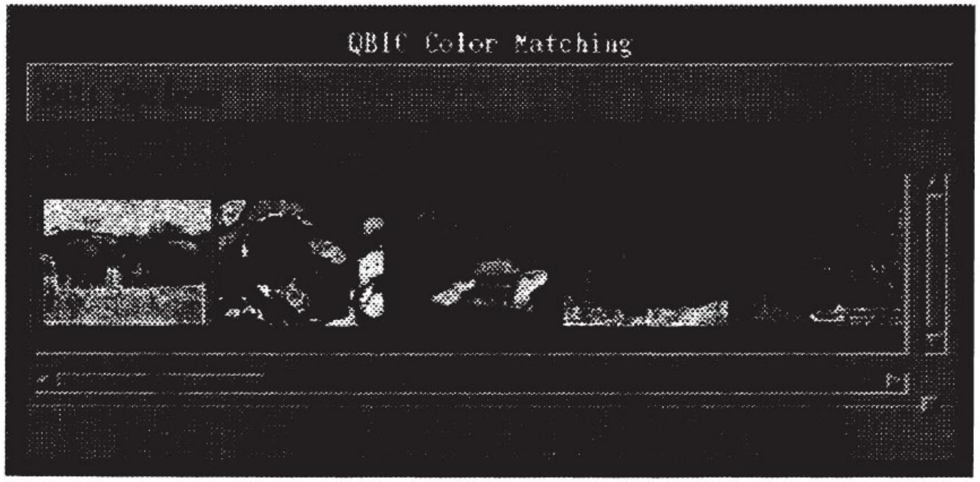 Figure 7: Result of query by example using color only.

Element	'529 Claim Recitation	Exemplary Citations in Reference
		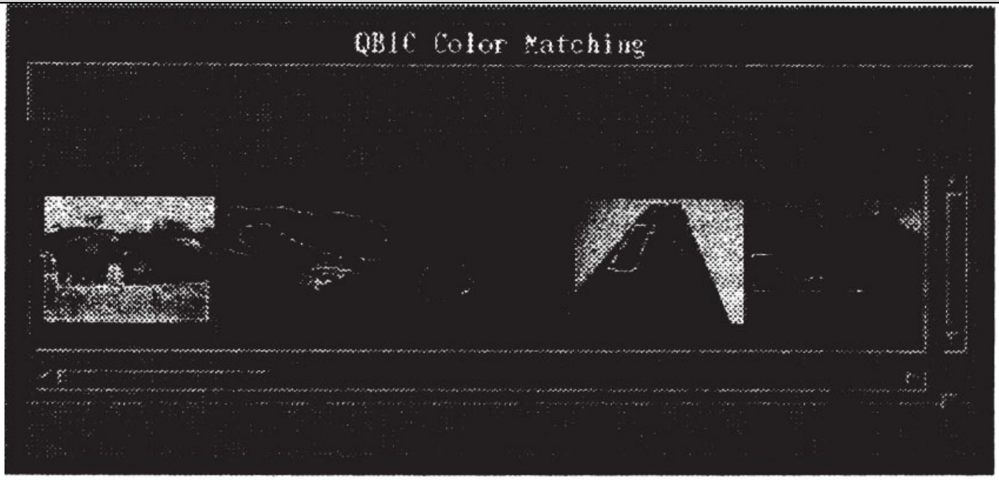 Figure 8: Result of query by example using texture only.

Element	'529 Claim Recitation	Exemplary Citations in Reference
		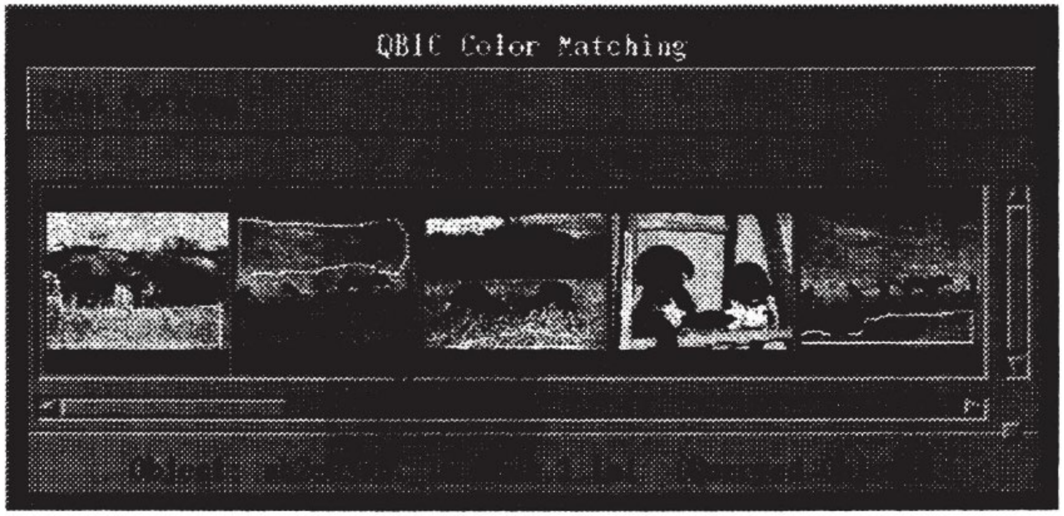 Figure 9: Result of query by example using color and texture. Query by Image and Video Content: The QBIC System (IBM 0000893–902): "To achieve this functionality, QBIC has two main components: database population (the process of creating an image database) and database query. During the population, images and videos are processed to extract features describing their content-colors, textures, shapes, and camera and object motion-and the features are stored in a database. During the query, the user composes a query graphically. Features are generated from the graphical query and then input to a matching engine that finds images or videos from the database with similar features. Figure 2 shows the system architecture." *Id.* at p. 3.

Element	'529 Claim Recitation	Exemplary Citations in Reference
		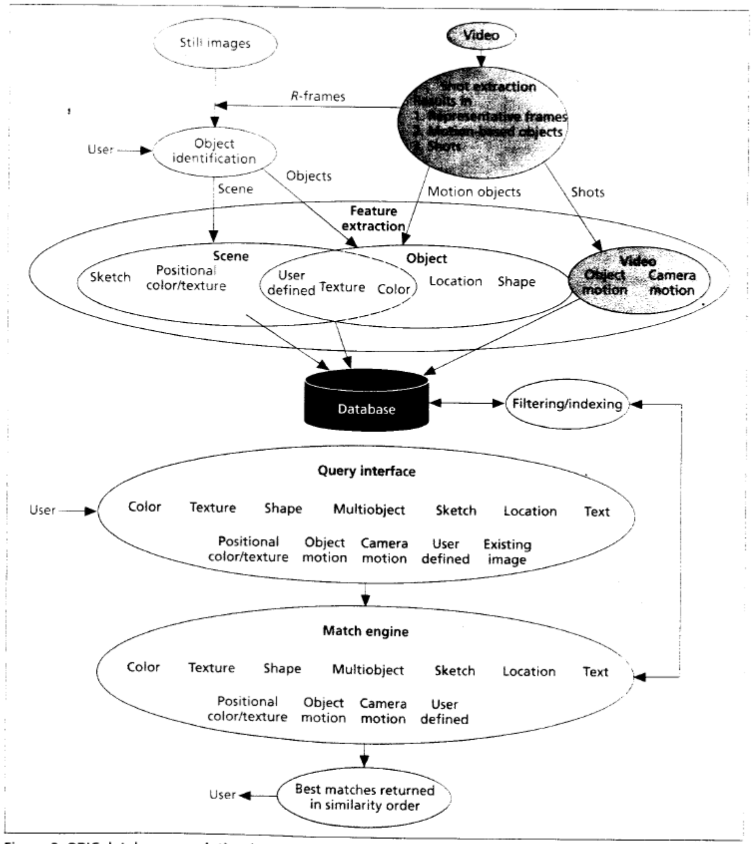 **Figure 2. QBIC database population (top) and query (bottom) architecture.** *Id*. at Fig. 2.

Element	'529 Claim Recitation	Exemplary Citations in Reference
		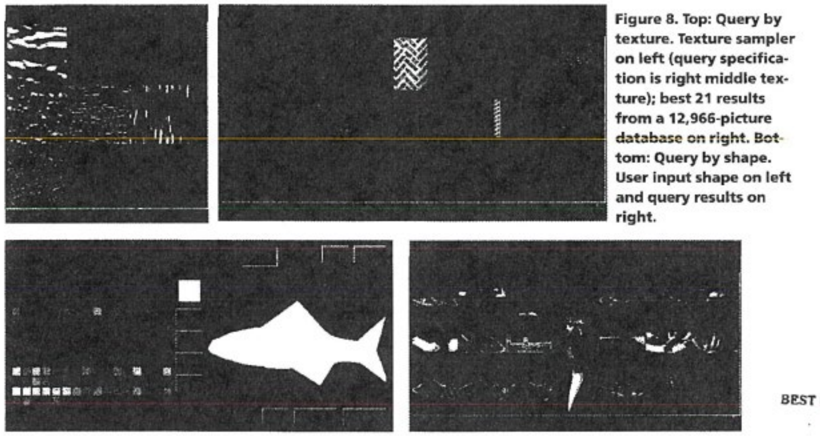 Figure 8. Top: Query by texture. Texture sampler on left (query specification is right middle texture); best 21 results from a 12,966-picture database on right. Bottom: Query by shape. User input shape on left and query results on right. *Id.* at Fig. 8 Image, Audio, and Video Extenders Administration and Programming (IBM 000001–617): "To make an image available for searching by content, you catalog the image. When you catalog an image, the Image Extender analyzes the image, by computing the feature values for the image, and stores the values in a QBIC catalog. When you search for an image by content, your query identifies one or more features for the search (such as average color), a source for each feature (such as an example image), and a target set of cataloged images. The Image Extender computes the feature value of the source and compares it to the cataloged feature values for the target images. It then computes a score that indicates how similar the feature values of the target images are to the source. You can have the Image Extender return the images whose features are most similar to the source. The Image Extender will return the handle of each image and the image score. You can also have the Image Extender return only the score of a single image." *Id*. at pp. 20–21.

Element	'529 Claim Recitation	Exemplary Citations in Reference
		QBIC Demo (IBM000747): 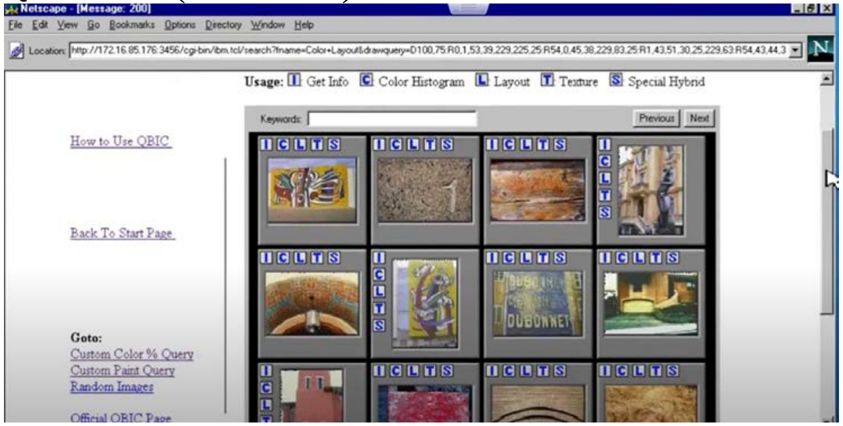 Location: http://172.16.85.176:3456/cgi-bin/ibm.tcl/search?fname=Color+Layout&drawquery=D100,75:R0,1,53,39,229,225,25:R54,0,45,38,229,83,25:R1,43,51,30,25,229,63:R54,43,44,3 QBIC Demo (IBM000747):  Location: http://172.16.85.176:3456/cgi-bin/ibm.tcl/info?france50%2Ejpg Image france50.jpg

Element	'529 Claim Recitation	Exemplary Citations in Reference
		Flickner Depo at 41-42: 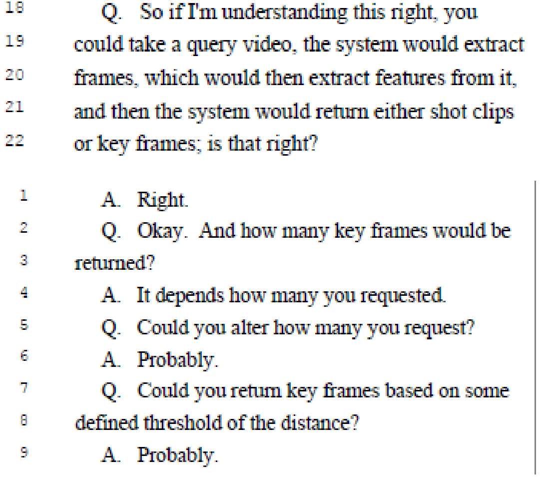 Q. So if I'm understanding this right, you could take a query video, the system would extract frames, which would then extract features from it, and then the system would return either shot clips or key frames; is that right? A. Right. Q. Okay. And how many key frames would be returned? A. It depends how many you requested. Q. Could you alter how many you request? A. Probably. Q. Could you return key frames based on some defined threshold of the distance? A. Probably. Flickner Depo at 50: 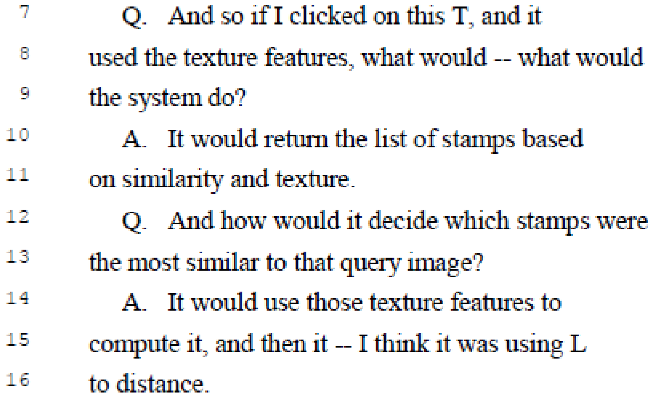 Q. And so if I clicked on this T, and it used the texture features, what would -- what would the system do? A. It would return the list of stamps based on similarity and texture. Q. And how would it decide which stamps were the most similar to that query image? A. It would use those texture features to compute it, and then it -- I think it was using L to distance. Flickner Depo at 162-163:

Element	'529 Claim Recitation	Exemplary Citations in Reference
		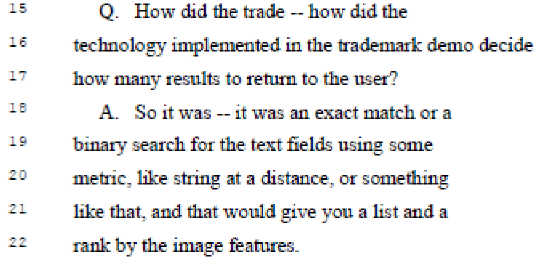 Q. How did the trade -- how did the technology implemented in the trademark demo decide how many results to return to the user? A. So it was -- it was an exact match or a binary search for the text fields using some metric, like string at a distance, or something like that, and that would give you a list and a rank by the image features. 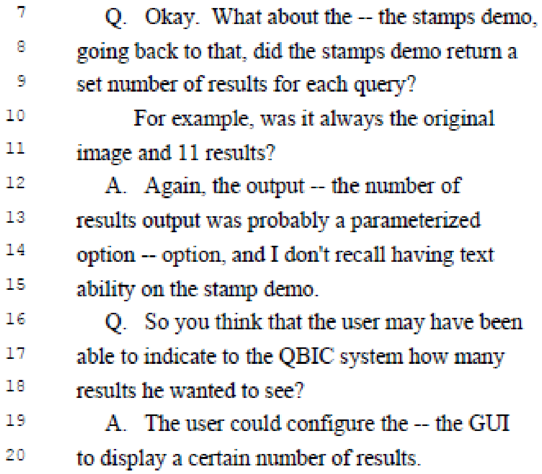 Q. Okay. What about the -- the stamps demo, going back to that, did the stamps demo return a set number of results for each query? For example, was it always the original image and 11 results? A. Again, the output -- the number of results output was probably a parameterized option -- option, and I don't recall having text ability on the stamp demo. Q. So you think that the user may have been able to indicate to the QBIC system how many results he wanted to see? A. The user could configure the -- the GUI to display a certain number of results. *See also* Flickner Depo at 25-26, 54-55, 78-79, 131-132, 201-202, 209-210. Flickner 8/29/23 Depo at 68–72:

		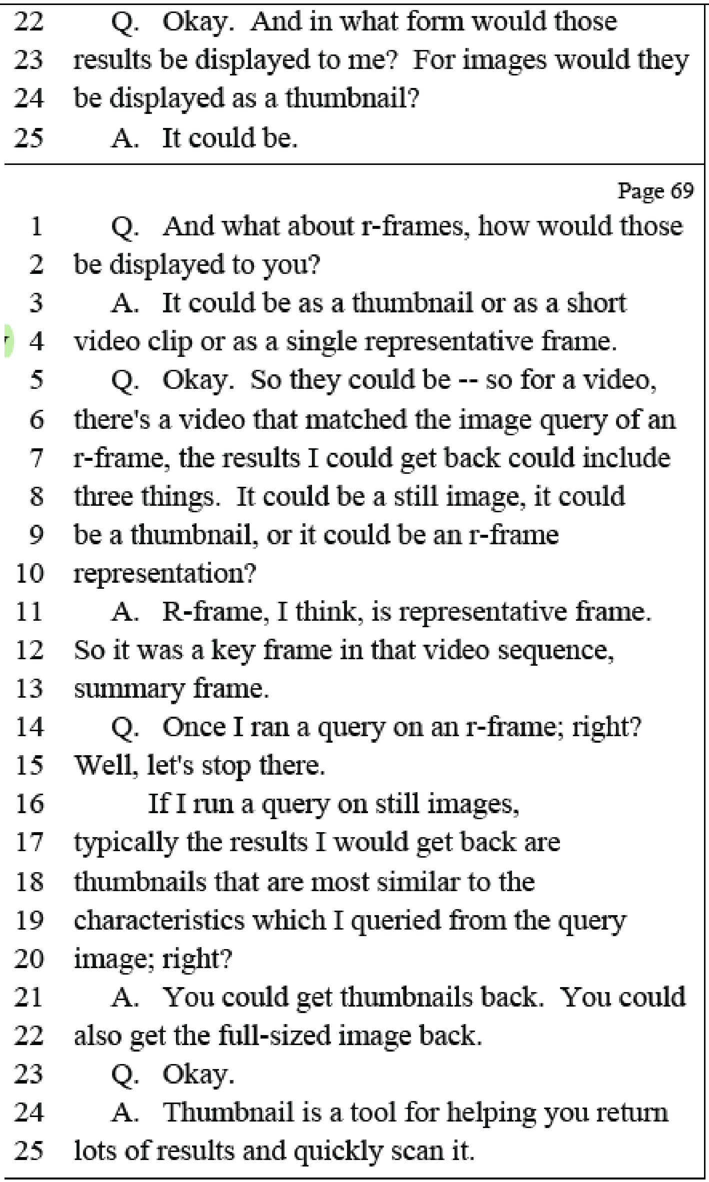 Q. Okay. And in what form would those results be displayed to me? For images would they be displayed as a thumbnail? A. It could be. Page 69 Q. And what about r-frames, how would those be displayed to you? A. It could be as a thumbnail or as a short video clip or as a single representative frame. Q. Okay. So they could be -- so for a video, there's a video that matched the image query of an r-frame, the results I could get back could include three things. It could be a still image, it could be a thumbnail, or it could be an r-frame representation? A. R-frame, I think, is representative frame. So it was a key frame in that video sequence, summary frame. Q. Once I ran a query on an r-frame; right? Well, let's stop there. If I run a query on still images, typically the results I would get back are thumbnails that are most similar to the characteristics which I queried from the query image; right? A. You could get thumbnails back. You could also get the full-sized image back. Q. Okay. A. Thumbnail is a tool for helping you return lots of results and quickly scan it.

Element	'529 Claim Recitation	Exemplary Citations in Reference
		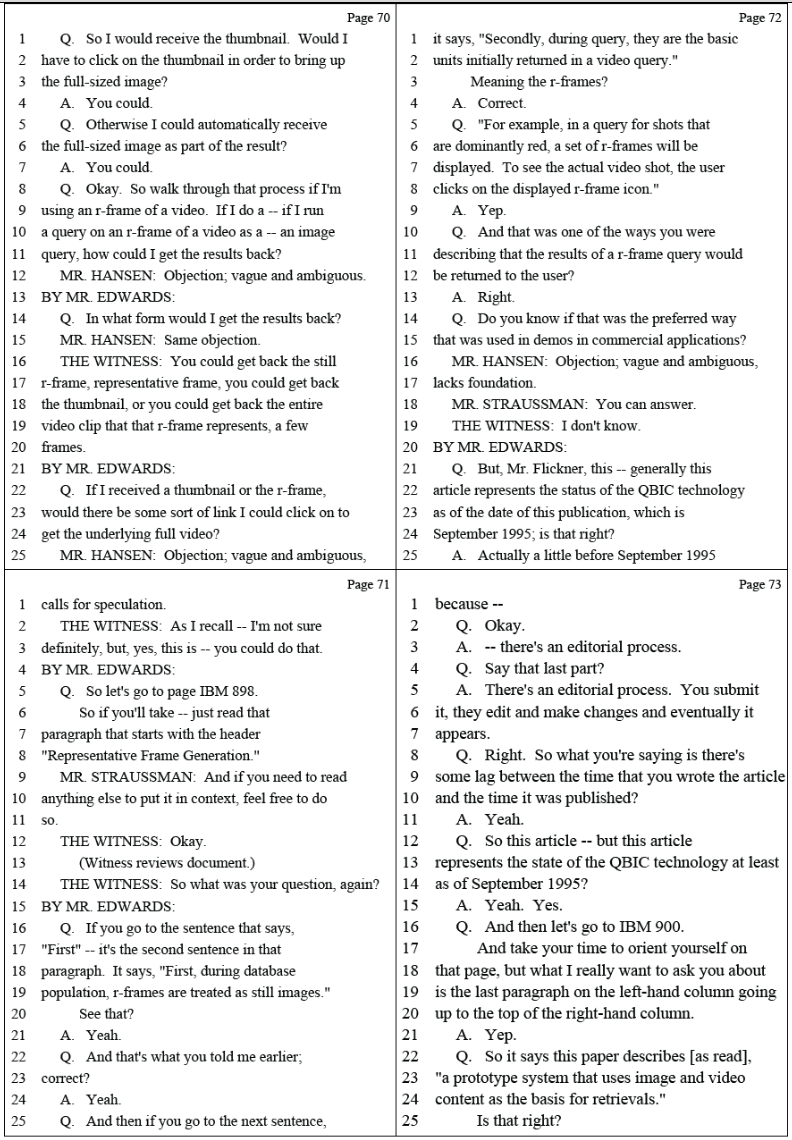 Page 70 Q. So I would receive the thumbnail. Would I have to click on the thumbnail in order to bring up the full-sized image? A. You could. Q. Otherwise I could automatically receive the full-sized image as part of the result? A. You could. Q. Okay. So walk through that process if I'm using an r-frame of a video. If I do a -- if I run a query on an r-frame of a video as a -- an image query, how could I get the results back? MR. HANSEN: Objection; vague and ambiguous. BY MR. EDWARDS: Q. In what form would I get the results back? MR. HANSEN: Same objection. THE WITNESS: You could get back the still r-frame, representative frame, you could get back the thumbnail, or you could get back the entire video clip that that r-frame represents, a few frames. BY MR. EDWARDS: Q. If I received a thumbnail or the r-frame, would there be some sort of link I could click on to get the underlying full video? MR. HANSEN: Objection; vague and ambiguous, Page 71 calls for speculation. THE WITNESS: As I recall -- I'm not sure definitely, but, yes, this is -- you could do that. BY MR. EDWARDS: Q. So let's go to page IBM 898. So if you'll take -- just read that paragraph that starts with the header "Representative Frame Generation." MR. STRAUSSMAN: And if you need to read anything else to put it in context, feel free to do so. THE WITNESS: Okay. (Witness reviews document.) THE WITNESS: So what was your question, again? BY MR. EDWARDS: Q. If you go to the sentence that says, "First" -- it's the second sentence in that paragraph. It says, "First, during database population, r-frames are treated as still images." See that? A. Yeah. Q. And that's what you told me earlier; correct? A. Yeah. Q. And then if you go to the next sentence, Page 72 it says, "Secondly, during query, they are the basic units initially returned in a video query." Meaning the r-frames? A. Correct. Q. "For example, in a query for shots that are dominantly red, a set of r-frames will be displayed. To see the actual video shot, the user clicks on the displayed r-frame icon." A. Yep. Q. And that was one of the ways you were describing that the results of a r-frame query would be returned to the user? A. Right. Q. Do you know if that was the preferred way that was used in demos in commercial applications? MR. HANSEN: Objection; vague and ambiguous, lacks foundation. MR. STRAUSSMAN: You can answer. THE WITNESS: I don't know. BY MR. EDWARDS: Q. But, Mr. Flickner, this -- generally this article represents the status of the QBIC technology as of the date of this publication, which is September 1995; is that right? A. Actually a little before September 1995 Page 73 because -- Q. Okay. A. -- there's an editorial process. Q. Say that last part? A. There's an editorial process. You submit it, they edit and make changes and eventually it appears. Q. Right. So what you're saying is there's some lag between the time that you wrote the article and the time it was published? A. Yeah. Q. So this article -- but this article represents the state of the QBIC technology at least as of September 1995? A. Yeah. Yes. Q. And then let's go to IBM 900. And take your time to orient yourself on that page, but what I really want to ask you about is the last paragraph on the left-hand column going up to the top of the right-hand column. A. Yep. Q. So it says this paper describes [as read], "a prototype system that uses image and video content as the basis for retrievals." Is that right?

Element	'529 Claim Recitation	Exemplary Citations in Reference
		Id. at 129 – 131: 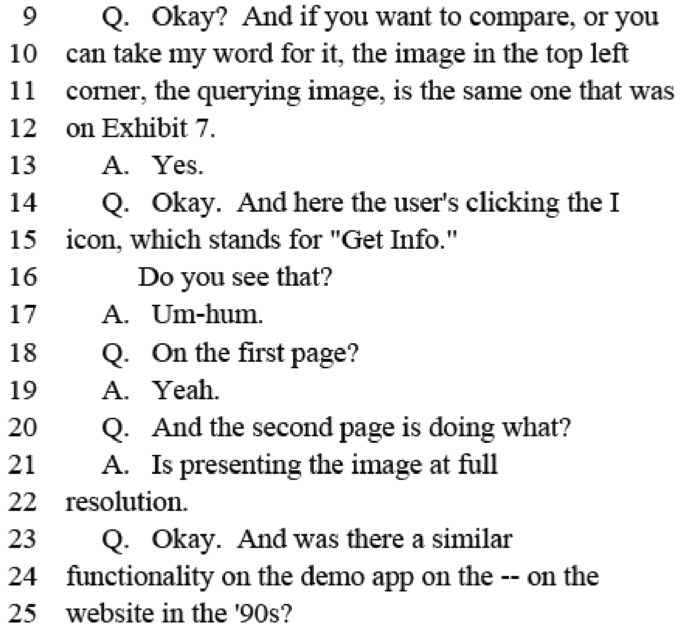 Q. Okay? And if you want to compare, or you can take my word for it, the image in the top left corner, the querying image, is the same one that was on Exhibit 7. A. Yes. Q. Okay. And here the user's clicking the I icon, which stands for "Get Info." Do you see that? A. Um-hum. Q. On the first page? A. Yeah. Q. And the second page is doing what? A. Is presenting the image at full resolution. Q. Okay. And was there a similar functionality on the demo app on the -- on the website in the '90s?

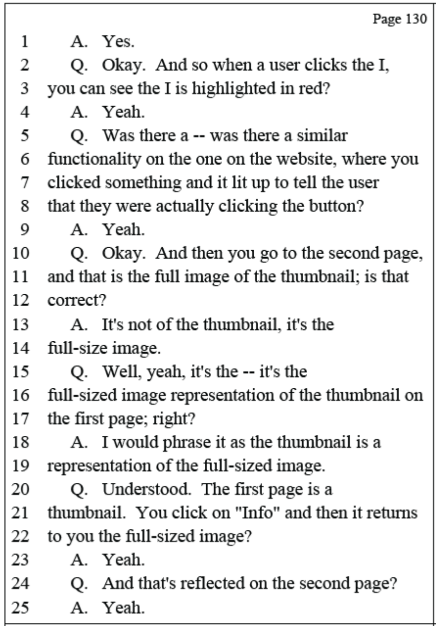

Page 130

A. Yes.
Q. Okay. And so when a user clicks the I, you can see the I is highlighted in red?
A. Yeah.
Q. Was there a -- was there a similar functionality on the one on the website, where you clicked something and it lit up to tell the user that they were actually clicking the button?
A. Yeah.
Q. Okay. And then you go to the second page, and that is the full image of the thumbnail; is that correct?
A. It's not of the thumbnail, it's the full-size image.
Q. Well, yeah, it's the -- it's the full-sized image representation of the thumbnail on the first page; right?
A. I would phrase it as the thumbnail is a representation of the full-sized image.
Q. Understood. The first page is a thumbnail. You click on "Info" and then it returns to you the full-sized image?
A. Yeah.
Q. And that's reflected on the second page?
A. Yeah.

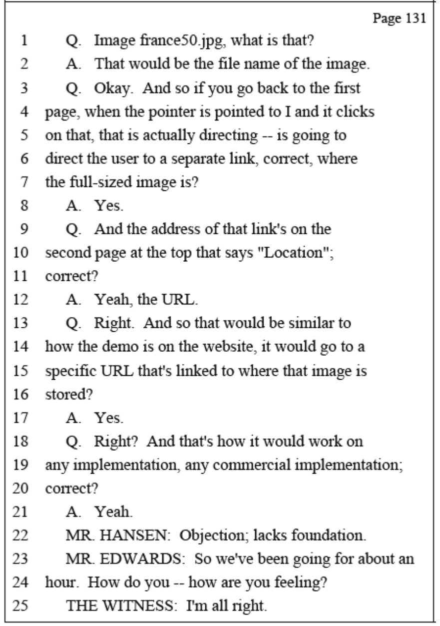

Page 131

Q. Image france50.jpg, what is that?
A. That would be the file name of the image.
Q. Okay. And so if you go back to the first page, when the pointer is pointed to I and it clicks on that, that is actually directing -- is going to direct the user to a separate link, correct, where the full-sized image is?
A. Yes.
Q. And the address of that link's on the second page at the top that says "Location"; correct?
A. Yeah, the URL.
Q. Right. And so that would be similar to how the demo is on the website, it would go to a specific URL that's linked to where that image is stored?
A. Yes.
Q. Right? And that's how it would work on any implementation, any commercial implementation; correct?
A. Yeah.
MR. HANSEN: Objection; lacks foundation.
MR. EDWARDS: So we've been going for about an hour. How do you -- how are you feeling?
THE WITNESS: I'm all right.

Element	'529 Claim Recitation	Exemplary Citations in Reference
		Id. at 161: 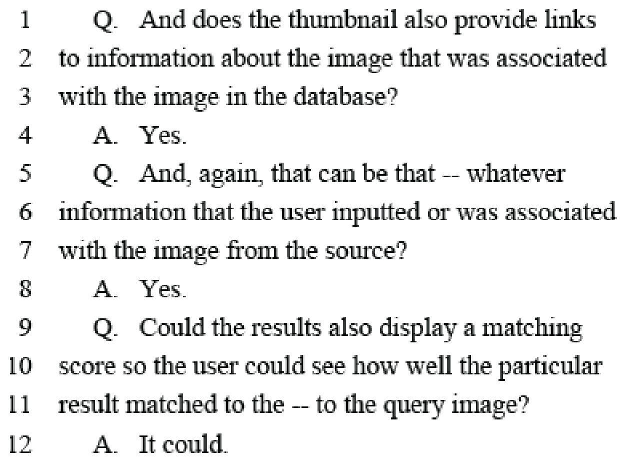 Q. And does the thumbnail also provide links to information about the image that was associated with the image in the database? A. Yes. Q. And, again, that can be that -- whatever information that the user inputted or was associated with the image from the source? A. Yes. Q. Could the results also display a matching score so the user could see how well the particular result matched to the -- to the query image? A. It could. *Id.* at 165:

Element	'529 Claim Recitation	Exemplary Citations in Reference
		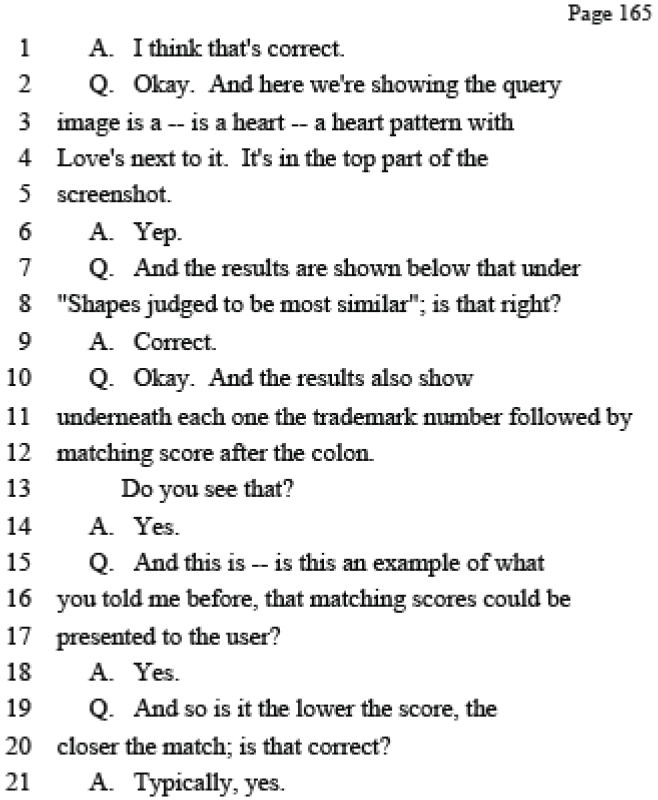 Page 165 A. I think that's correct. Q. Okay. And here we're showing the query image is a -- is a heart -- a heart pattern with Love's next to it. It's in the top part of the screenshot. A. Yep. Q. And the results are shown below that under "Shapes judged to be most similar"; is that right? A. Correct. Q. Okay. And the results also show underneath each one the trademark number followed by matching score after the colon. Do you see that? A. Yes. Q. And this is -- is this an example of what you told me before, that matching scores could be presented to the user? A. Yes. Q. And so is it the lower the score, the closer the match; is that correct? A. Typically, yes. *Id*. at 180-181:

Element	'529 Claim Recitation	Exemplary Citations in Reference
		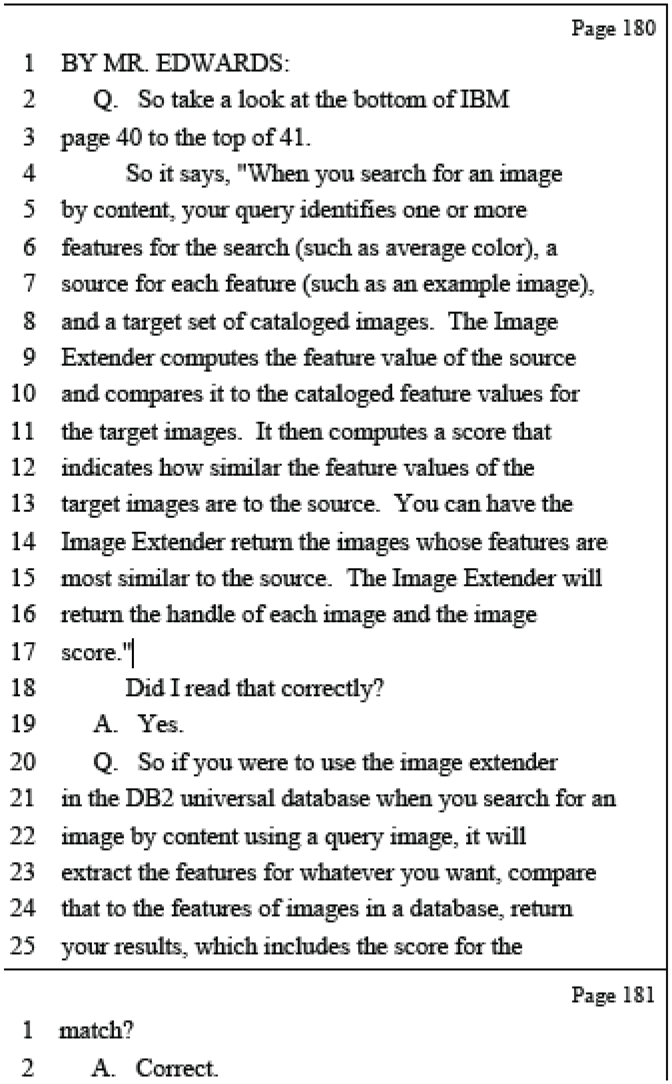 Page 180 BY MR. EDWARDS: Q. So take a look at the bottom of IBM page 40 to the top of 41. So it says, "When you search for an image by content, your query identifies one or more features for the search (such as average color), a source for each feature (such as an example image), and a target set of cataloged images. The Image Extender computes the feature value of the source and compares it to the cataloged feature values for the target images. It then computes a score that indicates how similar the feature values of the target images are to the source. You can have the Image Extender return the images whose features are most similar to the source. The Image Extender will return the handle of each image and the image score." Did I read that correctly? A. Yes. Q. So if you were to use the image extender in the DB2 universal database when you search for an image by content using a query image, it will extract the features for whatever you want, compare that to the features of images in a database, return your results, which includes the score for the Page 181 match? A. Correct.

Element	'529 Claim Recitation	Exemplary Citations in Reference
		Id. at 305: 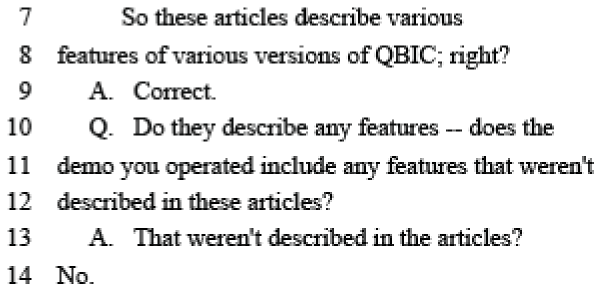 So these articles describe various features of various versions of QBIC; right? A. Correct. Q. Do they describe any features -- does the demo you operated include any features that weren't described in these articles? A. That weren't described in the articles? No. *Id.* at 35: 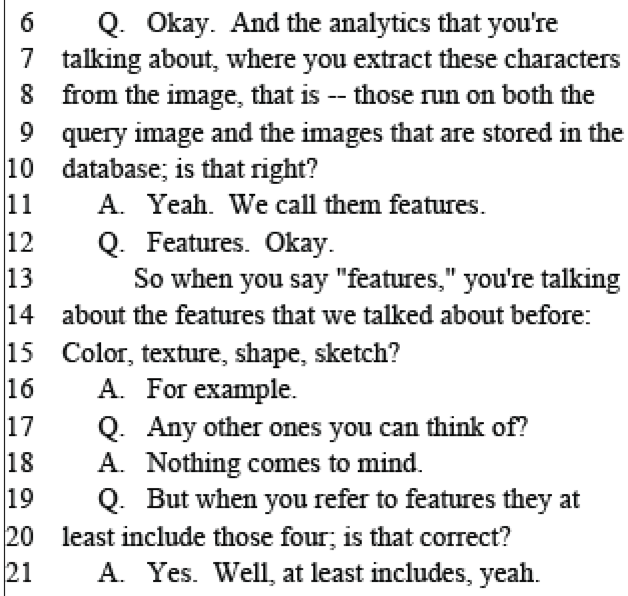 Q. Okay. And the analytics that you're talking about, where you extract these characters from the image, that is -- those run on both the query image and the images that are stored in the database; is that right? A. Yeah. We call them features. Q. Features. Okay. So when you say "features," you're talking about the features that we talked about before: Color, texture, shape, sketch? A. For example. Q. Any other ones you can think of? A. Nothing comes to mind. Q. But when you refer to features they at least include those four; is that correct? A. Yes. Well, at least includes, yeah.

Element	'529 Claim Recitation	Exemplary Citations in Reference
		Id. at 109-110: 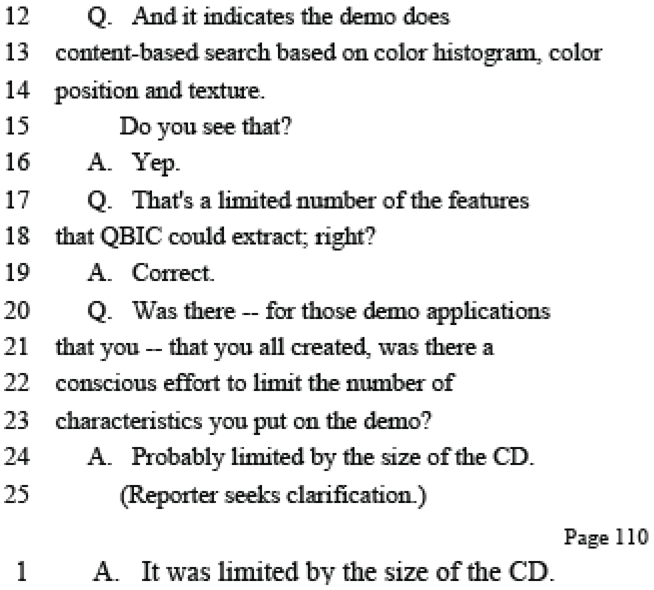 Q. And it indicates the demo does content-based search based on color histogram, color position and texture. Do you see that? A. Yep. Q. That's a limited number of the features that QBIC could extract; right? A. Correct. Q. Was there -- for those demo applications that you -- that you all created, was there a conscious effort to limit the number of characteristics you put on the demo? A. Probably limited by the size of the CD. (Reporter seeks clarification.) Page 110 A. It was limited by the size of the CD. *Id.* at 311: 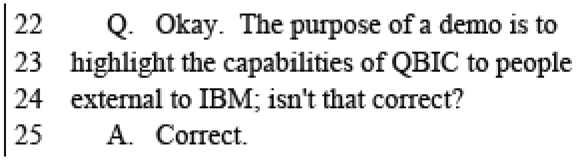 Q. Okay. The purpose of a demo is to highlight the capabilities of QBIC to people external to IBM; isn't that correct? A. Correct. *Id.* at 313-317:

Element	'529 Claim Recitation	Exemplary Citations in Reference
		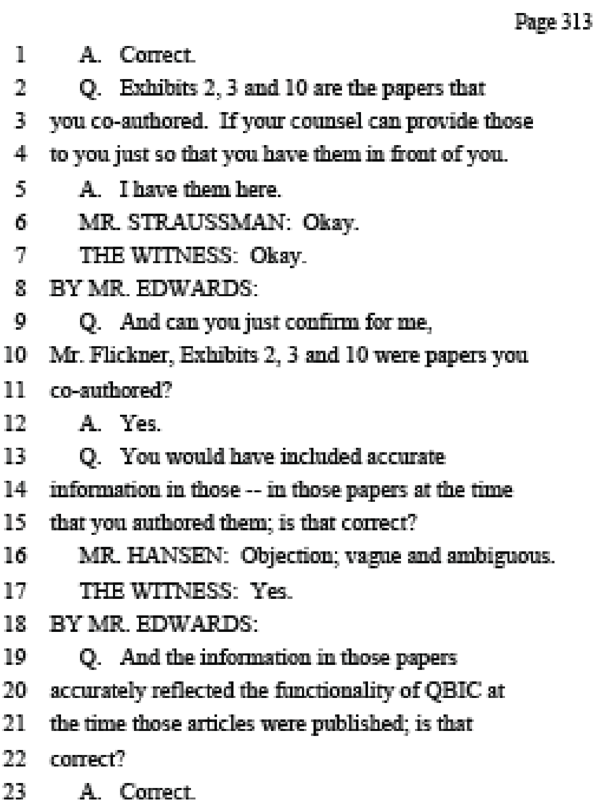Page 313 A. Correct. Q. Exhibits 2, 3 and 10 are the papers that you co-authored. If your counsel can provide those to you just so that you have them in front of you. A. I have them here. MR. STRAUSSMAN: Okay. THE WITNESS: Okay. BY MR. EDWARDS: Q. And can you just confirm for me, Mr. Flickner, Exhibits 2, 3 and 10 were papers you co-authored? A. Yes. Q. You would have included accurate information in those -- in those papers at the time that you authored them; is that correct? MR. HANSEN: Objection; vague and ambiguous. THE WITNESS: Yes. BY MR. EDWARDS: Q. And the information in those papers accurately reflected the functionality of QBIC at the time those articles were published; is that correct? A. Correct.

		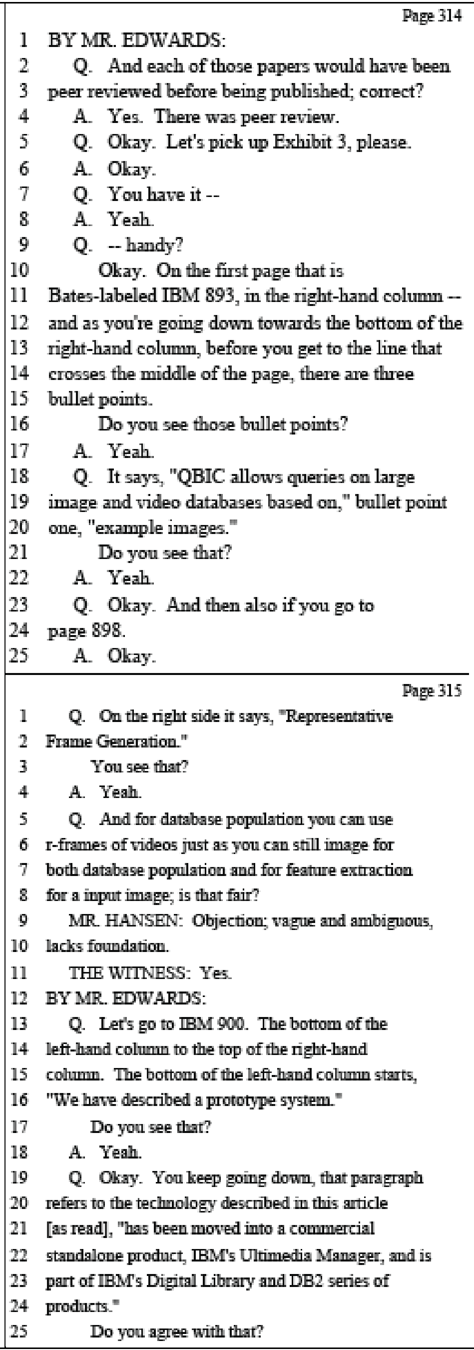 Page 314 BY MR. EDWARDS: Q. And each of those papers would have been peer reviewed before being published; correct? A. Yes. There was peer review. Q. Okay. Let's pick up Exhibit 3, please. A. Okay. Q. You have it -- A. Yeah. Q. -- handy? Okay. On the first page that is Bates-labeled IBM 893, in the right-hand column -- and as you're going down towards the bottom of the right-hand column, before you get to the line that crosses the middle of the page, there are three bullet points. Do you see those bullet points? A. Yeah. Q. It says, "QBIC allows queries on large image and video databases based on," bullet point one, "example images." Do you see that? A. Yeah. Q. Okay. And then also if you go to page 898. A. Okay. Page 315 Q. On the right side it says, "Representative Frame Generation." You see that? A. Yeah. Q. And for database population you can use r-frames of videos just as you can still image for both database population and for feature extraction for a input image; is that fair? MR. HANSEN: Objection; vague and ambiguous, lacks foundation. THE WITNESS: Yes. BY MR. EDWARDS: Q. Let's go to IBM 900. The bottom of the left-hand column to the top of the right-hand column. The bottom of the left-hand column starts, "We have described a prototype system." Do you see that? A. Yeah. Q. Okay. You keep going down, that paragraph refers to the technology described in this article [as read], "has been moved into a commercial standalone product, IBM's Ultimedia Manager, and is part of IBM's Digital Library and DB2 series of products." Do you agree with that?

Element	'529 Claim Recitation	Exemplary Citations in Reference
		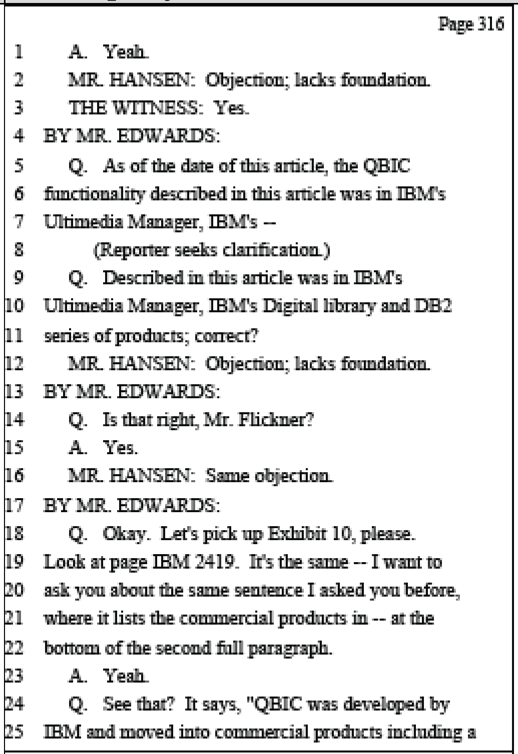 Page 316 A. Yeah. MR. HANSEN: Objection; lacks foundation. THE WITNESS: Yes. BY MR. EDWARDS: Q. As of the date of this article, the QBIC functionality described in this article was in IBM's Ultimedia Manager, IBM's -- (Reporter seeks clarification.) Q. Described in this article was in IBM's Ultimedia Manager, IBM's Digital library and DB2 series of products; correct? MR. HANSEN: Objection; lacks foundation. BY MR. EDWARDS: Q. Is that right, Mr. Flickner? A. Yes. MR. HANSEN: Same objection. BY MR. EDWARDS: Q. Okay. Let's pick up Exhibit 10, please. Look at page IBM 2419. It's the same -- I want to ask you about the same sentence I asked you before, where it lists the commercial products in -- at the bottom of the second full paragraph. A. Yeah. Q. See that? It says, "QBIC was developed by IBM and moved into commercial products including a 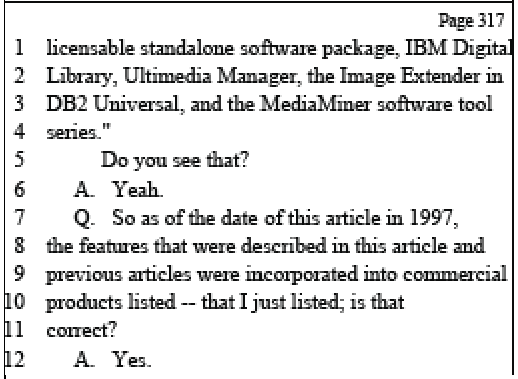 Page 317 licensable standalone software package, IBM Digital Library, Ultimedia Manager, the Image Extender in DB2 Universal, and the MediaMiner software tool series." Do you see that? A. Yeah. Q. So as of the date of this article in 1997, the features that were described in this article and previous articles were incorporated into commercial products listed -- that I just listed; is that correct? A. Yes.

Element	'529 Claim Recitation	Exemplary Citations in Reference
		Id. at 321: 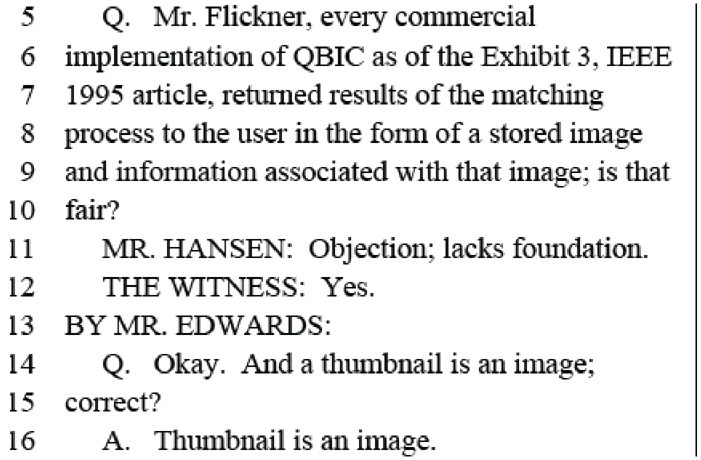 Q. Mr. Flickner, every commercial implementation of QBIC as of the Exhibit 3, IEEE 1995 article, returned results of the matching process to the user in the form of a stored image and information associated with that image; is that fair? MR. HANSEN: Objection; lacks foundation. THE WITNESS: Yes. BY MR. EDWARDS: Q. Okay. And a thumbnail is an image; correct? A. Thumbnail is an image. Flickner 8/29/23 Depo Ex. 7 at 1:

Element	'529 Claim Recitation	Exemplary Citations in Reference
		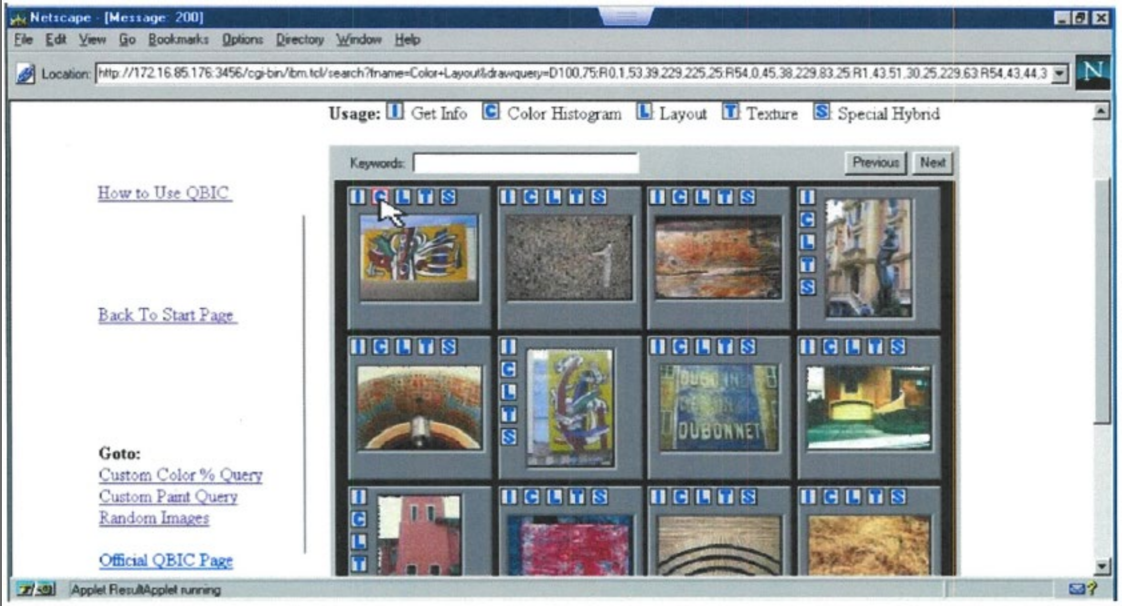 Location: http://172.16.85.176:3456/cgi-bin/ibm.tcl/search?fname=Color+Layout&drawquery=D100,75:R0,1,53,39,229,225,25:R54,0,45,38,229,83,25:R1,43,51,30,25,229,63:R54,43,44,3 Usage: I Get Info C Color Histogram L Layout T Texture S Special Hybrid Keywords: Goto: Custom Color % Query Custom Paint Query Random Images Official QBIC Page Applet ResultApplet running Flickner 8/29/23 Depo Ex. 8 at 1:

Element	'529 Claim Recitation	Exemplary Citations in Reference
		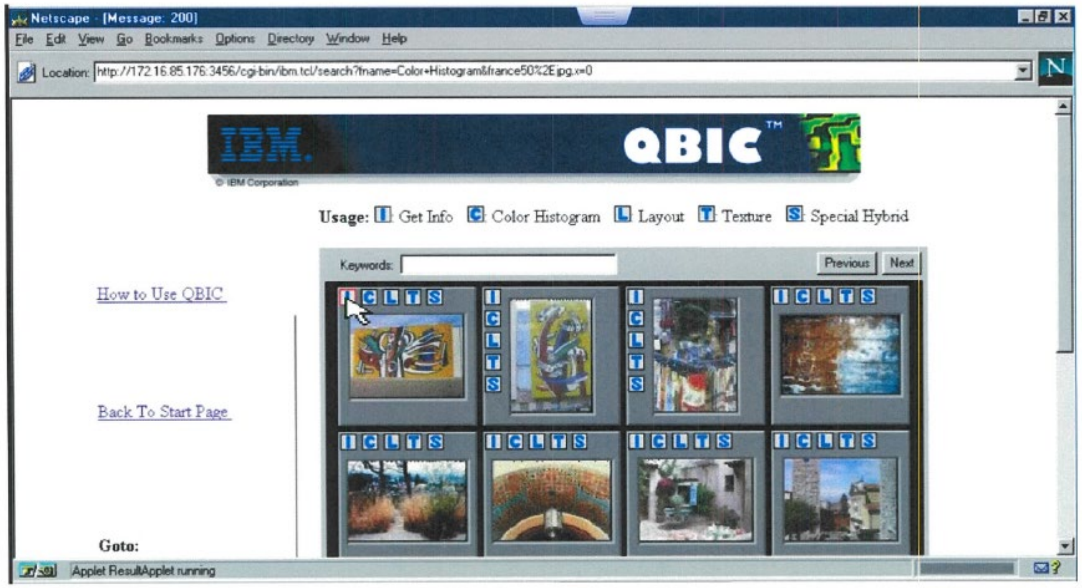 Location: http://172.16.85.176:3456/cgi-bin/ibm.tcl/search?fname=Color+Histogram&france50%2Ejpg.x=0 Applet ResultApplet running *Id*. at 2:

Element	'529 Claim Recitation	Exemplary Citations in Reference
		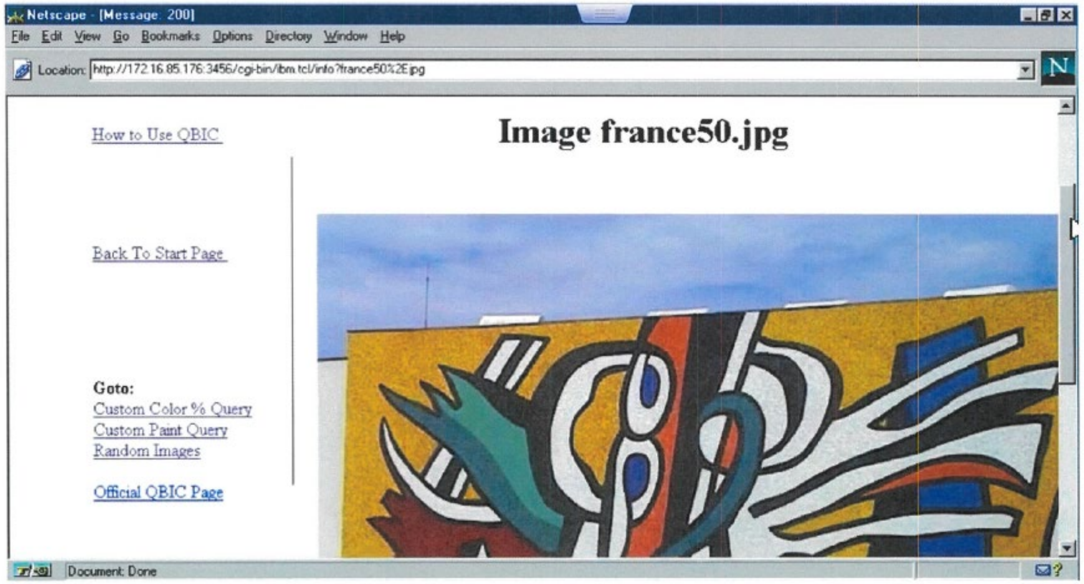 http://172.16.85.176:3456/cgi-bin/ibm.tcl/info?france50%2Ejpg To the extent it is found that QBIC does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that QBIC does not anticipate this claim, QBIC renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine QBIC with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1G]	the network-enabled device further programmed to present the information related to the	QBIC discloses the network-enabled device further programmed to present the information related to the object to a user.

Element	'529 Claim Recitation	Exemplary Citations in Reference
	object to a user.	*See* Limitation 1F. The QBIC Project: Querying Images by Content Using Color, Texture, and Shape (IBM 0002390–404) at 2396 4.1. **Object/element identification** Our user interface for object identification, shown in figure 4, allows for user assisted object identification. This interface is made up of four parts (listed from top to bottom): the window menu, the tool selection buttons, the image display area, and the object entry area. From the File menu selection the user has the ability to save the current image and objects and load another image. The Options menu contains selections for tuning panels for Snakes, Flood-Fill, as well as a Zoom selection, useful for working with very large or small images. The QBIC Project: Querying Images by Content Using Color, Texture, and Shape (IBM 0002390–404) at 2400, Fig. 4 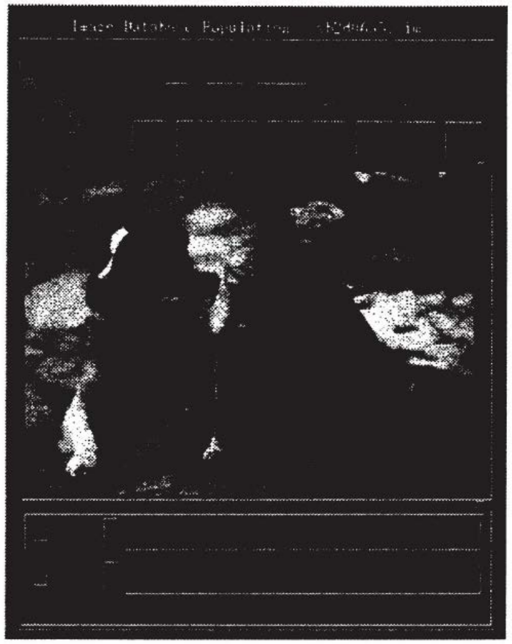 Figure 4: Object Identification Tool. Note for example, the right dog's ear, which has been outlined.

Element	'529 Claim Recitation	Exemplary Citations in Reference
		The QBIC Project: Querying Images by Content Using Color, Texture, and Shape (IBM 0002390–404) at 8 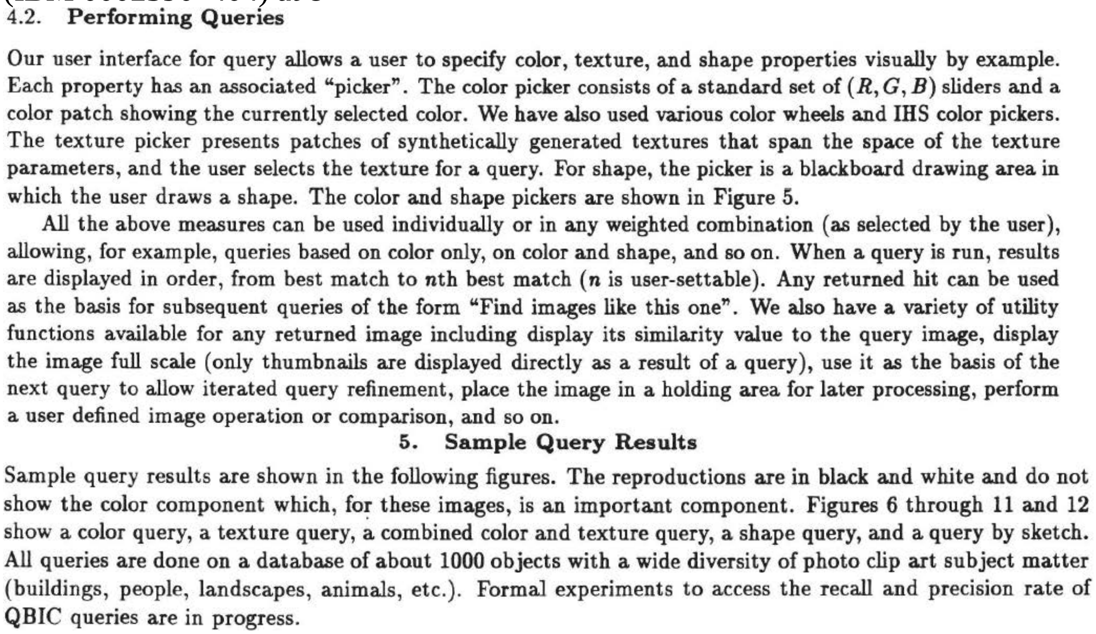 4.2. **Performing Queries** Our user interface for query allows a user to specify color, texture, and shape properties visually by example. Each property has an associated "picker". The color picker consists of a standard set of (R, G, B) sliders and a color patch showing the currently selected color. We have also used various color wheels and IHS color pickers. The texture picker presents patches of synthetically generated textures that span the space of the texture parameters, and the user selects the texture for a query. For shape, the picker is a blackboard drawing area in which the user draws a shape. The color and shape pickers are shown in Figure 5. All the above measures can be used individually or in any weighted combination (as selected by the user), allowing, for example, queries based on color only, on color and shape, and so on. When a query is run, results are displayed in order, from best match to nth best match (n is user-settable). Any returned hit can be used as the basis for subsequent queries of the form "Find images like this one". We also have a variety of utility functions available for any returned image including display its similarity value to the query image, display the image full scale (only thumbnails are displayed directly as a result of a query), use it as the basis of the next query to allow iterated query refinement, place the image in a holding area for later processing, perform a user defined image operation or comparison, and so on. 5. **Sample Query Results** Sample query results are shown in the following figures. The reproductions are in black and white and do not show the color component which, for these images, is an important component. Figures 6 through 11 and 12 show a color query, a texture query, a combined color and texture query, a shape query, and a query by sketch. All queries are done on a database of about 1000 objects with a wide diversity of photo clip art subject matter (buildings, people, landscapes, animals, etc.). Formal experiments to access the recall and precision rate of QBIC queries are in progress. The QBIC Project: Querying Images by Content Using Color, Texture, and Shape (IBM 0002390–404) at Figs. 7–9

Element	'529 Claim Recitation	Exemplary Citations in Reference
		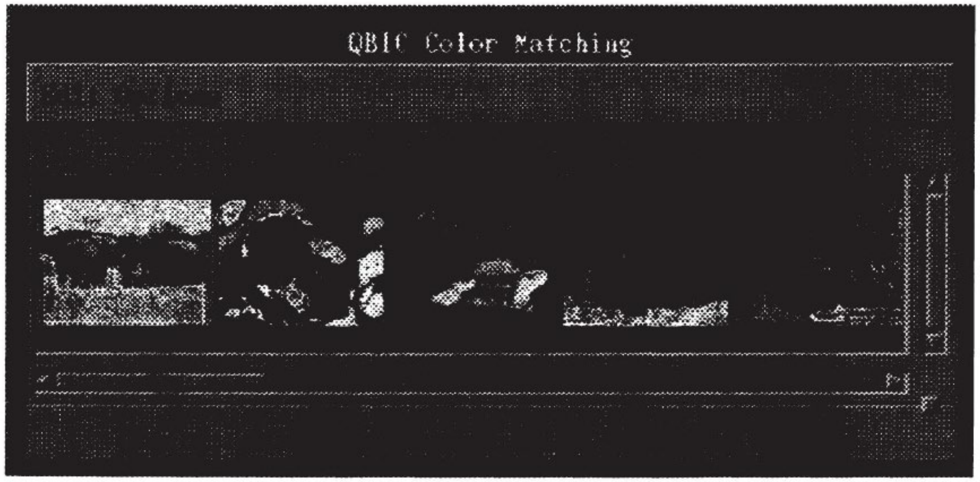 Figure 7: Result of query by example using color only.

Element	'529 Claim Recitation	Exemplary Citations in Reference
		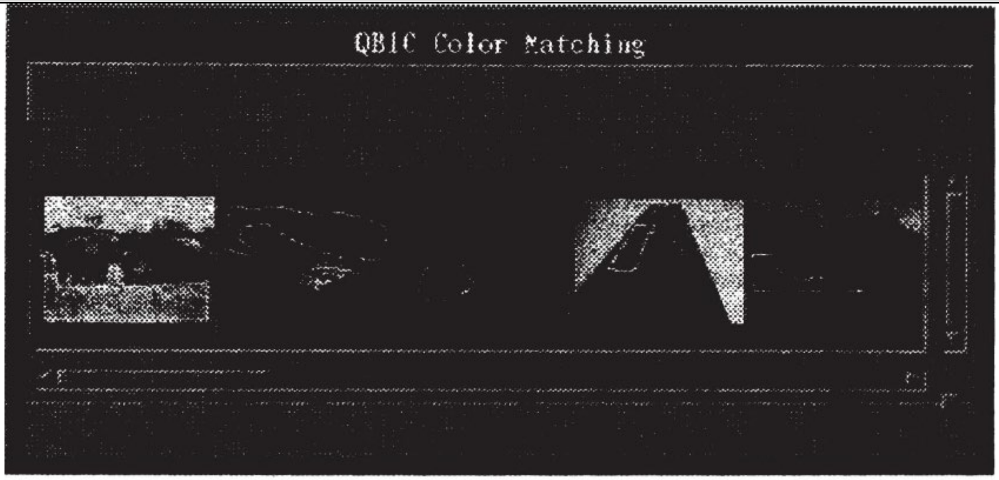 Figure 8: Result of query by example using texture only.

Element	'529 Claim Recitation	Exemplary Citations in Reference
		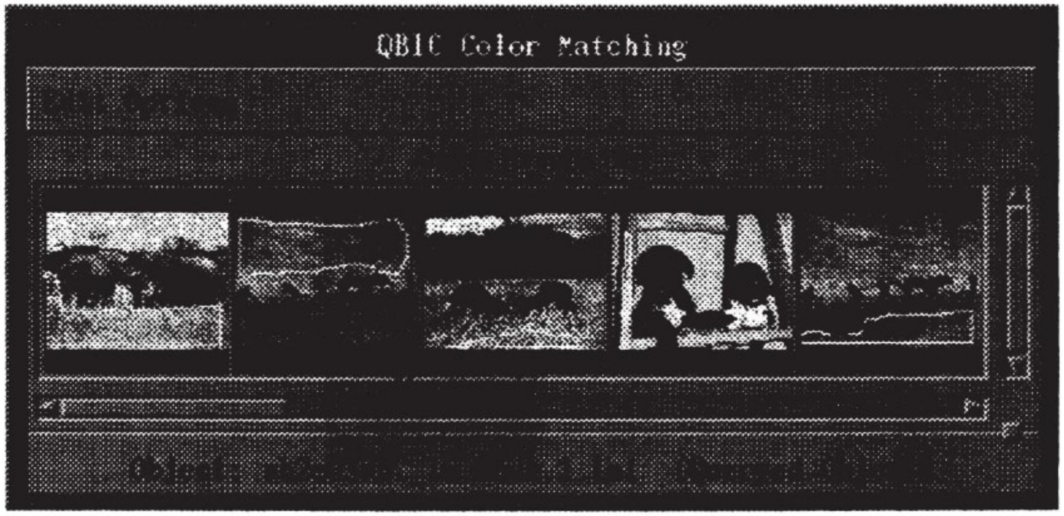 Figure 9: Result of query by example using color and texture.

Element	'529 Claim Recitation	Exemplary Citations in Reference
		QBIC Demo (IBM000747): 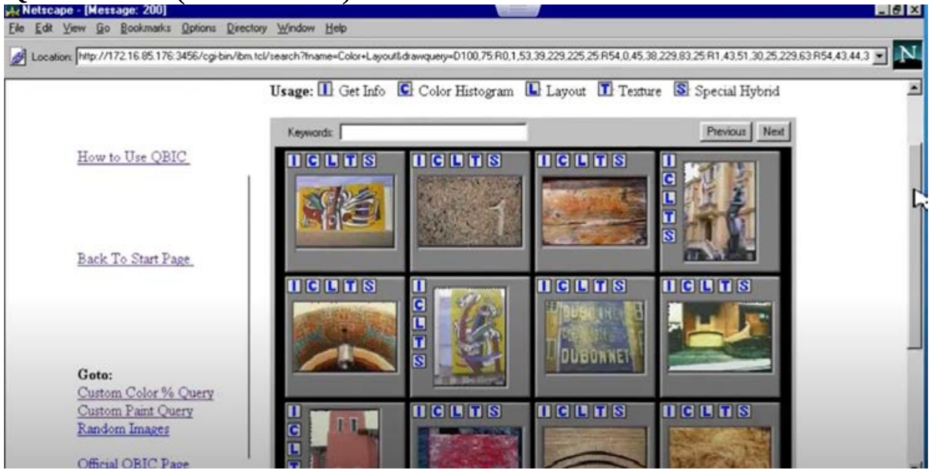 Location: http://172.16.85.176:3456/cgi-bin/ibm.tcl/search?fname=Color+Layout&drawquery=D100,75:R0,1,53,39,229,225,25:R54,0,45,38,229,83,25:R1,43,51,30,25,229,63:R54,43,44,3 QBIC Demo (IBM000747):  Location: http://172.16.85.176:3456/cgi-bin/ibm.tcl/info?france50%2Ejpg Image france50.jpg

Element	'529 Claim Recitation	Exemplary Citations in Reference
		Flickner Depo at 63: 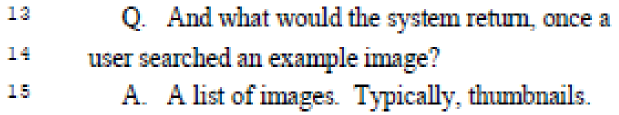 Q. And what would the system return, once a user searched an example image? A. A list of images. Typically, thumbnails. Flickner 8/29/23 Depo at 68–72:

		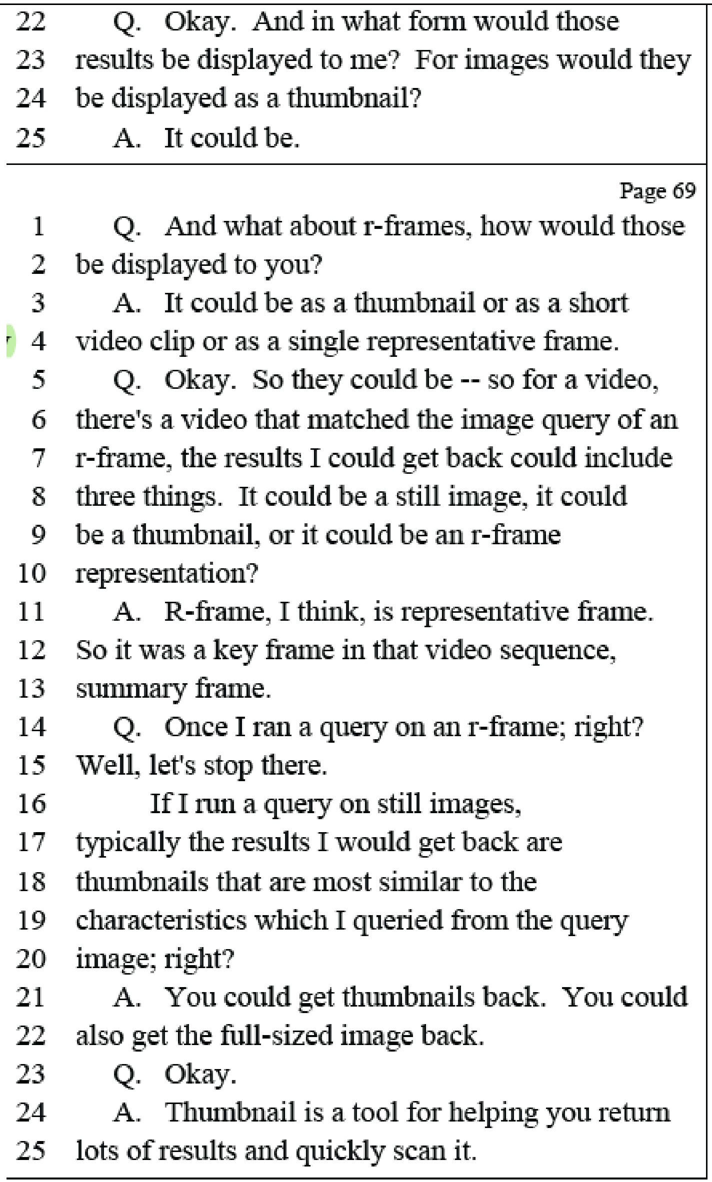 Q. Okay. And in what form would those results be displayed to me? For images would they be displayed as a thumbnail? A. It could be. Page 69 Q. And what about r-frames, how would those be displayed to you? A. It could be as a thumbnail or as a short video clip or as a single representative frame. Q. Okay. So they could be -- so for a video, there's a video that matched the image query of an r-frame, the results I could get back could include three things. It could be a still image, it could be a thumbnail, or it could be an r-frame representation? A. R-frame, I think, is representative frame. So it was a key frame in that video sequence, summary frame. Q. Once I ran a query on an r-frame; right? Well, let's stop there. If I run a query on still images, typically the results I would get back are thumbnails that are most similar to the characteristics which I queried from the query image; right? A. You could get thumbnails back. You could also get the full-sized image back. Q. Okay. A. Thumbnail is a tool for helping you return lots of results and quickly scan it.	

Element	’529 Claim Recitation	Exemplary Citations in Reference
		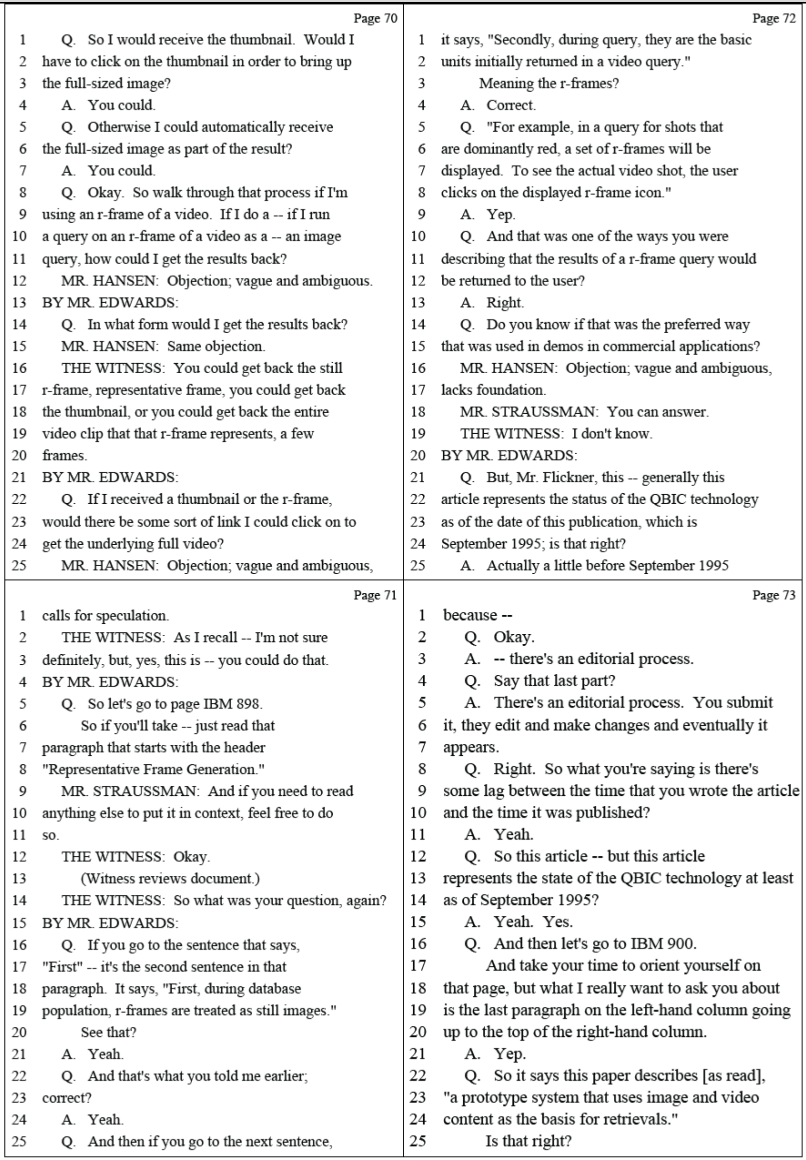

Page 70

Q. So I would receive the thumbnail. Would I have to click on the thumbnail in order to bring up the full-sized image?
A. You could.
Q. Otherwise I could automatically receive the full-sized image as part of the result?
A. You could.
Q. Okay. So walk through that process if I'm using an r-frame of a video. If I do a -- if I run a query on an r-frame of a video as a -- an image query, how could I get the results back?
MR. HANSEN: Objection; vague and ambiguous.
BY MR. EDWARDS:
Q. In what form would I get the results back?
MR. HANSEN: Same objection.
THE WITNESS: You could get back the still r-frame, representative frame, you could get back the thumbnail, or you could get back the entire video clip that that r-frame represents, a few frames.
BY MR. EDWARDS:
Q. If I received a thumbnail or the r-frame, would there be some sort of link I could click on to get the underlying full video?
MR. HANSEN: Objection; vague and ambiguous,

Page 71

calls for speculation.
THE WITNESS: As I recall -- I'm not sure definitely, but, yes, this is -- you could do that.
BY MR. EDWARDS:
Q. So let's go to page IBM 898.
So if you'll take -- just read that paragraph that starts with the header "Representative Frame Generation."
MR. STRAUSSMAN: And if you need to read anything else to put it in context, feel free to do so.
THE WITNESS: Okay.
(Witness reviews document.)
THE WITNESS: So what was your question, again?
BY MR. EDWARDS:
Q. If you go to the sentence that says, "First" -- it's the second sentence in that paragraph. It says, "First, during database population, r-frames are treated as still images."
See that?
A. Yeah.
Q. And that's what you told me earlier; correct?
A. Yeah.
Q. And then if you go to the next sentence,

Page 72

it says, "Secondly, during query, they are the basic units initially returned in a video query."
Meaning the r-frames?
A. Correct.
Q. "For example, in a query for shots that are dominantly red, a set of r-frames will be displayed. To see the actual video shot, the user clicks on the displayed r-frame icon."
A. Yep.
Q. And that was one of the ways you were describing that the results of a r-frame query would be returned to the user?
A. Right.
Q. Do you know if that was the preferred way that was used in demos in commercial applications?
MR. HANSEN: Objection; vague and ambiguous, lacks foundation.
MR. STRAUSSMAN: You can answer.
THE WITNESS: I don't know.
BY MR. EDWARDS:
Q. But, Mr. Flickner, this -- generally this article represents the status of the QBIC technology as of the date of this publication, which is September 1995; is that right?
A. Actually a little before September 1995

Page 73

because --
Q. Okay.
A. -- there's an editorial process.
Q. Say that last part?
A. There's an editorial process. You submit it, they edit and make changes and eventually it appears.
Q. Right. So what you're saying is there's some lag between the time that you wrote the article and the time it was published?
A. Yeah.
Q. So this article -- but this article represents the state of the QBIC technology at least as of September 1995?
A. Yeah. Yes.
Q. And then let's go to IBM 900.
And take your time to orient yourself on that page, but what I really want to ask you about is the last paragraph on the left-hand column going up to the top of the right-hand column.
A. Yep.
Q. So it says this paper describes [as read], "a prototype system that uses image and video content as the basis for retrievals."
Is that right?

Element	'529 Claim Recitation	Exemplary Citations in Reference
		Id. at 129 – 131: 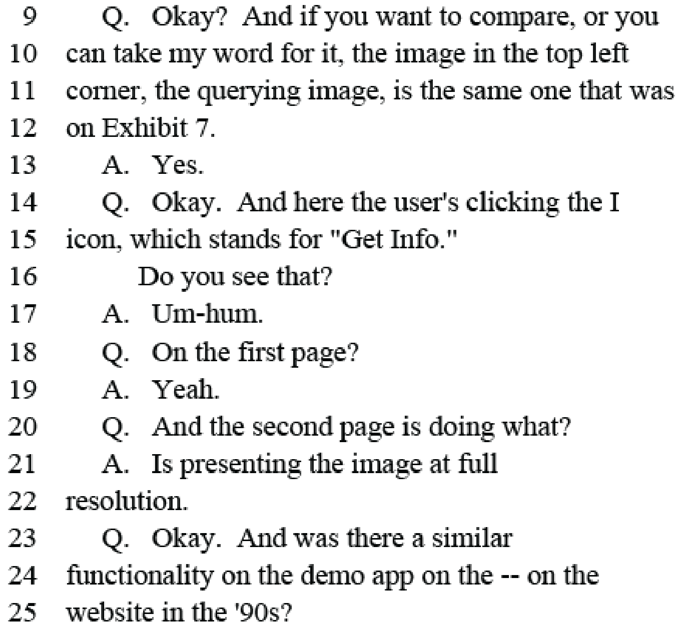 Q. Okay? And if you want to compare, or you can take my word for it, the image in the top left corner, the querying image, is the same one that was on Exhibit 7. A. Yes. Q. Okay. And here the user's clicking the I icon, which stands for "Get Info." Do you see that? A. Um-hum. Q. On the first page? A. Yeah. Q. And the second page is doing what? A. Is presenting the image at full resolution. Q. Okay. And was there a similar functionality on the demo app on the -- on the website in the '90s?

			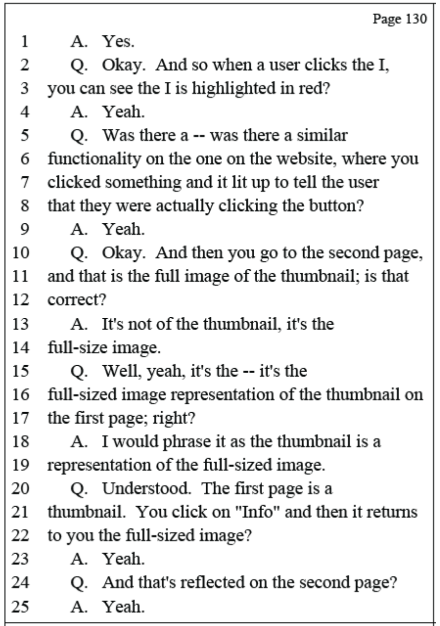 Page 130 A. Yes. Q. Okay. And so when a user clicks the I, you can see the I is highlighted in red? A. Yeah. Q. Was there a -- was there a similar functionality on the one on the website, where you clicked something and it lit up to tell the user that they were actually clicking the button? A. Yeah. Q. Okay. And then you go to the second page, and that is the full image of the thumbnail; is that correct? A. It's not of the thumbnail, it's the full-size image. Q. Well, yeah, it's the -- it's the full-sized image representation of the thumbnail on the first page; right? A. I would phrase it as the thumbnail is a representation of the full-sized image. Q. Understood. The first page is a thumbnail. You click on "Info" and then it returns to you the full-sized image? A. Yeah. Q. And that's reflected on the second page? A. Yeah. 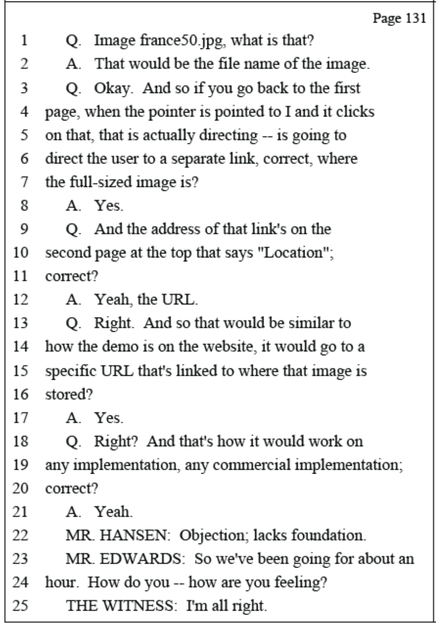 Page 131 Q. Image france50.jpg, what is that? A. That would be the file name of the image. Q. Okay. And so if you go back to the first page, when the pointer is pointed to I and it clicks on that, that is actually directing -- is going to direct the user to a separate link, correct, where the full-sized image is? A. Yes. Q. And the address of that link's on the second page at the top that says "Location"; correct? A. Yeah, the URL. Q. Right. And so that would be similar to how the demo is on the website, it would go to a specific URL that's linked to where that image is stored? A. Yes. Q. Right? And that's how it would work on any implementation, any commercial implementation; correct? A. Yeah. MR. HANSEN: Objection; lacks foundation. MR. EDWARDS: So we've been going for about an hour. How do you -- how are you feeling? THE WITNESS: I'm all right.

Element	**'529 Claim Recitation**	**Exemplary Citations in Reference**
		Id. at 161: 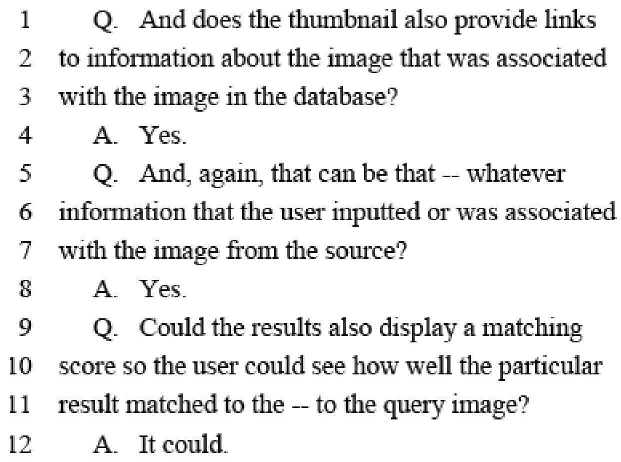 Q. And does the thumbnail also provide links to information about the image that was associated with the image in the database? A. Yes. Q. And, again, that can be that -- whatever information that the user inputted or was associated with the image from the source? A. Yes. Q. Could the results also display a matching score so the user could see how well the particular result matched to the -- to the query image? A. It could. *Id.* at 165:

Element	'529 Claim Recitation	Exemplary Citations in Reference
		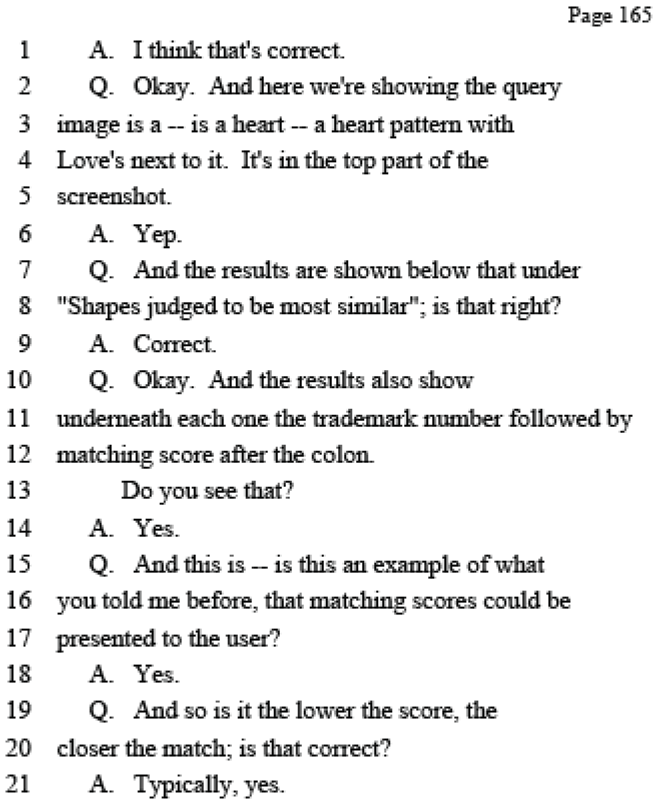 Page 165 A. I think that's correct. Q. Okay. And here we're showing the query image is a -- is a heart -- a heart pattern with Love's next to it. It's in the top part of the screenshot. A. Yep. Q. And the results are shown below that under "Shapes judged to be most similar"; is that right? A. Correct. Q. Okay. And the results also show underneath each one the trademark number followed by matching score after the colon. Do you see that? A. Yes. Q. And this is -- is this an example of what you told me before, that matching scores could be presented to the user? A. Yes. Q. And so is it the lower the score, the closer the match; is that correct? A. Typically, yes. *Id.* at 180-181:

Element	'529 Claim Recitation	Exemplary Citations in Reference
		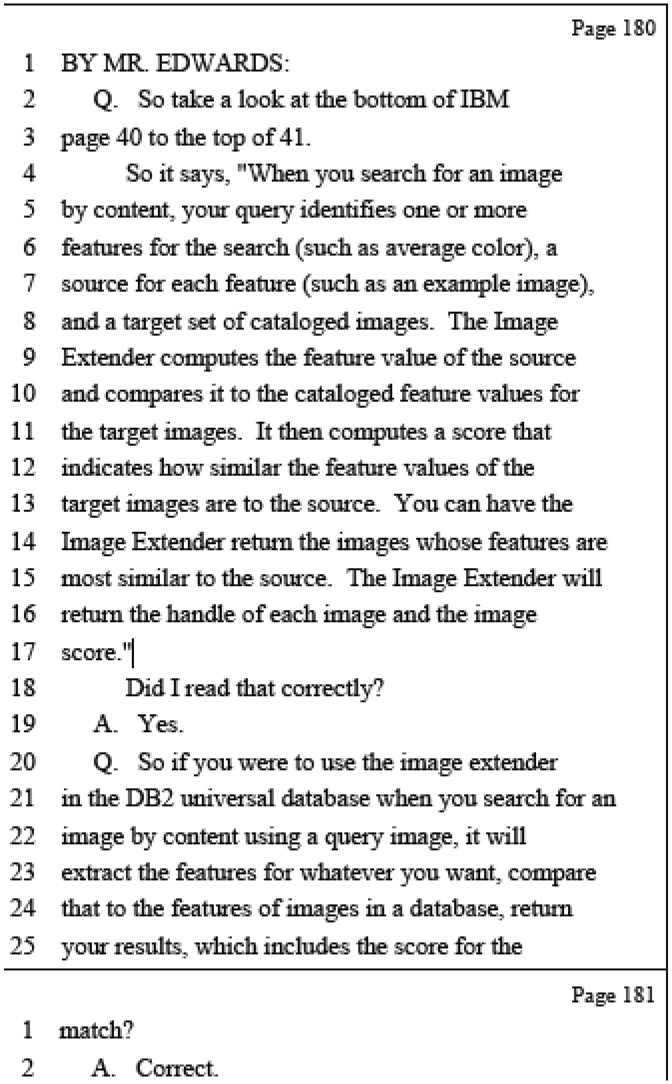 Page 180 BY MR. EDWARDS: Q. So take a look at the bottom of IBM page 40 to the top of 41. So it says, "When you search for an image by content, your query identifies one or more features for the search (such as average color), a source for each feature (such as an example image), and a target set of cataloged images. The Image Extender computes the feature value of the source and compares it to the cataloged feature values for the target images. It then computes a score that indicates how similar the feature values of the target images are to the source. You can have the Image Extender return the images whose features are most similar to the source. The Image Extender will return the handle of each image and the image score." Did I read that correctly? A. Yes. Q. So if you were to use the image extender in the DB2 universal database when you search for an image by content using a query image, it will extract the features for whatever you want, compare that to the features of images in a database, return your results, which includes the score for the Page 181 match? A. Correct.

Element	’529 Claim Recitation	Exemplary Citations in Reference
		Id. at 305: 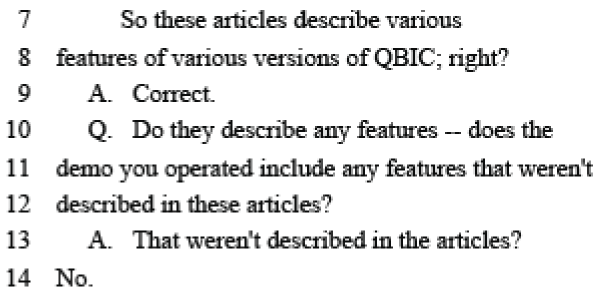 So these articles describe various features of various versions of QBIC; right? A. Correct. Q. Do they describe any features -- does the demo you operated include any features that weren't described in these articles? A. That weren't described in the articles? No. *Id.* at 35: 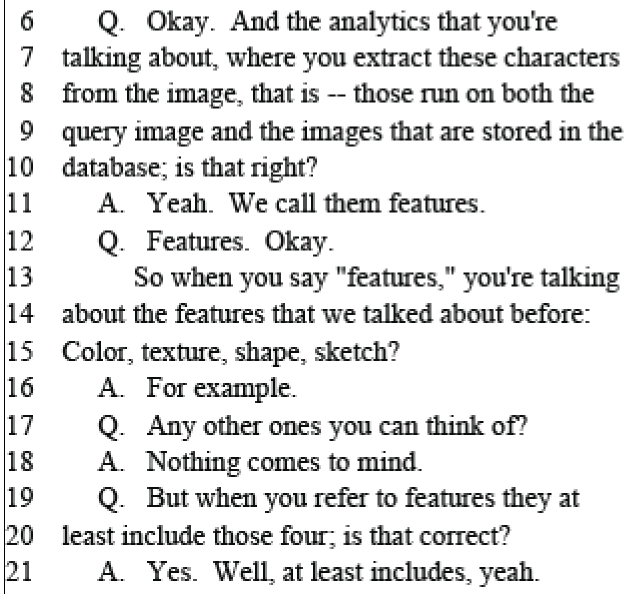 Q. Okay. And the analytics that you're talking about, where you extract these characters from the image, that is -- those run on both the query image and the images that are stored in the database; is that right? A. Yeah. We call them features. Q. Features. Okay. So when you say "features," you're talking about the features that we talked about before: Color, texture, shape, sketch? A. For example. Q. Any other ones you can think of? A. Nothing comes to mind. Q. But when you refer to features they at least include those four; is that correct? A. Yes. Well, at least includes, yeah.

Element	'529 Claim Recitation	Exemplary Citations in Reference
		Id. at 109-110: 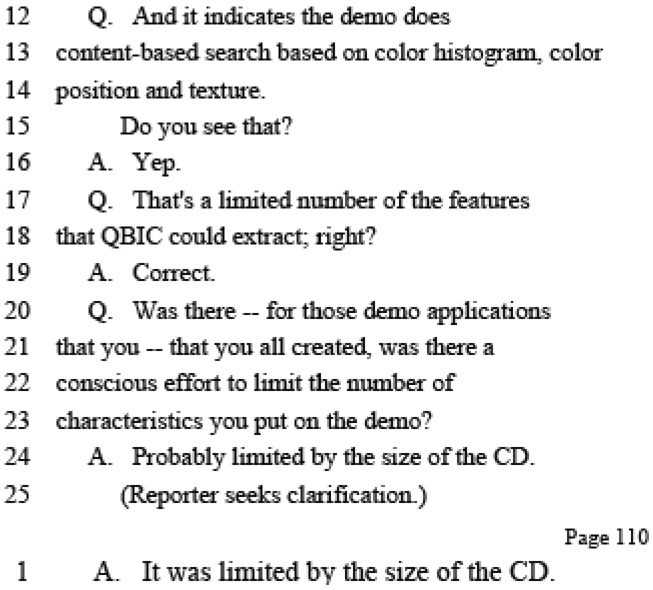 Q. And it indicates the demo does content-based search based on color histogram, color position and texture. Do you see that? A. Yep. Q. That's a limited number of the features that QBIC could extract; right? A. Correct. Q. Was there -- for those demo applications that you -- that you all created, was there a conscious effort to limit the number of characteristics you put on the demo? A. Probably limited by the size of the CD. (Reporter seeks clarification.) Page 110 A. It was limited by the size of the CD. *Id.* at 311: 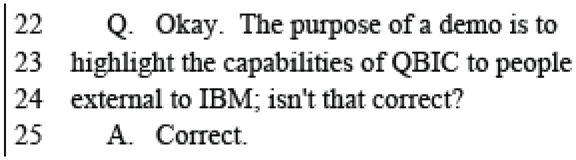 Q. Okay. The purpose of a demo is to highlight the capabilities of QBIC to people external to IBM; isn't that correct? A. Correct. *Id.* at 313-317:

Element	'529 Claim Recitation	Exemplary Citations in Reference
		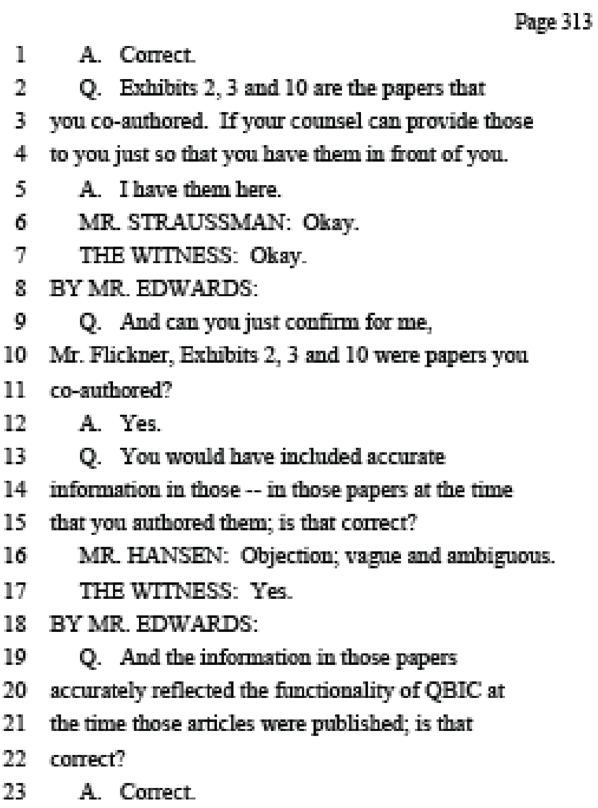 Page 313 A. Correct. Q. Exhibits 2, 3 and 10 are the papers that you co-authored. If your counsel can provide those to you just so that you have them in front of you. A. I have them here. MR. STRAUSSMAN: Okay. THE WITNESS: Okay. BY MR. EDWARDS: Q. And can you just confirm for me, Mr. Flickner, Exhibits 2, 3 and 10 were papers you co-authored? A. Yes. Q. You would have included accurate information in those -- in those papers at the time that you authored them; is that correct? MR. HANSEN: Objection; vague and ambiguous. THE WITNESS: Yes. BY MR. EDWARDS: Q. And the information in those papers accurately reflected the functionality of QBIC at the time those articles were published; is that correct? A. Correct.

		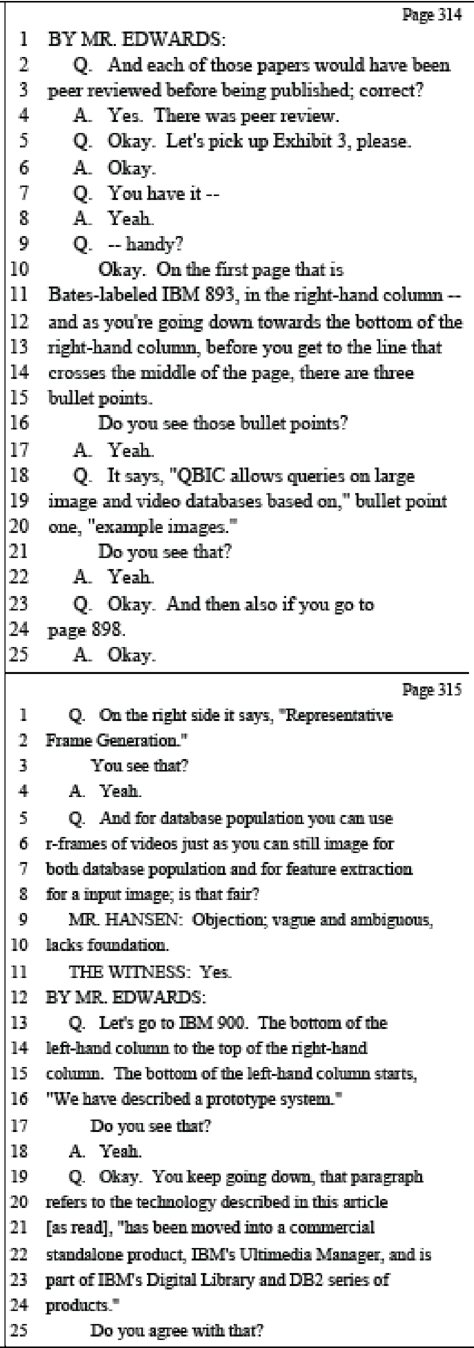 Page 314 BY MR. EDWARDS: Q. And each of those papers would have been peer reviewed before being published; correct? A. Yes. There was peer review. Q. Okay. Let's pick up Exhibit 3, please. A. Okay. Q. You have it -- A. Yeah. Q. -- handy? Okay. On the first page that is Bates-labeled IBM 893, in the right-hand column -- and as you're going down towards the bottom of the right-hand column, before you get to the line that crosses the middle of the page, there are three bullet points. Do you see those bullet points? A. Yeah. Q. It says, "QBIC allows queries on large image and video databases based on," bullet point one, "example images." Do you see that? A. Yeah. Q. Okay. And then also if you go to page 898. A. Okay. Page 315 Q. On the right side it says, "Representative Frame Generation." You see that? A. Yeah. Q. And for database population you can use r-frames of videos just as you can still image for both database population and for feature extraction for a input image; is that fair? MR. HANSEN: Objection; vague and ambiguous, lacks foundation. THE WITNESS: Yes. BY MR. EDWARDS: Q. Let's go to IBM 900. The bottom of the left-hand column to the top of the right-hand column. The bottom of the left-hand column starts, "We have described a prototype system." Do you see that? A. Yeah. Q. Okay. You keep going down, that paragraph refers to the technology described in this article [as read], "has been moved into a commercial standalone product, IBM's Ultimedia Manager, and is part of IBM's Digital Library and DB2 series of products." Do you agree with that?

Element	'529 Claim Recitation	Exemplary Citations in Reference
		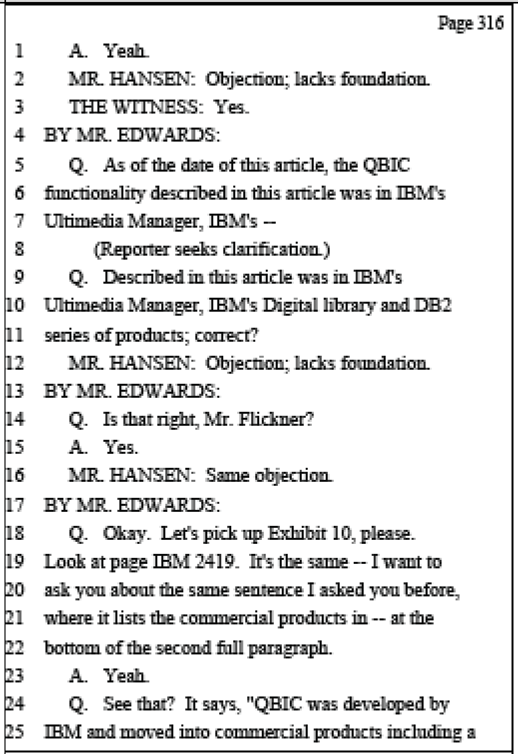 Page 316 A. Yeah. MR. HANSEN: Objection; lacks foundation. THE WITNESS: Yes. BY MR. EDWARDS: Q. As of the date of this article, the QBIC functionality described in this article was in IBM's Ultimedia Manager, IBM's -- (Reporter seeks clarification.) Q. Described in this article was in IBM's Ultimedia Manager, IBM's Digital library and DB2 series of products; correct? MR. HANSEN: Objection; lacks foundation. BY MR. EDWARDS: Q. Is that right, Mr. Flickner? A. Yes. MR. HANSEN: Same objection. BY MR. EDWARDS: Q. Okay. Let's pick up Exhibit 10, please. Look at page IBM 2419. It's the same -- I want to ask you about the same sentence I asked you before, where it lists the commercial products in -- at the bottom of the second full paragraph. A. Yeah. Q. See that? It says, "QBIC was developed by IBM and moved into commercial products including a 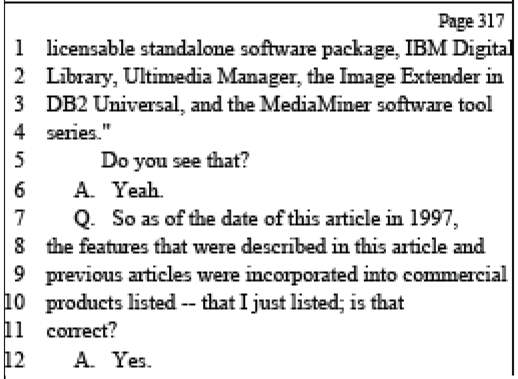 Page 317 licensable standalone software package, IBM Digital Library, Ultimedia Manager, the Image Extender in DB2 Universal, and the MediaMiner software tool series." Do you see that? A. Yeah. Q. So as of the date of this article in 1997, the features that were described in this article and previous articles were incorporated into commercial products listed -- that I just listed; is that correct? A. Yes.

Element	'529 Claim Recitation	Exemplary Citations in Reference
		Id. at 321: 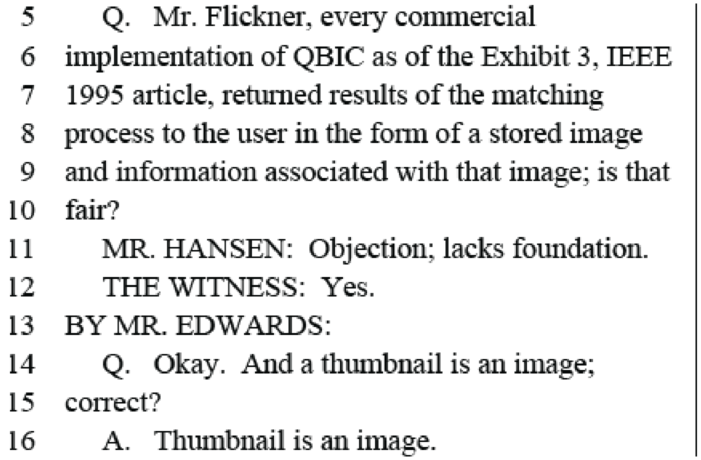 Q. Mr. Flickner, every commercial implementation of QBIC as of the Exhibit 3, IEEE 1995 article, returned results of the matching process to the user in the form of a stored image and information associated with that image; is that fair? MR. HANSEN: Objection; lacks foundation. THE WITNESS: Yes. BY MR. EDWARDS: Q. Okay. And a thumbnail is an image; correct? A. Thumbnail is an image. Flickner 8/29/23 Depo Ex. 7 at 1:

Element	**'529 Claim Recitation**	**Exemplary Citations in Reference**
		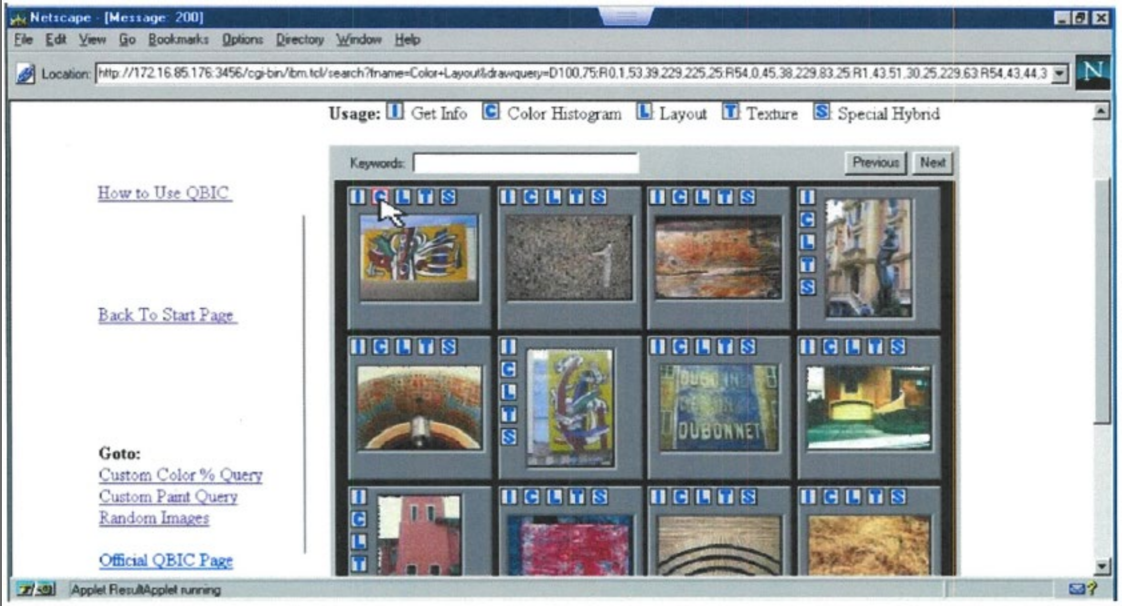 Location: http://172.16.85.176:3456/cgi-bin/ibm.tcl/search?fname=Color+Layout&drawquery=D100,75:R0,1,53,39,229,225,25:R54,0,45,38,229,83,25:R1,43,51,30,25,229,63:R54,43,44,3 *Id.* at 2:

Element	'529 Claim Recitation	Exemplary Citations in Reference
		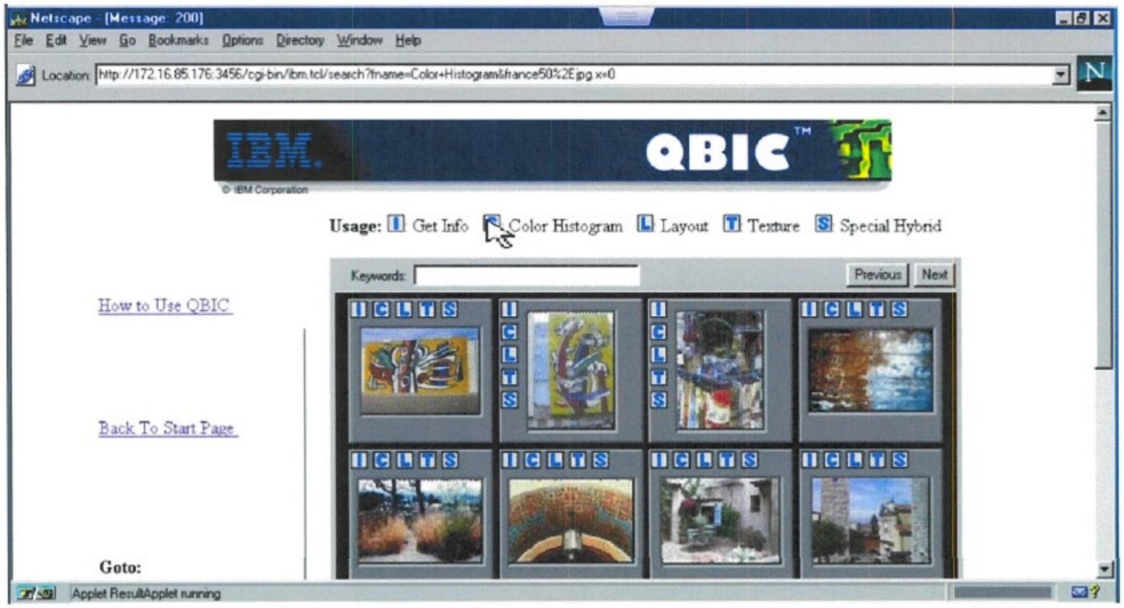 Flickner 8/29/23 Depo Ex. 8 at 1:

Element	'529 Claim Recitation	Exemplary Citations in Reference
		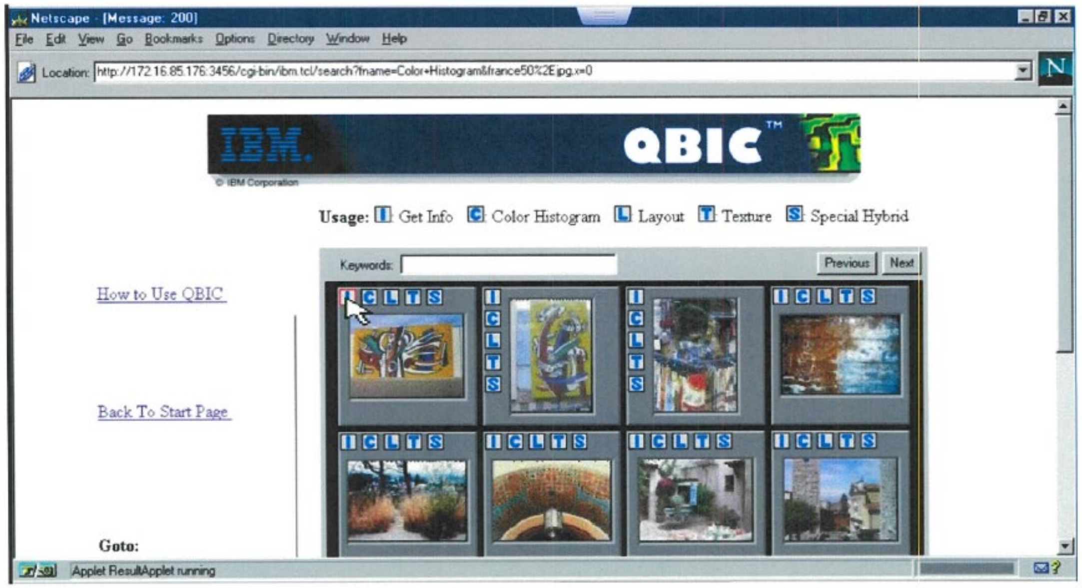 http://172.16.85.176:3456/cgi-bin/ibm.tcl/search?fname=Color+Histogram&france50%2Ejpg.x=0 *Id*. at 2:

Element	'529 Claim Recitation	Exemplary Citations in Reference
		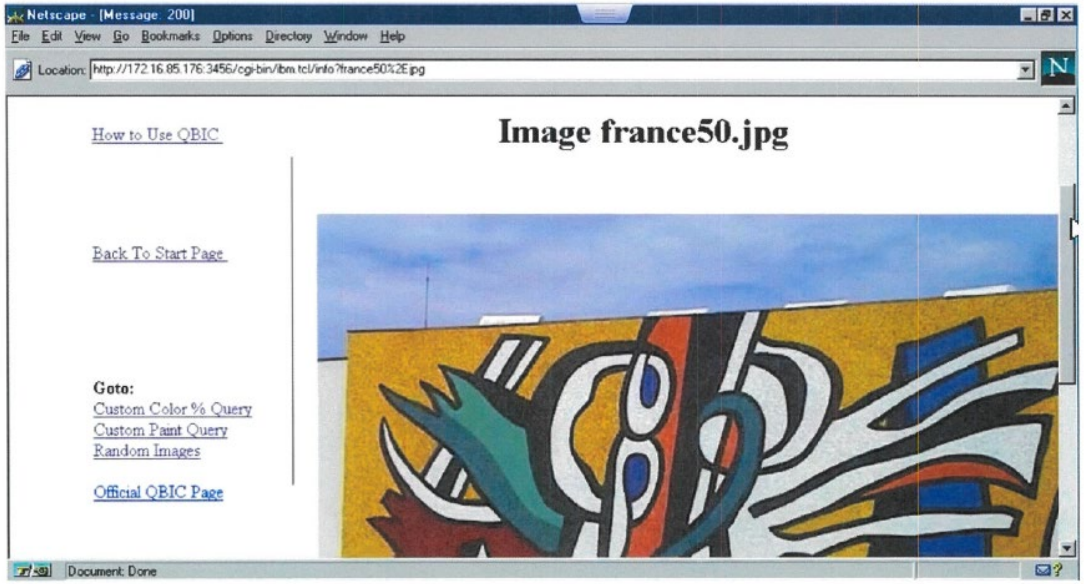 Netscape - [Message: 200] Location: http://172.16.85.176:3456/cgi-bin/ibm.tcl/info?france50%2Ejpg Image france50.jpg How to Use QBIC Back To Start Page Goto: Custom Color % Query Custom Paint Query Random Images Official QBIC Page File Edit View Go Bookmarks Options Directory Window Help Document: Done To the extent it is found that QBIC does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that QBIC does not anticipate this claim, QBIC renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine QBIC with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

2. Claim 4

Claim	'529 Claim Recitation	Exemplary Citations in Reference
4	The system of claim 1, wherein the service is further programmed to identify the object by applying multiple algorithms to at least a portion of the image or of the data, and identify the object as a function of a confidence levels associated with results of the multiple algorithms.	The QBIC Project: Querying Images by Content Using Color, Texture, and Shape (IBM 0002390–404) at 3 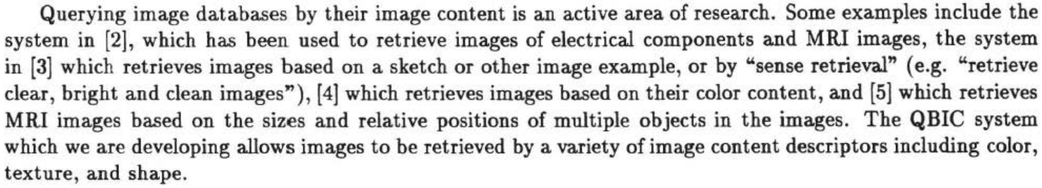 Querying image databases by their image content is an active area of research. Some examples include the system in [2], which has been used to retrieve images of electrical components and MRI images, the system in [3] which retrieves images based on a sketch or other image example, or by "sense retrieval" (e.g. "retrieve clear, bright and clean images"), [4] which retrieves images based on their color content, and [5] which retrieves MRI images based on the sizes and relative positions of multiple objects in the images. The QBIC system which we are developing allows images to be retrieved by a variety of image content descriptors including color, texture, and shape. The QBIC Project: Querying Images by Content Using Color, Texture, and Shape (IBM 0002390–404) at 2391-92. 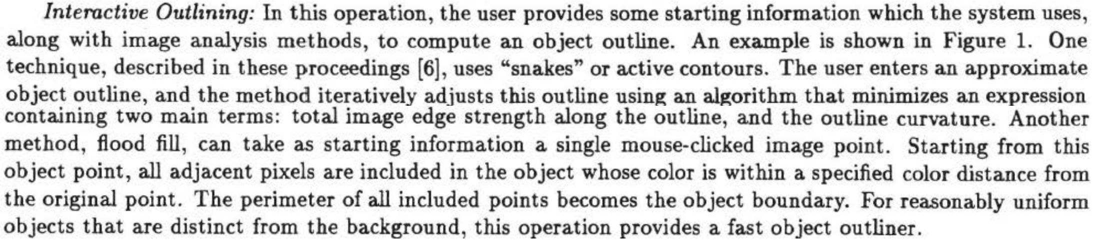 *Interactive Outlining:* In this operation, the user provides some starting information which the system uses, along with image analysis methods, to compute an object outline. An example is shown in Figure 1. One technique, described in these proceedings [6], uses "snakes" or active contours. The user enters an approximate object outline, and the method iteratively adjusts this outline using an algorithm that minimizes an expression containing two main terms: total image edge strength along the outline, and the outline curvature. Another method, flood fill, can take as starting information a single mouse-clicked image point. Starting from this object point, all adjacent pixels are included in the object whose color is within a specified color distance from the original point. The perimeter of all included points becomes the object boundary. For reasonably uniform objects that are distinct from the background, this operation provides a fast object outliner. The QBIC Project: Querying Images by Content Using Color, Texture, and Shape (IBM 0002390–404) at 2392-93 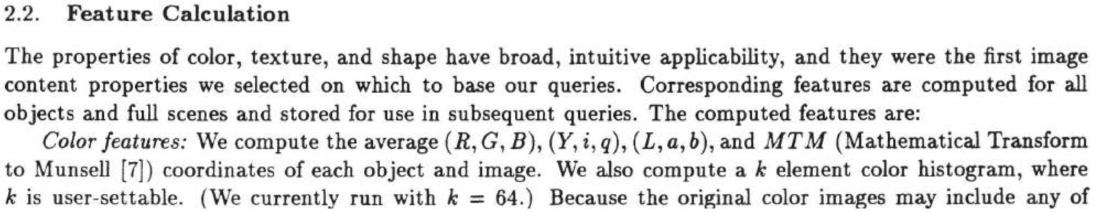 2.2. **Feature Calculation** The properties of color, texture, and shape have broad, intuitive applicability, and they were the first image content properties we selected on which to base our queries. Corresponding features are computed for all objects and full scenes and stored for use in subsequent queries. The computed features are: *Color features:* We compute the average (R, G, B), (Y, i, q), (L, a, b), and MTM (Mathematical Transform to Munsell [7]) coordinates of each object and image. We also compute a k element color histogram, where k is user-settable. (We currently run with $k = 64$.) Because the original color images may include any of

		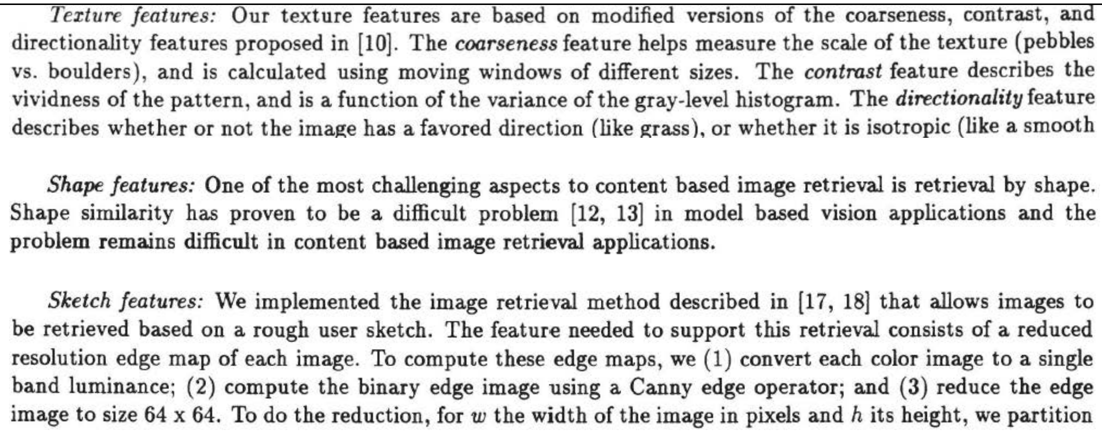 *Texture features:* Our texture features are based on modified versions of the coarseness, contrast, and directionality features proposed in [10]. The *coarseness* feature helps measure the scale of the texture (pebbles vs. boulders), and is calculated using moving windows of different sizes. The *contrast* feature describes the vividness of the pattern, and is a function of the variance of the gray-level histogram. The *directionality* feature describes whether or not the image has a favored direction (like grass), or whether it is isotropic (like a smooth *Shape features:* One of the most challenging aspects to content based image retrieval is retrieval by shape. Shape similarity has proven to be a difficult problem [12, 13] in model based vision applications and the problem remains difficult in content based image retrieval applications. *Sketch features:* We implemented the image retrieval method described in [17, 18] that allows images to be retrieved based on a rough user sketch. The feature needed to support this retrieval consists of a reduced resolution edge map of each image. To compute these edge maps, we (1) convert each color image to a single band luminance; (2) compute the binary edge image using a Canny edge operator; and (3) reduce the edge image to size 64 x 64. To do the reduction, for w the width of the image in pixels and h its height, we partition The QBIC Project: Querying Images by Content Using Color, Texture, and Shape (IBM 0002390–404) at 5–6 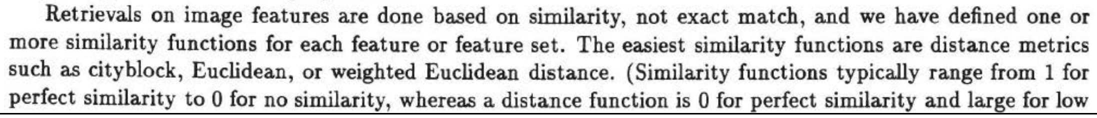 Retrievals on image features are done based on similarity, not exact match, and we have defined one or more similarity functions for each feature or feature set. The easiest similarity functions are distance metrics such as cityblock, Euclidean, or weighted Euclidean distance. (Similarity functions typically range from 1 for perfect similarity to 0 for no similarity, whereas a distance function is 0 for perfect similarity and large for low

		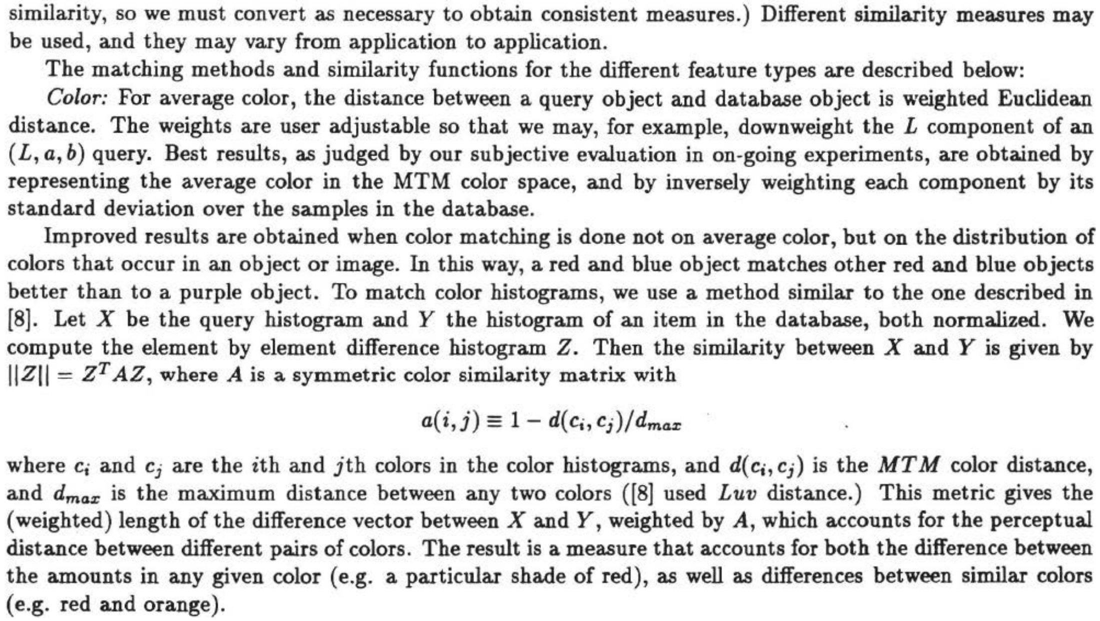 similarity, so we must convert as necessary to obtain consistent measures.) Different similarity measures may be used, and they may vary from application to application. The matching methods and similarity functions for the different feature types are described below: *Color:* For average color, the distance between a query object and database object is weighted Euclidean distance. The weights are user adjustable so that we may, for example, downweight the L component of an (L, a, b) query. Best results, as judged by our subjective evaluation in on-going experiments, are obtained by representing the average color in the MTM color space, and by inversely weighting each component by its standard deviation over the samples in the database. Improved results are obtained when color matching is done not on average color, but on the distribution of colors that occur in an object or image. In this way, a red and blue object matches other red and blue objects better than to a purple object. To match color histograms, we use a method similar to the one described in [8]. Let X be the query histogram and Y the histogram of an item in the database, both normalized. We compute the element by element difference histogram Z. Then the similarity between X and Y is given by $\|Z\| = Z^T A Z$, where A is a symmetric color similarity matrix with $a(i,j) \equiv 1 - d(c_i, c_j)/d_{max}$ where c_i and c_j are the ith and jth colors in the color histograms, and $d(c_i, c_j)$ is the MTM color distance, and d_{max} is the maximum distance between any two colors ([8] used Luv distance.) This metric gives the (weighted) length of the difference vector between X and Y, weighted by A, which accounts for the perceptual distance between different pairs of colors. The result is a measure that accounts for both the difference between the amounts in any given color (e.g. a particular shade of red), as well as differences between similar colors (e.g. red and orange). Query by Image and Video Content: The QBIC System (IBM 0000893–902): "QBIC has two main components: database population (the process of creating an image database) and database query. During the population, images and videos are processed to extract features describing their content-colors, textures, shapes, and camera and object motion-and the features are stored in a database. During the query, the user composes a query graphically. Features are generated from the graphical query and then input to a matching engine that finds images or videos from the database with similar features. Figure 2 shows the system architecture." *Id*. at p. 3.

		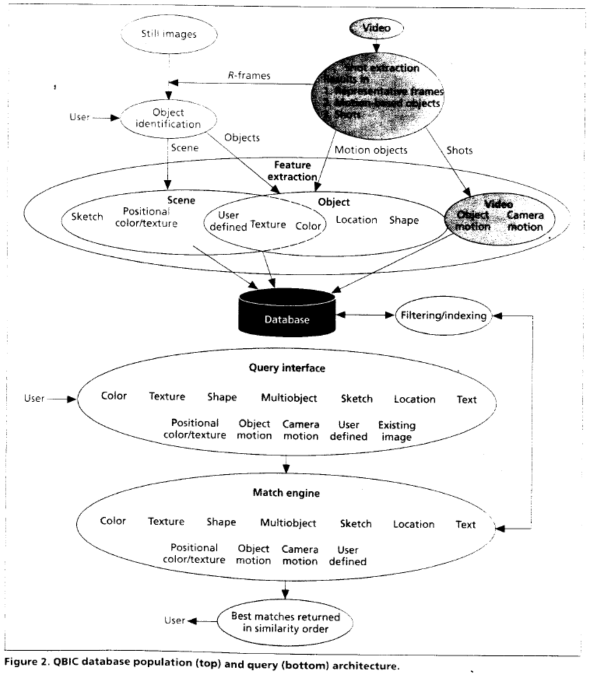 **Figure 2. QBIC database population (top) and query (bottom) architecture.** *Id*. at Fig. 2. "Ideally, object identification would be automatic, but this is generally difficult. The alternative-manual identification-is tedious and can inhibit query-by-content applications. As a result, we have devoted considerable effort to developing tools to aid in this step. In recent work, we have successfully used fully automatic unsupervised segmentation methods along with a foreground/background model to identify objects in a restricted class of images. The images, typical of museums i and retail catalogs, have a small number of foreground objects on a generally separable background. Figure 4 shows example results. Even in this domain, robust algorithms are required because of the textured and variegated backgrounds." *Id*. at pp. 3–4. "ANNOTATION AND DATABASE POPULATION TOOLS. Automatic methods (such as our Positional Color query) that don't rely on object identification,

		methods that identify objects automatically as in the museum image example, fast and easy-to-use semiautomatic outlining tools, and motion-based segmentation algorithms will enable additional application areas." *Id.* at p. 8. "INTEGRATION WITH TEXT AND PARAMETRIC ANNOTATION. Query by visual content complements and extends existing query methods. Systems must be able to integrate queries combining date, subject matter, price, and availability with content properties such as color, texture, and shape." *Id.* at p. 8. Image, Audio, and Video Extenders Administration and Programming (IBM 000001–617): "To make an image available for searching by content, you catalog the image. When you catalog an image, the Image Extender analyzes the image, by computing the feature values for the image, and stores the values in a QBIC catalog. When you search for an image by content, your query identifies one or more features for the search (such as average color), a source for each feature (such as an example image), and a target set of cataloged images. The Image Extender computes the feature value of the source and compares it to the cataloged feature values for the target images. It then computes a score that indicates how similar the feature values of the target images are to the source. You can have the Image Extender return the images whose features are most similar to the source. The Image Extender will return the handle of each image and the image score. You can also have the Image Extender return only the score of a single image." *Id.* at pp. 20–21.

		QBIC Demo (IBM000747): 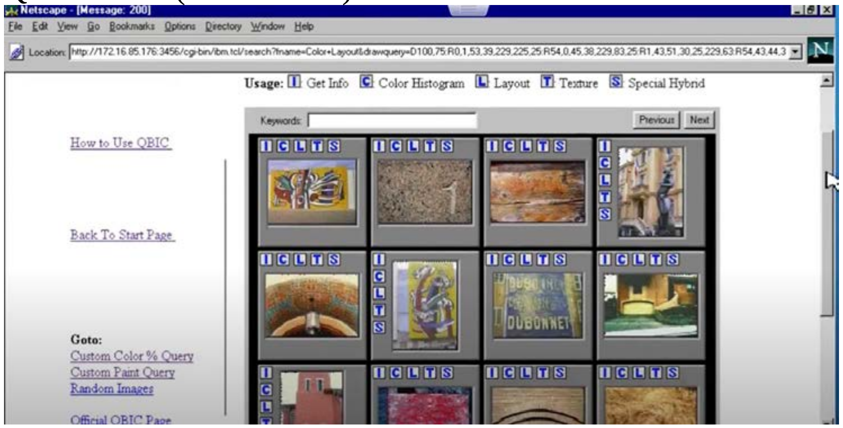 Location: http://172.16.85.176:3456/cgi-bin/ibm.tcl/search?fname=Color+Layout&drawquery=D100,75:R0,1,53,39,229,225,25:R54,0,45,38,229,83,25:R1,43,51,30,25,229,63:R54,43,44,3 QBIC Demo (IBM000747):  Location: http://172.16.85.176:3456/cgi-bin/ibm.tcl/search?fname=Color+Layout&france77%2Ejpg.x=0

Flickner Depo at 28-29:

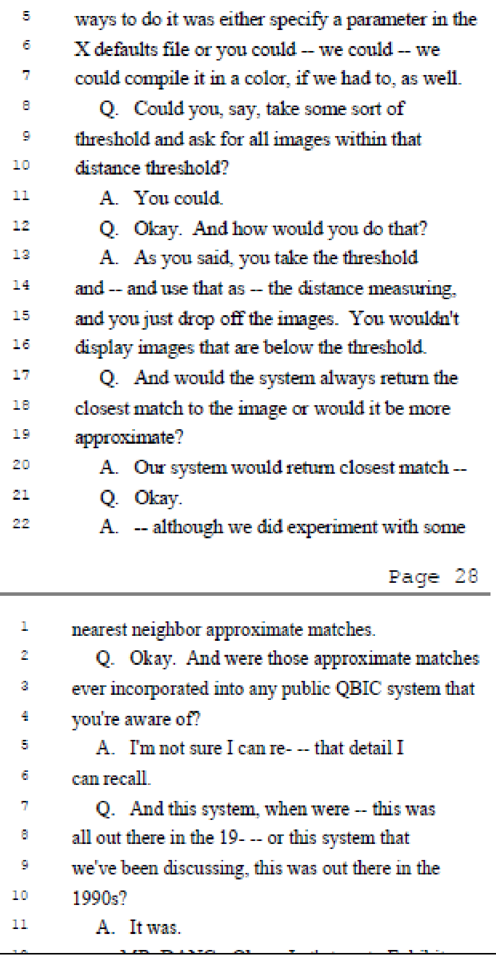

Q. Okay. And how would you alter it then?
A. I don't remember exactly how we -- we ended up doing it in the actual system. But the ways to do it was either specify a parameter in the X defaults file or you could -- we could -- we could compile it in a color, if we had to, as well.
Q. Could you, say, take some sort of threshold and ask for all images within that distance threshold?
A. You could.
Q. Okay. And how would you do that?
A. As you said, you take the threshold and -- and use that as -- the distance measuring, and you just drop off the images. You wouldn't display images that are below the threshold.
Q. And would the system always return the closest match to the image or would it be more approximate?
A. Our system would return closest match --
Q. Okay.
A. -- although we did experiment with some

Page 28

nearest neighbor approximate matches.
Q. Okay. And were those approximate matches ever incorporated into any public QBIC system that you're aware of?
A. I'm not sure I can re- -- that detail I can recall.
Q. And this system, when were -- this was all out there in the 19- -- or this system that we've been discussing, this was out there in the 1990s?
A. It was.

		See also Flickner Depo at 25-26, 41-42, 131-132, 162-163. Flickner 8/29/23 Depo at 85–99: 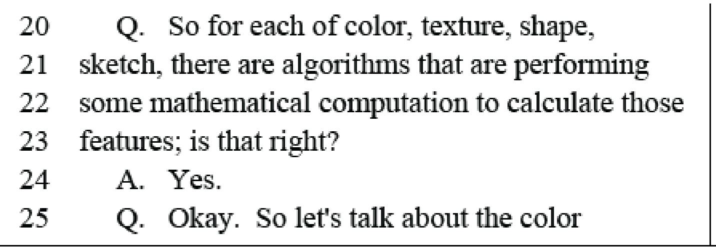 Q. So for each of color, texture, shape, sketch, there are algorithms that are performing some mathematical computation to calculate those features; is that right? A. Yes. Q. Okay. So let's talk about the color

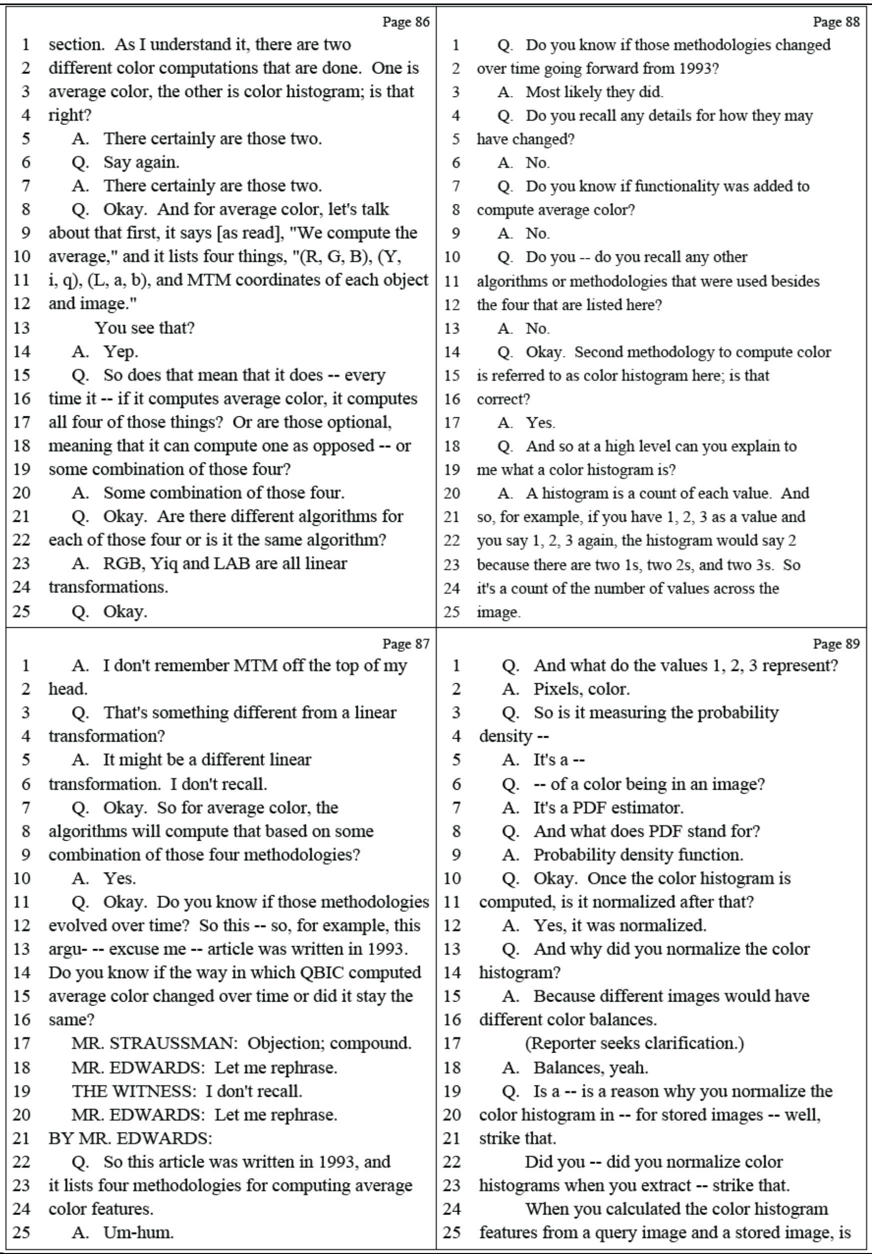

Page 86

section. As I understand it, there are two different color computations that are done. One is average color, the other is color histogram; is that right?
A. There certainly are those two.
Q. Say again.
A. There certainly are those two.
Q. Okay. And for average color, let's talk about that first, it says [as read], "We compute the average," and it lists four things, "(R, G, B), (Y, i, q), (L, a, b), and MTM coordinates of each object and image."
You see that?
A. Yep.
Q. So does that mean that it does -- every time it -- if it computes average color, it computes all four of those things? Or are those optional, meaning that it can compute one as opposed -- or some combination of those four?
A. Some combination of those four.
Q. Okay. Are there different algorithms for each of those four or is it the same algorithm?
A. RGB, Yiq and LAB are all linear transformations.
Q. Okay.

Page 87

A. I don't remember MTM off the top of my head.
Q. That's something different from a linear transformation?
A. It might be a different linear transformation. I don't recall.
Q. Okay. So for average color, the algorithms will compute that based on some combination of those four methodologies?
A. Yes.
Q. Okay. Do you know if those methodologies evolved over time? So this -- so, for example, this argu- -- excuse me -- article was written in 1993. Do you know if the way in which QBIC computed average color changed over time or did it stay the same?
MR. STRAUSSMAN: Objection; compound.
MR. EDWARDS: Let me rephrase.
THE WITNESS: I don't recall.
MR. EDWARDS: Let me rephrase.
BY MR. EDWARDS:
Q. So this article was written in 1993, and it lists four methodologies for computing average color features.
A. Um-hum.

Page 88

Q. Do you know if those methodologies changed over time going forward from 1993?
A. Most likely they did.
Q. Do you recall any details for how they may have changed?
A. No.
Q. Do you know if functionality was added to compute average color?
A. No.
Q. Do you -- do you recall any other algorithms or methodologies that were used besides the four that are listed here?
A. No.
Q. Okay. Second methodology to compute color is referred to as color histogram here; is that correct?
A. Yes.
Q. And so at a high level can you explain to me what a color histogram is?
A. A histogram is a count of each value. And so, for example, if you have 1, 2, 3 as a value and you say 1, 2, 3 again, the histogram would say 2 because there are two 1s, two 2s, and two 3s. So it's a count of the number of values across the image.

Page 89

Q. And what do the values 1, 2, 3 represent?
A. Pixels, color.
Q. So is it measuring the probability density --
A. It's a --
Q. -- of a color being in an image?
A. It's a PDF estimator.
Q. And what does PDF stand for?
A. Probability density function.
Q. Okay. Once the color histogram is computed, is it normalized after that?
A. Yes, it was normalized.
Q. And why did you normalize the color histogram?
A. Because different images would have different color balances.
(Reporter seeks clarification.)
A. Balances, yeah.
Q. Is a -- is a reason why you normalize the color histogram in -- for stored images -- well, strike that.
Did you -- did you normalize color histograms when you extract -- strike that.
When you calculated the color histogram features from a query image and a stored image, is

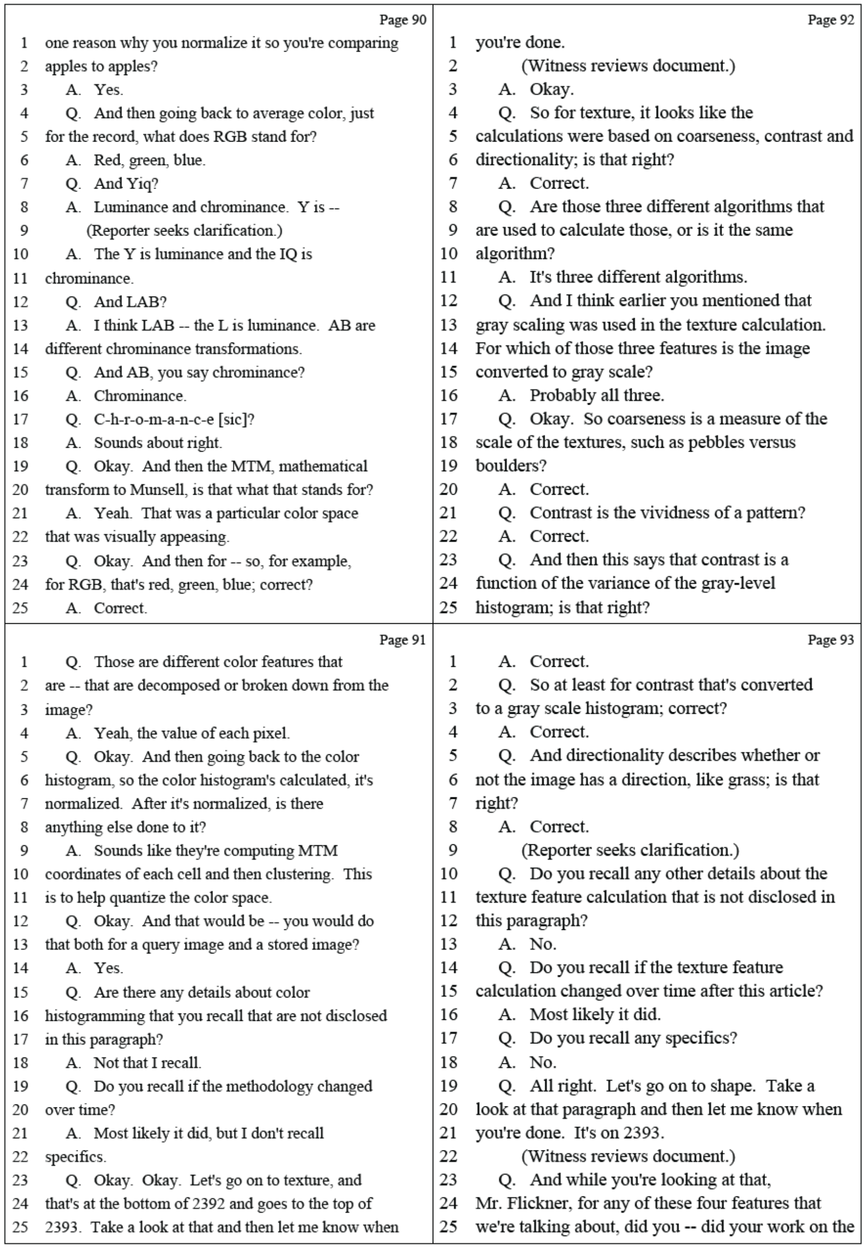

Page 90

one reason why you normalize it so you're comparing apples to apples?
A. Yes.
Q. And then going back to average color, just for the record, what does RGB stand for?
A. Red, green, blue.
Q. And Yiq?
A. Luminance and chrominance. Y is --
(Reporter seeks clarification.)
A. The Y is luminance and the IQ is chrominance.
Q. And LAB?
A. I think LAB -- the L is luminance. AB are different chrominance transformations.
Q. And AB, you say chrominance?
A. Chrominance.
Q. C-h-r-o-m-a-n-c-e [sic]?
A. Sounds about right.
Q. Okay. And then the MTM, mathematical transform to Munsell, is that what that stands for?
A. Yeah. That was a particular color space that was visually appeasing.
Q. Okay. And then for -- so, for example, for RGB, that's red, green, blue; correct?
A. Correct.

Page 91

Q. Those are different color features that are -- that are decomposed or broken down from the image?
A. Yeah, the value of each pixel.
Q. Okay. And then going back to the color histogram, so the color histogram's calculated, it's normalized. After it's normalized, is there anything else done to it?
A. Sounds like they're computing MTM coordinates of each cell and then clustering. This is to help quantize the color space.
Q. Okay. And that would be -- you would do that both for a query image and a stored image?
A. Yes.
Q. Are there any details about color histogramming that you recall that are not disclosed in this paragraph?
A. Not that I recall.
Q. Do you recall if the methodology changed over time?
A. Most likely it did, but I don't recall specifics.
Q. Okay. Okay. Let's go on to texture, and that's at the bottom of 2392 and goes to the top of 2393. Take a look at that and then let me know when

Page 92

you're done.
(Witness reviews document.)
A. Okay.
Q. So for texture, it looks like the calculations were based on coarseness, contrast and directionality; is that right?
A. Correct.
Q. Are those three different algorithms that are used to calculate those, or is it the same algorithm?
A. It's three different algorithms.
Q. And I think earlier you mentioned that gray scaling was used in the texture calculation. For which of those three features is the image converted to gray scale?
A. Probably all three.
Q. Okay. So coarseness is a measure of the scale of the textures, such as pebbles versus boulders?
A. Correct.
Q. Contrast is the vividness of a pattern?
A. Correct.
Q. And then this says that contrast is a function of the variance of the gray-level histogram; is that right?

Page 93

A. Correct.
Q. So at least for contrast that's converted to a gray scale histogram; correct?
A. Correct.
Q. And directionality describes whether or not the image has a direction, like grass; is that right?
A. Correct.
(Reporter seeks clarification.)
Q. Do you recall any other details about the texture feature calculation that is not disclosed in this paragraph?
A. No.
Q. Do you recall if the texture feature calculation changed over time after this article?
A. Most likely it did.
Q. Do you recall any specifics?
A. No.
Q. All right. Let's go on to shape. Take a look at that paragraph and then let me know when you're done. It's on 2393.
(Witness reviews document.)
Q. And while you're looking at that, Mr. Flickner, for any of these four features that we're talking about, did you -- did your work on the

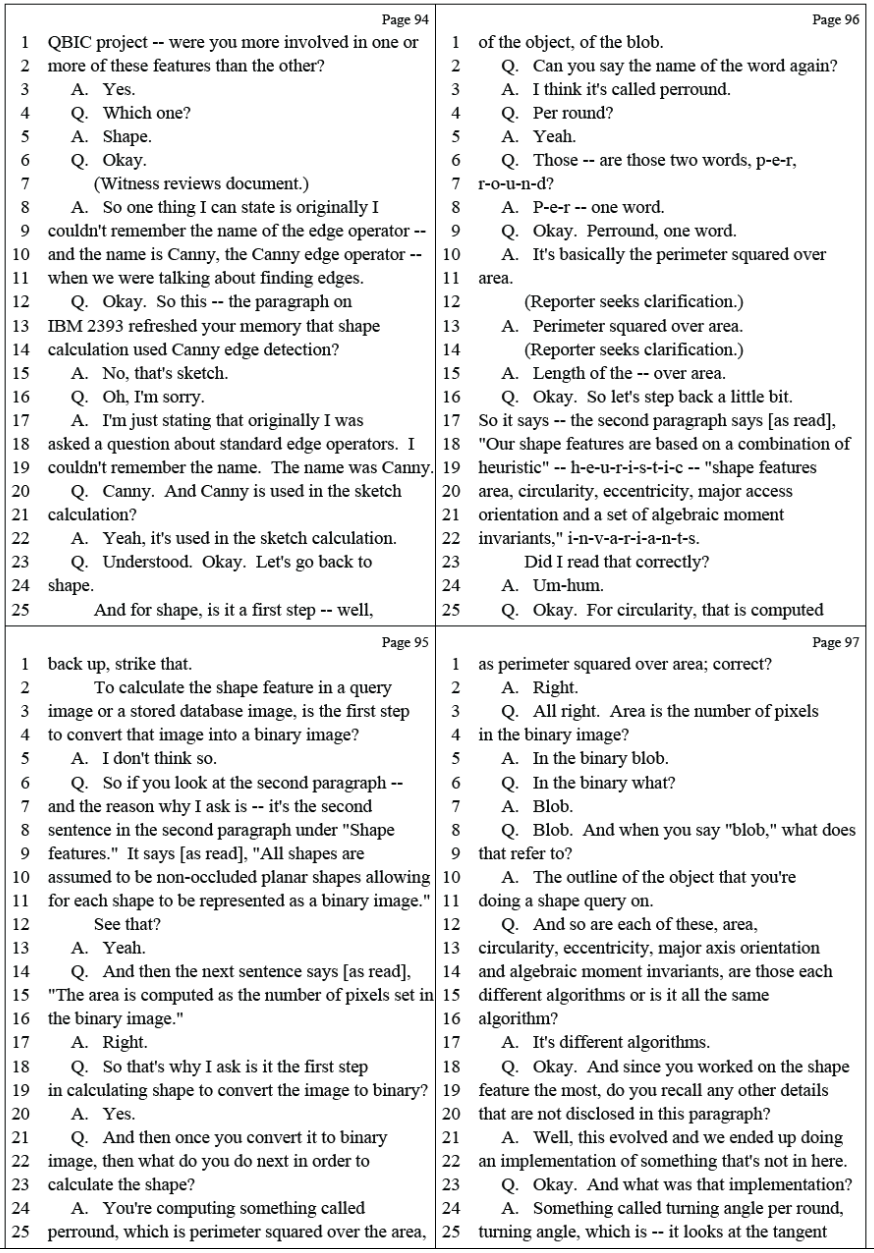

Page 94

QBIC project -- were you more involved in one or more of these features than the other?
A. Yes.
Q. Which one?
A. Shape.
Q. Okay.
(Witness reviews document.)
A. So one thing I can state is originally I couldn't remember the name of the edge operator -- and the name is Canny, the Canny edge operator -- when we were talking about finding edges.
Q. Okay. So this -- the paragraph on IBM 2393 refreshed your memory that shape calculation used Canny edge detection?
A. No, that's sketch.
Q. Oh, I'm sorry.
A. I'm just stating that originally I was asked a question about standard edge operators. I couldn't remember the name. The name was Canny.
Q. Canny. And Canny is used in the sketch calculation?
A. Yeah, it's used in the sketch calculation.
Q. Understood. Okay. Let's go back to shape.
And for shape, is it a first step -- well,

Page 95

back up, strike that.
To calculate the shape feature in a query image or a stored database image, is the first step to convert that image into a binary image?
A. I don't think so.
Q. So if you look at the second paragraph -- and the reason why I ask is -- it's the second sentence in the second paragraph under "Shape features." It says [as read], "All shapes are assumed to be non-occluded planar shapes allowing for each shape to be represented as a binary image."
See that?
A. Yeah.
Q. And then the next sentence says [as read], "The area is computed as the number of pixels set in the binary image."
A. Right.
Q. So that's why I ask is it the first step in calculating shape to convert the image to binary?
A. Yes.
Q. And then once you convert it to binary image, then what do you do next in order to calculate the shape?
A. You're computing something called perround, which is perimeter squared over the area,

Page 96

of the object, of the blob.
Q. Can you say the name of the word again?
A. I think it's called perround.
Q. Per round?
A. Yeah.
Q. Those -- are those two words, p-e-r, r-o-u-n-d?
A. P-e-r -- one word.
Q. Okay. Perround, one word.
A. It's basically the perimeter squared over area.
(Reporter seeks clarification.)
A. Perimeter squared over area.
(Reporter seeks clarification.)
A. Length of the -- over area.
Q. Okay. So let's step back a little bit. So it says -- the second paragraph says [as read], "Our shape features are based on a combination of heuristic" -- h-e-u-r-i-s-t-i-c -- "shape features area, circularity, eccentricity, major access orientation and a set of algebraic moment invariants," i-n-v-a-r-i-a-n-t-s.
Did I read that correctly?
A. Um-hum.
Q. Okay. For circularity, that is computed

Page 97

as perimeter squared over area; correct?
A. Right.
Q. All right. Area is the number of pixels in the binary image?
A. In the binary blob.
Q. In the binary what?
A. Blob.
Q. Blob. And when you say "blob," what does that refer to?
A. The outline of the object that you're doing a shape query on.
Q. And so are each of these, area, circularity, eccentricity, major axis orientation and algebraic moment invariants, are those each different algorithms or is it all the same algorithm?
A. It's different algorithms.
Q. Okay. And since you worked on the shape feature the most, do you recall any other details that are not disclosed in this paragraph?
A. Well, this evolved and we ended up doing an implementation of something that's not in here.
Q. Okay. And what was that implementation?
A. Something called turning angle per round, turning angle, which is -- it looks at the tangent

			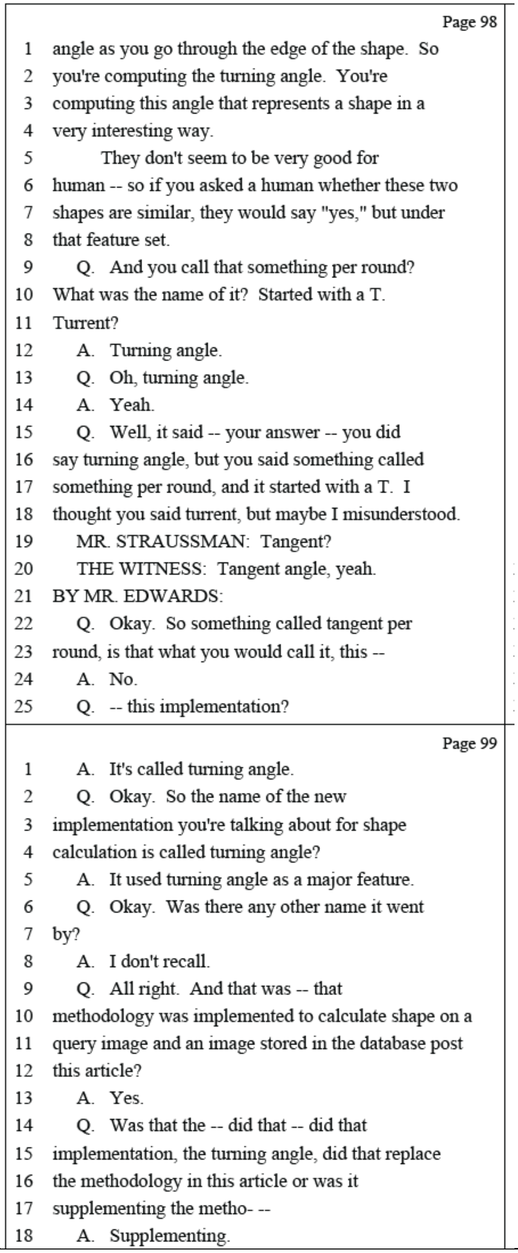 Page 98 angle as you go through the edge of the shape. So you're computing the turning angle. You're computing this angle that represents a shape in a very interesting way. They don't seem to be very good for human -- so if you asked a human whether these two shapes are similar, they would say "yes," but under that feature set. Q. And you call that something per round? What was the name of it? Started with a T. Turrent? A. Turning angle. Q. Oh, turning angle. A. Yeah. Q. Well, it said -- your answer -- you did say turning angle, but you said something called something per round, and it started with a T. I thought you said turrent, but maybe I misunderstood. MR. STRAUSSMAN: Tangent? THE WITNESS: Tangent angle, yeah. BY MR. EDWARDS: Q. Okay. So something called tangent per round, is that what you would call it, this -- A. No. Q. -- this implementation? Page 99 A. It's called turning angle. Q. Okay. So the name of the new implementation you're talking about for shape calculation is called turning angle? A. It used turning angle as a major feature. Q. Okay. Was there any other name it went by? A. I don't recall. Q. All right. And that was -- that methodology was implemented to calculate shape on a query image and an image stored in the database post this article? A. Yes. Q. Was that the -- did that -- did that implementation, the turning angle, did that replace the methodology in this article or was it supplementing the metho- -- A. Supplementing.

		Id. at 123 – 124: 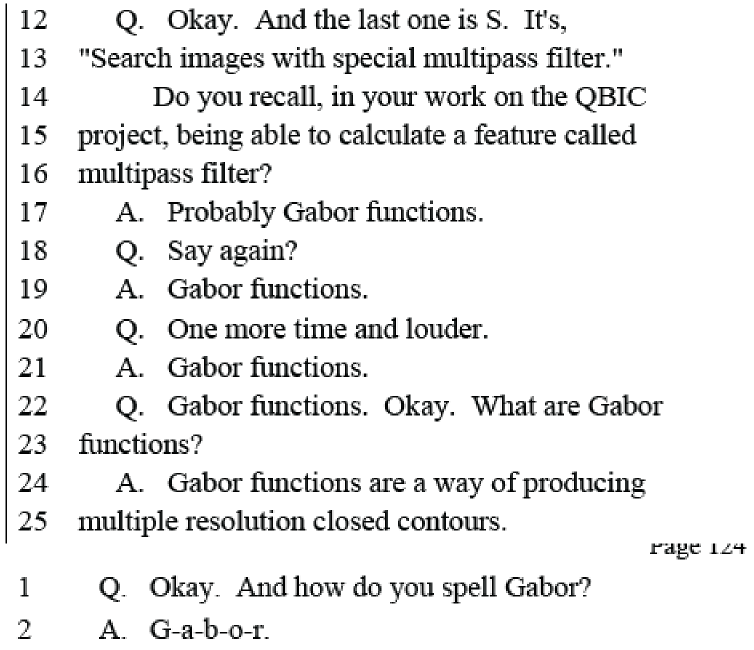 Q. Okay. And the last one is S. It's, "Search images with special multipass filter." Do you recall, in your work on the QBIC project, being able to calculate a feature called multipass filter? A. Probably Gabor functions. Q. Say again? A. Gabor functions. Q. One more time and louder. A. Gabor functions. Q. Gabor functions. Okay. What are Gabor functions? A. Gabor functions are a way of producing multiple resolution closed contours. Page 124 Q. Okay. And how do you spell Gabor? A. G-a-b-o-r. *Id.* at 127 – 128:

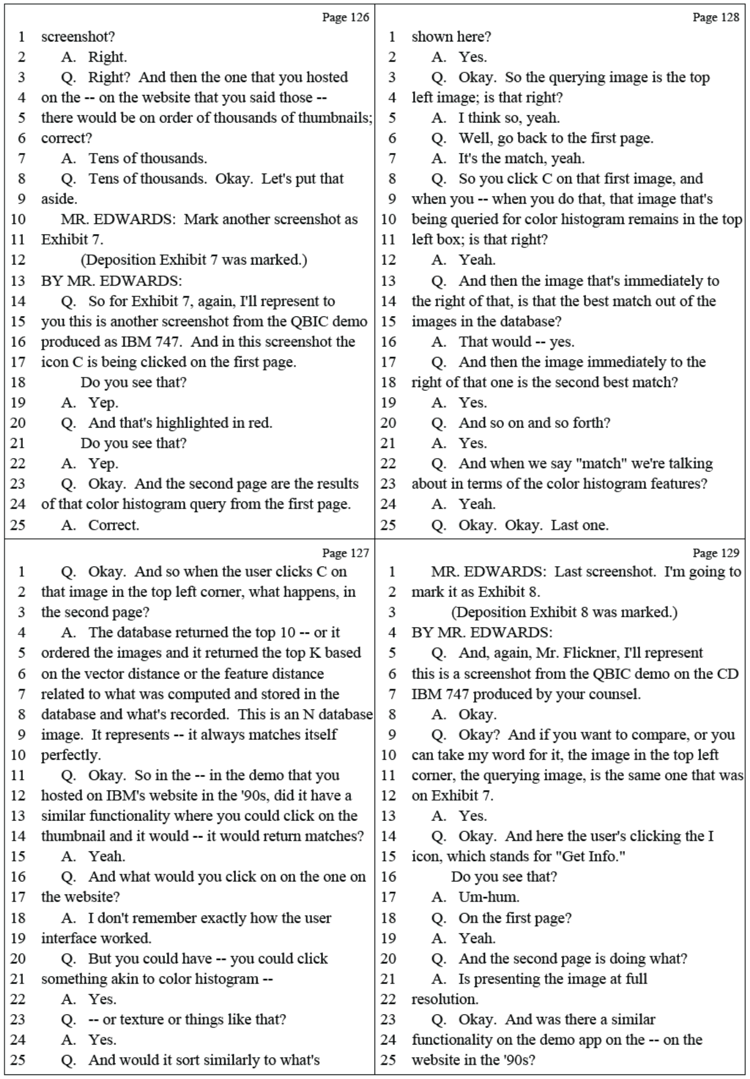

Page 126

screenshot?
A. Right.
Q. Right? And then the one that you hosted on the -- on the website that you said those -- there would be on order of thousands of thumbnails; correct?
A. Tens of thousands.
Q. Tens of thousands. Okay. Let's put that aside.
MR. EDWARDS: Mark another screenshot as Exhibit 7.
(Deposition Exhibit 7 was marked.)
BY MR. EDWARDS:
Q. So for Exhibit 7, again, I'll represent to you this is another screenshot from the QBIC demo produced as IBM 747. And in this screenshot the icon C is being clicked on the first page.
Do you see that?
A. Yep.
Q. And that's highlighted in red.
Do you see that?
A. Yep.
Q. Okay. And the second page are the results of that color histogram query from the first page.
A. Correct.

Page 127

Q. Okay. And so when the user clicks C on that image in the top left corner, what happens, in the second page?
A. The database returned the top 10 -- or it ordered the images and it returned the top K based on the vector distance or the feature distance related to what was computed and stored in the database and what's recorded. This is an N database image. It represents -- it always matches itself perfectly.
Q. Okay. So in the -- in the demo that you hosted on IBM's website in the '90s, did it have a similar functionality where you could click on the thumbnail and it would -- it would return matches?
A. Yeah.
Q. And what would you click on on the one on the website?
A. I don't remember exactly how the user interface worked.
Q. But you could have -- you could click something akin to color histogram --
A. Yes.
Q. -- or texture or things like that?
A. Yes.
Q. And would it sort similarly to what's

Page 128

shown here?
A. Yes.
Q. Okay. So the querying image is the top left image; is that right?
A. I think so, yeah.
Q. Well, go back to the first page.
A. It's the match, yeah.
Q. So you click C on that first image, and when you -- when you do that, that image that's being queried for color histogram remains in the top left box; is that right?
A. Yeah.
Q. And then the image that's immediately to the right of that, is that the best match out of the images in the database?
A. That would -- yes.
Q. And then the image immediately to the right of that one is the second best match?
A. Yes.
Q. And so on and so forth?
A. Yes.
Q. And when we say "match" we're talking about in terms of the color histogram features?
A. Yeah.
Q. Okay. Okay. Last one.

Page 129

MR. EDWARDS: Last screenshot. I'm going to mark it as Exhibit 8.
(Deposition Exhibit 8 was marked.)
BY MR. EDWARDS:
Q. And, again, Mr. Flickner, I'll represent this is a screenshot from the QBIC demo on the CD IBM 747 produced by your counsel.
A. Okay.
Q. Okay? And if you want to compare, or you can take my word for it, the image in the top left corner, the querying image, is the same one that was on Exhibit 7.
A. Yes.
Q. Okay. And here the user's clicking the I icon, which stands for "Get Info."
Do you see that?
A. Um-hum.
Q. On the first page?
A. Yeah.
Q. And the second page is doing what?
A. Is presenting the image at full resolution.
Q. Okay. And was there a similar functionality on the demo app on the -- on the website in the '90s?

Id. at 135 – 136:

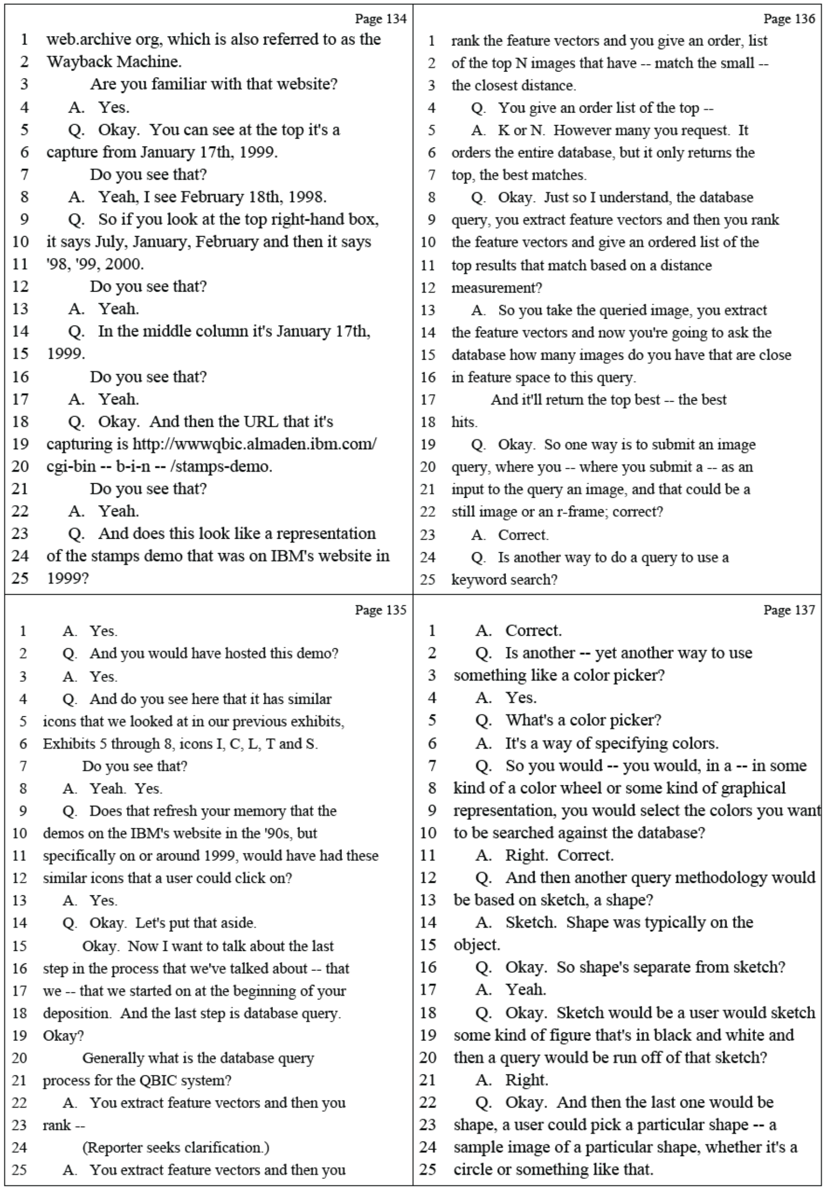

Page 134

web.archive org, which is also referred to as the Wayback Machine.
Are you familiar with that website?
A. Yes.
Q. Okay. You can see at the top it's a capture from January 17th, 1999.
Do you see that?
A. Yeah, I see February 18th, 1998.
Q. So if you look at the top right-hand box, it says July, January, February and then it says '98, '99, 2000.
Do you see that?
A. Yeah.
Q. In the middle column it's January 17th, 1999.
Do you see that?
A. Yeah.
Q. Okay. And then the URL that it's capturing is http://wwwqbic.almaden.ibm.com/ cgi-bin -- b-i-n -- /stamps-demo.
Do you see that?
A. Yeah.
Q. And does this look like a representation of the stamps demo that was on IBM's website in 1999?

Page 135

A. Yes.
Q. And you would have hosted this demo?
A. Yes.
Q. And do you see here that it has similar icons that we looked at in our previous exhibits, Exhibits 5 through 8, icons I, C, L, T and S.
Do you see that?
A. Yeah. Yes.
Q. Does that refresh your memory that the demos on the IBM's website in the '90s, but specifically on or around 1999, would have had these similar icons that a user could click on?
A. Yes.
Q. Okay. Let's put that aside.
Okay. Now I want to talk about the last step in the process that we've talked about -- that we -- that we started on at the beginning of your deposition. And the last step is database query. Okay?
Generally what is the database query process for the QBIC system?
A. You extract feature vectors and then you rank --
(Reporter seeks clarification.)
A. You extract feature vectors and then you

Page 136

rank the feature vectors and you give an order, list of the top N images that have -- match the small -- the closest distance.
Q. You give an order list of the top --
A. K or N. However many you request. It orders the entire database, but it only returns the top, the best matches.
Q. Okay. Just so I understand, the database query, you extract feature vectors and then you rank the feature vectors and give an ordered list of the top results that match based on a distance measurement?
A. So you take the queried image, you extract the feature vectors and now you're going to ask the database how many images do you have that are close in feature space to this query.
And it'll return the top best -- the best hits.
Q. Okay. So one way is to submit an image query, where you -- where you submit a -- as an input to the query an image, and that could be a still image or an r-frame; correct?
A. Correct.
Q. Is another way to do a query to use a keyword search?

Page 137

A. Correct.
Q. Is another -- yet another way to use something like a color picker?
A. Yes.
Q. What's a color picker?
A. It's a way of specifying colors.
Q. So you would -- you would, in a -- in some kind of a color wheel or some kind of graphical representation, you would select the colors you want to be searched against the database?
A. Right. Correct.
Q. And then another query methodology would be based on sketch, a shape?
A. Sketch. Shape was typically on the object.
Q. Okay. So shape's separate from sketch?
A. Yeah.
Q. Okay. Sketch would be a user would sketch some kind of figure that's in black and white and then a query would be run off of that sketch?
A. Right.
Q. Okay. And then the last one would be shape, a user could pick a particular shape -- a sample image of a particular shape, whether it's a circle or something like that.

Id. at 143 – 149:

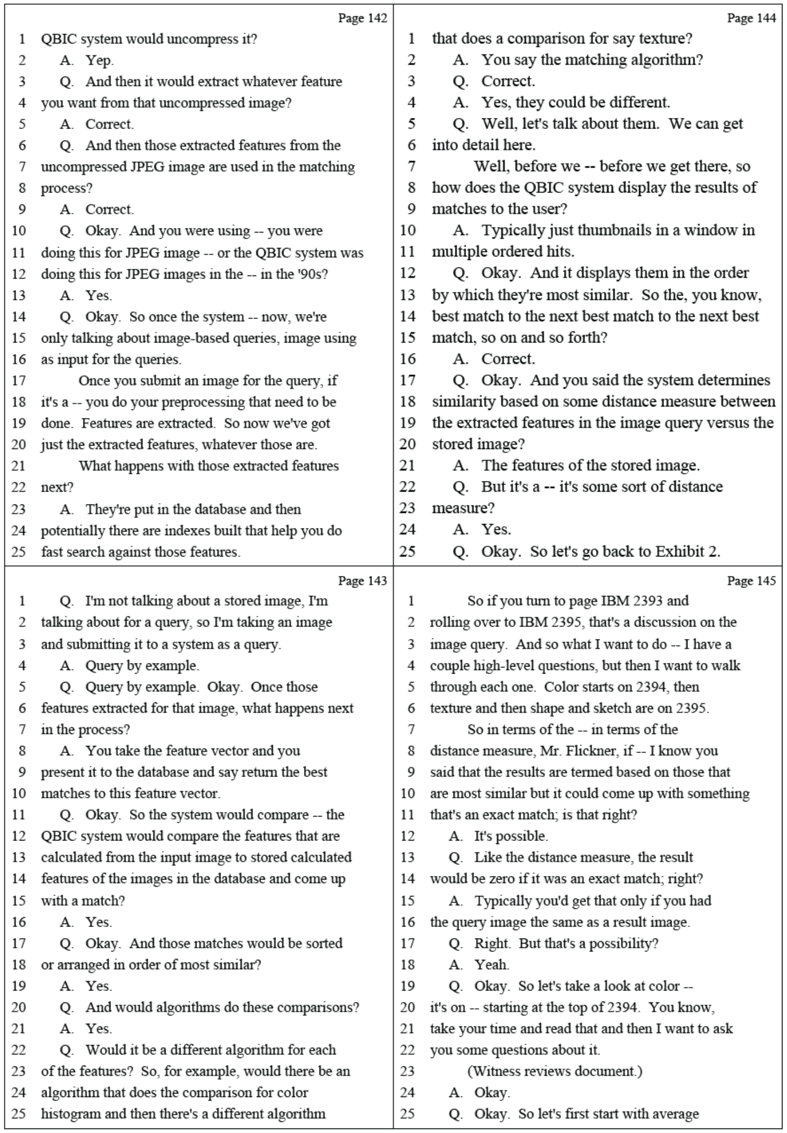

Page 142

QBIC system would uncompress it?
A. Yep.
Q. And then it would extract whatever feature you want from that uncompressed image?
A. Correct.
Q. And then those extracted features from the uncompressed JPEG image are used in the matching process?
A. Correct.
Q. Okay. And you were using -- you were doing this for JPEG image -- or the QBIC system was doing this for JPEG images in the -- in the '90s?
A. Yes.
Q. Okay. So once the system -- now, we're only talking about image-based queries, image using as input for the queries.
Once you submit an image for the query, if it's a -- you do your preprocessing that need to be done. Features are extracted. So now we've got just the extracted features, whatever those are.
What happens with those extracted features next?
A. They're put in the database and then potentially there are indexes built that help you do fast search against those features.

Page 143

Q. I'm not talking about a stored image, I'm talking about for a query, so I'm taking an image and submitting it to a system as a query.
A. Query by example.
Q. Query by example. Okay. Once those features extracted for that image, what happens next in the process?
A. You take the feature vector and you present it to the database and say return the best matches to this feature vector.
Q. Okay. So the system would compare -- the QBIC system would compare the features that are calculated from the input image to stored calculated features of the images in the database and come up with a match?
A. Yes.
Q. Okay. And those matches would be sorted or arranged in order of most similar?
A. Yes.
Q. And would algorithms do these comparisons?
A. Yes.
Q. Would it be a different algorithm for each of the features? So, for example, would there be an algorithm that does the comparison for color histogram and then there's a different algorithm

Page 144

that does a comparison for say texture?
A. You say the matching algorithm?
Q. Correct.
A. Yes, they could be different.
Q. Well, let's talk about them. We can get into detail here.
Well, before we -- before we get there, so how does the QBIC system display the results of matches to the user?
A. Typically just thumbnails in a window in multiple ordered hits.
Q. Okay. And it displays them in the order by which they're most similar. So the, you know, best match to the next best match to the next best match, so on and so forth?
A. Correct.
Q. Okay. And you said the system determines similarity based on some distance measure between the extracted features in the image query versus the stored image?
A. The features of the stored image.
Q. But it's a -- it's some sort of distance measure?
A. Yes.
Q. Okay. So let's go back to Exhibit 2.

Page 145

So if you turn to page IBM 2393 and rolling over to IBM 2395, that's a discussion on the image query. And so what I want to do -- I have a couple high-level questions, but then I want to walk through each one. Color starts on 2394, then texture and then shape and sketch are on 2395.
So in terms of the -- in terms of the distance measure, Mr. Flickner, if -- I know you said that the results are termed based on those that are most similar but it could come up with something that's an exact match; is that right?
A. It's possible.
Q. Like the distance measure, the result would be zero if it was an exact match; right?
A. Typically you'd get that only if you had the query image the same as a result image.
Q. Right. But that's a possibility?
A. Yeah.
Q. Okay. So let's take a look at color -- it's on -- starting at the top of 2394. You know, take your time and read that and then I want to ask you some questions about it.
(Witness reviews document.)
A. Okay.
Q. Okay. So let's first start with average

Page 146

color. Average color determines similarity between the query image and the database image using weighted Euclidean distance; is that right?

A. Yep.

Q. And Euclidean is E-u-c-l-i-d-e-a-n.

And can you just tell us at a high level what -- if you can -- what Euclidean distance is?

A. So if you have multiple features or multiple scaler features, you have to ask the question I'm trying to build a distance from that vector to a query vector.

And as you build that -- the distance, you have -- you may want to weight different elements of the vector differently. And we found out that using variance normalization, where weights of the variance of the feature, gave good aesthetic results.

(Reporter seeks clarification.)

A. Aesthetic results.

Q. And I apologize, I couldn't hear what you said. You said, "And we found out that using variance" . . .

A. So you normalize each feature with its variance.

Q. Normalize each feature with . . .

Page 147

A. With its variance.

Q. With its variance.

A. And then compute the average.

Q. Okay.

A. Or, I'm sorry, compute the weighted sum. So red divided by the variants of red, plus green divided by the variants of green --

(Reporter seeks clarification.)

A. Red divided by the variants of red, green divided by the variants of green, L divided by the variants of L, that's how you combine the various features.

Q. Understand. And L represents what?

A. Luminance.

Q. Okay. And this is -- you're talking specifically with respect to average color?

A. Well, the variance normalization is well known in terms of multidimensional feature merging.

Q. But for average color, you wouldn't use luminance, right, as a -- as a variant?

A. Luminance would just -- you wouldn't use just luminance. You could use luminance . . .

Q. Well, so the color features we talked about before would be RGB, Yiq, LAB --

A. Y is luminance.

Page 148

(Reporter seeks clarification.)

Q. Well, let me ask my -- let me get my question out first.

So we talked about the features for average color before are RGB, Yiq, LAB and MTM.

A. Correct.

Q. And so, for example, for the RGB, you're saying that you would take the R and you would divide it by the variants of red and green, for example --

A. Just variants of red.

Q. Okay. And so you could do the same for the others, for Yiq or LAB?

A. Exactly.

Q. Okay. You do that for both the image input query and the stored image; correct?

A. Correct.

Q. And then you're calculating the distance between them, so determining how similar they are based on the distance?

A. Yeah, or vector distance.

Q. Understood.

A. Two vectors of differences --

(Reporter seeks clarification.)

A. Two vectors with a distance calculation on

Page 149

them. So simply just -- in our product.

Q. You're going to have to repeat that and be just a little bit louder, Mr. Flickner. Sorry.

A. That's all right. Let me get my thoughts together.

So we have multiple dimensional features. It's not unusual to take each dimension and normalize it by the variance because that gives you -- what that does is if you have tight variances you divide it by, you get bigger numbers. So you have more information close by.

Q. What was the word you said before "close by"?

A. I don't remember.

THE REPORTER: Information?

THE WITNESS: Information close by?

BY MR. EDWARDS:

Q. Okay. Okay. So now let's talk about color histograms.

A. Okay.

Q. Okay? Well, stop. Let's go back to average color. So is there -- you described the weighted Euclidean distance that's calculated between the query image and the database image for average color.

		Id. at 159 – 160: 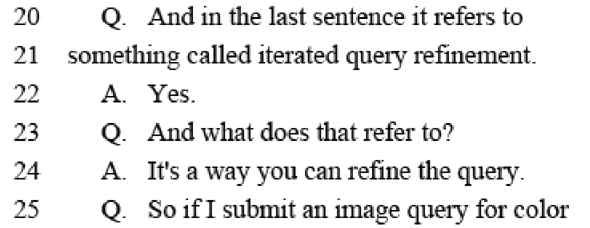 Q. And in the last sentence it refers to something called iterated query refinement. A. Yes. Q. And what does that refer to? A. It's a way you can refine the query. Q. So if I submit an image query for color 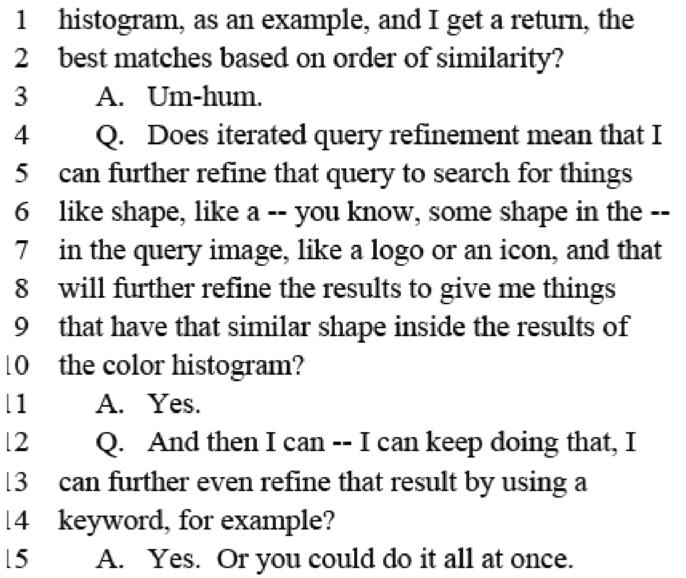 histogram, as an example, and I get a return, the best matches based on order of similarity? A. Um-hum. Q. Does iterated query refinement mean that I can further refine that query to search for things like shape, like a -- you know, some shape in the -- in the query image, like a logo or an icon, and that will further refine the results to give me things that have that similar shape inside the results of the color histogram? A. Yes. Q. And then I can -- I can keep doing that, I can further even refine that result by using a keyword, for example? A. Yes. Or you could do it all at once. *Id.* at 305:

		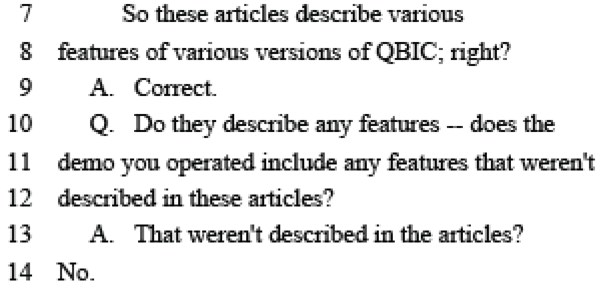 So these articles describe various features of various versions of QBIC; right? A. Correct. Q. Do they describe any features -- does the demo you operated include any features that weren't described in these articles? A. That weren't described in the articles? No. *Id.* at 35: 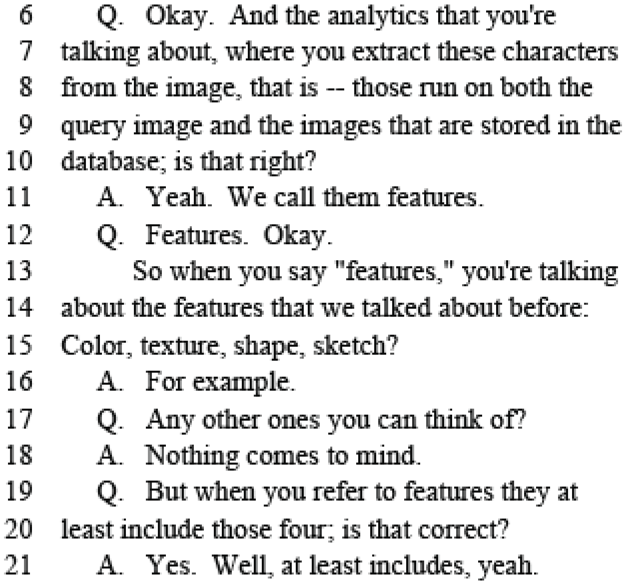 Q. Okay. And the analytics that you're talking about, where you extract these characters from the image, that is -- those run on both the query image and the images that are stored in the database; is that right? A. Yeah. We call them features. Q. Features. Okay. So when you say "features," you're talking about the features that we talked about before: Color, texture, shape, sketch? A. For example. Q. Any other ones you can think of? A. Nothing comes to mind. Q. But when you refer to features they at least include those four; is that correct? A. Yes. Well, at least includes, yeah. *Id.* at 109-110:

		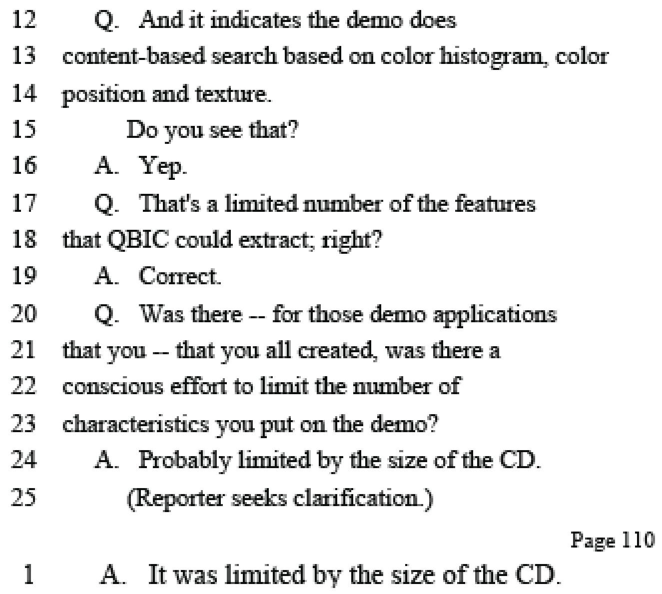 Q. And it indicates the demo does content-based search based on color histogram, color position and texture. Do you see that? A. Yep. Q. That's a limited number of the features that QBIC could extract; right? A. Correct. Q. Was there -- for those demo applications that you -- that you all created, was there a conscious effort to limit the number of characteristics you put on the demo? A. Probably limited by the size of the CD. (Reporter seeks clarification.) Page 110 A. It was limited by the size of the CD. *Id.* at 311: 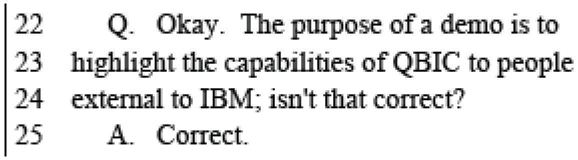 Q. Okay. The purpose of a demo is to highlight the capabilities of QBIC to people external to IBM; isn't that correct? A. Correct. *Id.* at 313-317:

		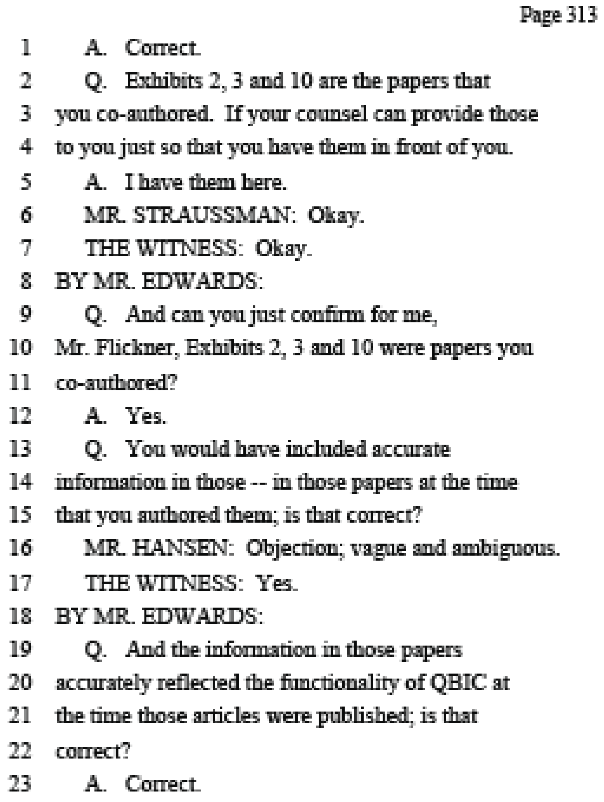 Page 313 A. Correct. Q. Exhibits 2, 3 and 10 are the papers that you co-authored. If your counsel can provide those to you just so that you have them in front of you. A. I have them here. MR. STRAUSSMAN: Okay. THE WITNESS: Okay. BY MR. EDWARDS: Q. And can you just confirm for me, Mr. Flickner, Exhibits 2, 3 and 10 were papers you co-authored? A. Yes. Q. You would have included accurate information in those -- in those papers at the time that you authored them; is that correct? MR. HANSEN: Objection; vague and ambiguous. THE WITNESS: Yes. BY MR. EDWARDS: Q. And the information in those papers accurately reflected the functionality of QBIC at the time those articles were published; is that correct? A. Correct.

		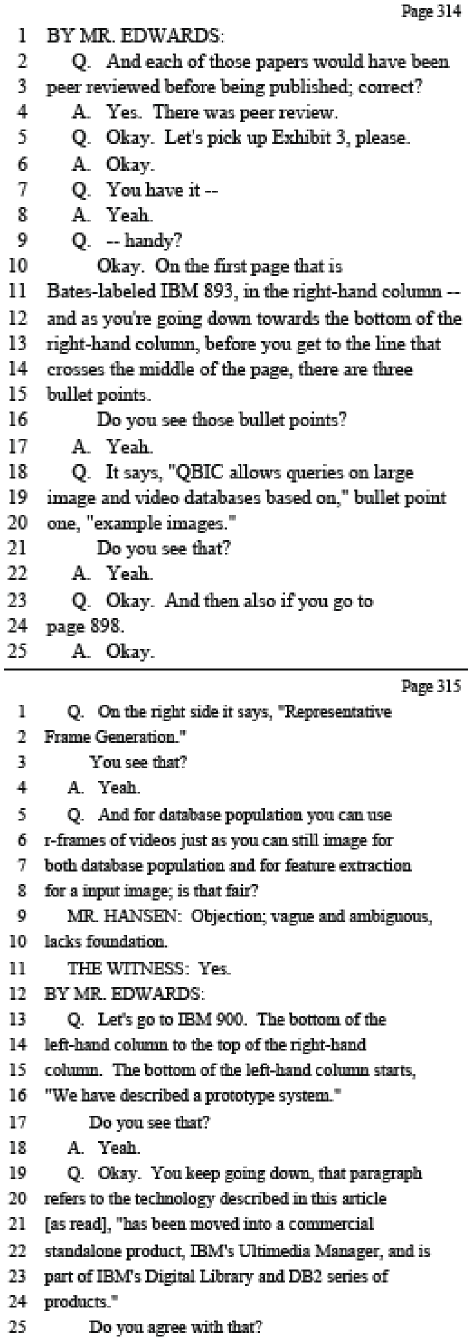 Page 314 BY MR. EDWARDS: Q. And each of those papers would have been peer reviewed before being published; correct? A. Yes. There was peer review. Q. Okay. Let's pick up Exhibit 3, please. A. Okay. Q. You have it -- A. Yeah. Q. -- handy? Okay. On the first page that is Bates-labeled IBM 893, in the right-hand column -- and as you're going down towards the bottom of the right-hand column, before you get to the line that crosses the middle of the page, there are three bullet points. Do you see those bullet points? A. Yeah. Q. It says, "QBIC allows queries on large image and video databases based on," bullet point one, "example images." Do you see that? A. Yeah. Q. Okay. And then also if you go to page 898. A. Okay. Page 315 Q. On the right side it says, "Representative Frame Generation." You see that? A. Yeah. Q. And for database population you can use r-frames of videos just as you can still image for both database population and for feature extraction for a input image; is that fair? MR. HANSEN: Objection; vague and ambiguous, lacks foundation. THE WITNESS: Yes. BY MR. EDWARDS: Q. Let's go to IBM 900. The bottom of the left-hand column to the top of the right-hand column. The bottom of the left-hand column starts, "We have described a prototype system." Do you see that? A. Yeah. Q. Okay. You keep going down, that paragraph refers to the technology described in this article [as read], "has been moved into a commercial standalone product, IBM's Ultimedia Manager, and is part of IBM's Digital Library and DB2 series of products." Do you agree with that?

		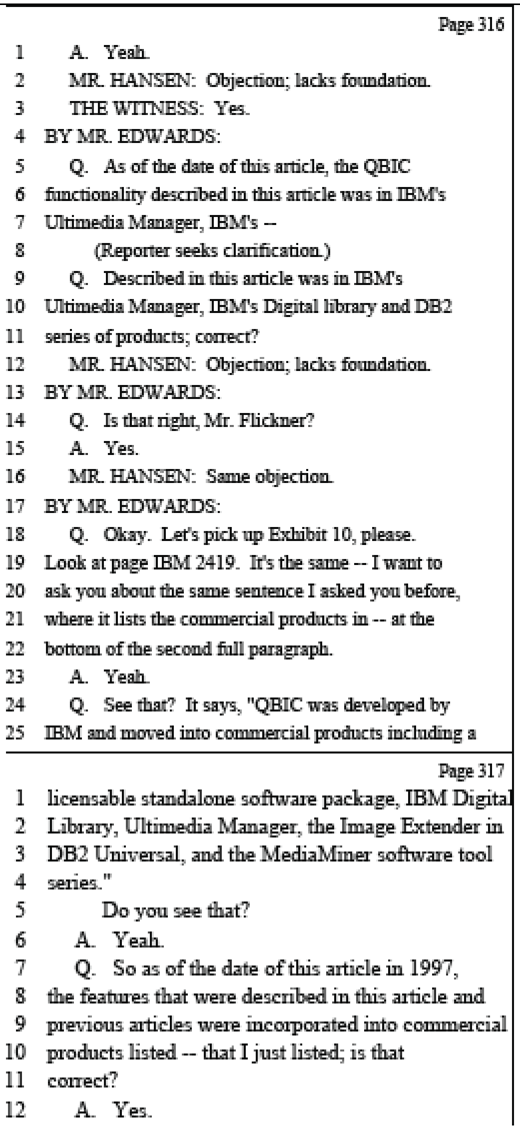 Page 316 A. Yeah. MR. HANSEN: Objection; lacks foundation. THE WITNESS: Yes. BY MR. EDWARDS: Q. As of the date of this article, the QBIC functionality described in this article was in IBM's Ultimedia Manager, IBM's -- (Reporter seeks clarification.) Q. Described in this article was in IBM's Ultimedia Manager, IBM's Digital library and DB2 series of products; correct? MR. HANSEN: Objection; lacks foundation. BY MR. EDWARDS: Q. Is that right, Mr. Flickner? A. Yes. MR. HANSEN: Same objection. BY MR. EDWARDS: Q. Okay. Let's pick up Exhibit 10, please. Look at page IBM 2419. It's the same -- I want to ask you about the same sentence I asked you before, where it lists the commercial products in -- at the bottom of the second full paragraph. A. Yeah. Q. See that? It says, "QBIC was developed by IBM and moved into commercial products including a Page 317 licensable standalone software package, IBM Digital Library, Ultimedia Manager, the Image Extender in DB2 Universal, and the MediaMiner software tool series." Do you see that? A. Yeah. Q. So as of the date of this article in 1997, the features that were described in this article and previous articles were incorporated into commercial products listed -- that I just listed; is that correct? A. Yes. *Id.* at 321 – 322:

		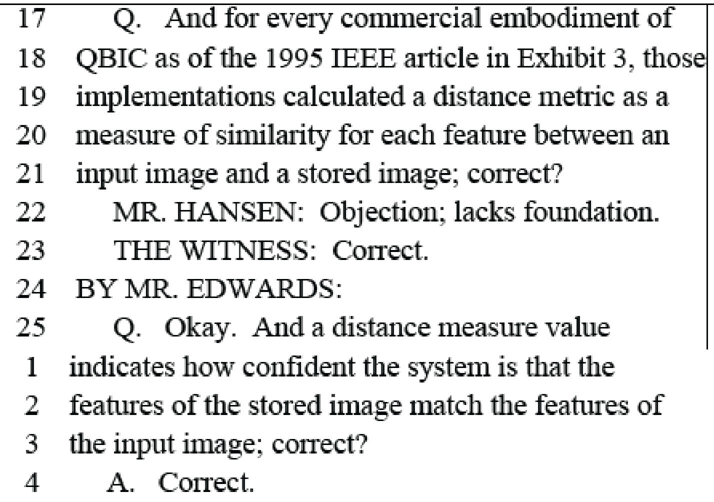 Q. And for every commercial embodiment of QBIC as of the 1995 IEEE article in Exhibit 3, those implementations calculated a distance metric as a measure of similarity for each feature between an input image and a stored image; correct? MR. HANSEN: Objection; lacks foundation. THE WITNESS: Correct. BY MR. EDWARDS: Q. Okay. And a distance measure value indicates how confident the system is that the features of the stored image match the features of the input image; correct? A. Correct. Flickner 8/29/23 Depo at Ex. 6 at 1:

		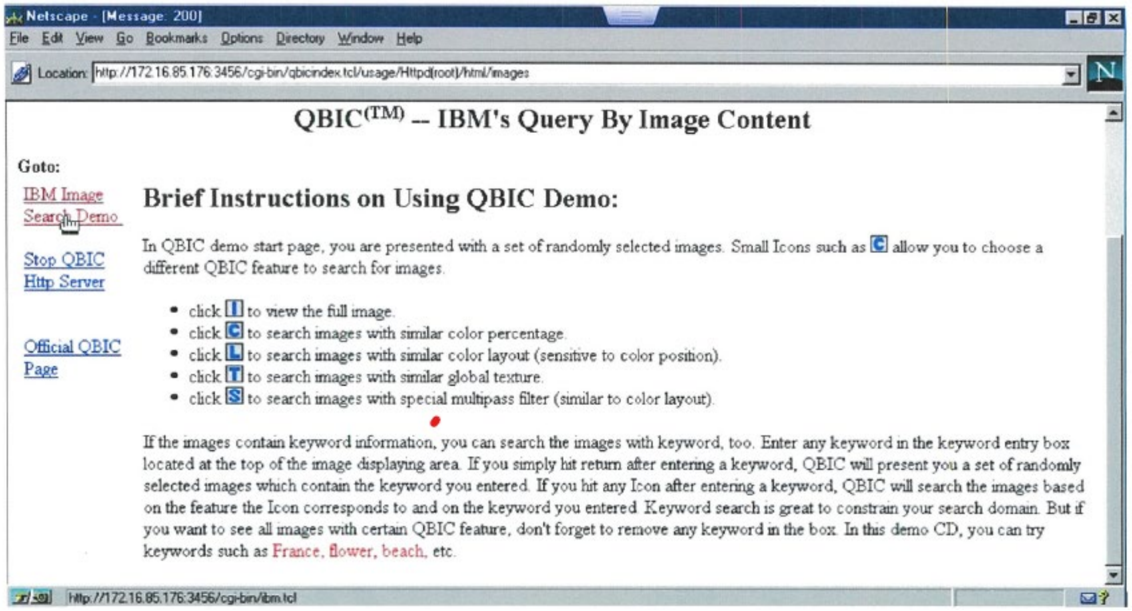 **QBIC(TM) -- IBM's Query By Image Content** **Brief Instructions on Using QBIC Demo:** In QBIC demo start page, you are presented with a set of randomly selected images. Small Icons such as C allow you to choose a different QBIC feature to search for images. - click I to view the full image. - click C to search images with similar color percentage. - click L to search images with similar color layout (sensitive to color position). - click T to search images with similar global texture. - click S to search images with special multipass filter (similar to color layout). If the images contain keyword information, you can search the images with keyword, too. Enter any keyword in the keyword entry box located at the top of the image displaying area. If you simply hit return after entering a keyword, QBIC will present you a set of randomly selected images which contain the keyword you entered. If you hit any Icon after entering a keyword, QBIC will search the images based on the feature the Icon corresponds to and on the keyword you entered. Keyword search is great to constrain your search domain. But if you want to see all images with certain QBIC feature, don't forget to remove any keyword in the box. In this demo CD, you can try keywords such as France, flower, beach, etc. *Id.* at 2:

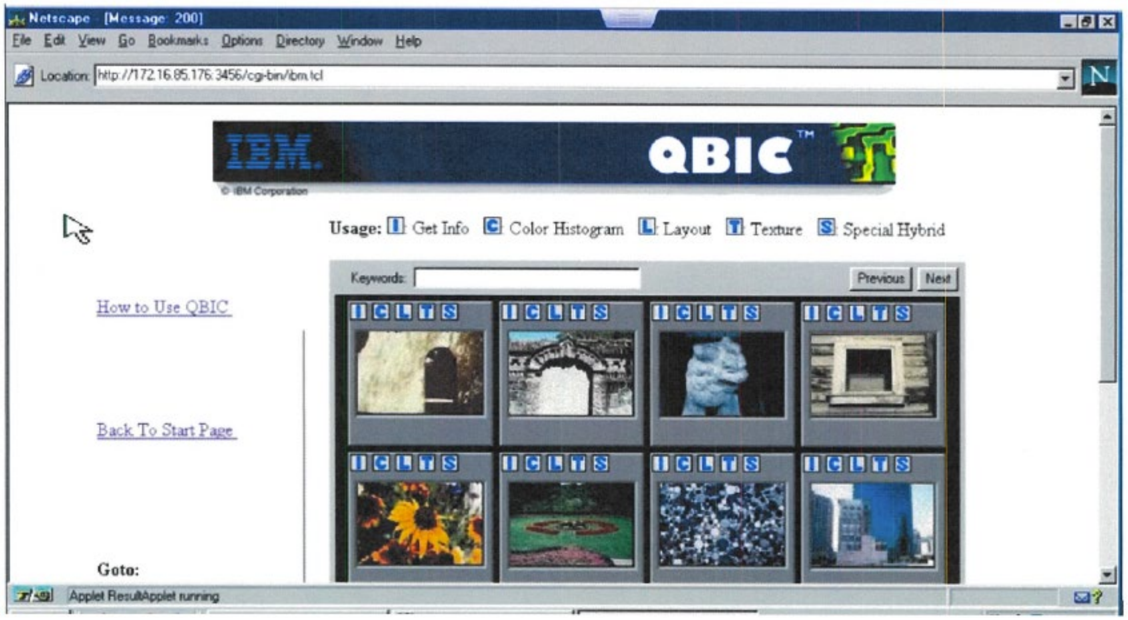

The QBIC system was implemented in IBM commercial products, including in IBM's Ultimedia Manager, the DB2 series of products such as the Image Extender in DB2 Universal Database, and the Enterprise Information Portal. All of the features referenced in this chart were incorporated into these commercial implementations before November 6, 2000, as reflected in a combination of scientific papers, IBM's internal documentation, and the QBIC demo software applications. *See, e.g.*, Flickner 8/29/23 Depo Tr. at 33:9-24, 71:21-72:10, 164:21-166:2, 166:11-175:24, 178:5-186:17, 186:21-192:4, 192:14-194:1, 195:16-196:4, 194:1-308:1; IBM 000001-617; IBM 000747; IBM 000777-880; IBM 000893-902; IBM 0002315-320; IBM 0002390-404; IBM 0002418-430.

To the extent it is found that QBIC does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that QBIC does not anticipate this claim, QBIC renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references

		identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine QBIC with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

3. **Claim 20**

Claim	'529 Claim Recitation	Exemplary Citations in Reference
20	The system of claim 1 wherein the data processing operation is selected from the group consisting of contrast enhancement, noise removal, and de-blurring.	To the extent it is found that QBIC does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that QBIC does not anticipate this claim, QBIC renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine QBIC with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

Exhibit G-31

Bank of America Corporation ("BofA" or "Defendant") provides this chart subject to all reservations, objections, statements, and disclaimers set forth herein and in BofA's Final Invalidity Contentions Cover Pleading, as well as any amendment, supplement, or modification thereof, which are incorporated herein by reference in their entirety.

On information and belief, CoolTown was made, known or used by others in the United States no later than July 10, 2000. *See generally* T. Kindberg et al., A Web-Based Nomadic Computing System; J. Barton et al., The Cooltown User Experience; cooltown developer's network, coolbase overview; cooltown developer's network, people, places, things: web presence for the real world. CoolTown is available as prior art at least under 35 U.S.C. § 102(a). On information and belief, HP was using CoolTown and made it available to the public in the United States by releasing the website on the Internet (and accessible to U.S. users) by at least March 2000, making the software available on the HP website by March 2000, and making available on the Internet (and accessible to U.S. users) the open source code, which reflects the functionality of CoolTown by at least March 2000. *See e.g.*, HP000043–044; HP000045–046; HP000060–061; *see also* https://www.zdnet.com/article/hp-brings-open-source-to-cooltown/. Therefore, HP Cooltown was at least known, used, and/or sold to the public by at least 2000. On information and belief, CoolTown is also available as prior art under 102(g). BofA is diligently investigating CoolTown and will supplement with additional information, if necessary.

As shown in the chart below, CoolTown anticipates asserted claims 4 and 20 of the '529 Patent. To the extent CoolTown is found not to expressly disclose certain limitations in the asserted claims, such limitations are inherent. To the extent CoolTown is found not to anticipate any asserted claims or claim elements of the '529 Patent, the reference nevertheless renders those claims or claim elements obvious under 35 U.S.C. § 103, either alone or in combination with other art identified herein or in the cover pleading of BofA's Invalidity Contentions, for the reasons identified in the cover pleading. These Final Invalidity Contentions are not an admission by BofA that the accused products, including any current or past versions of the accused products, are covered by, or infringe the asserted claims, but are based instead on the recognition that if the claims are interpreted to be broad enough to encompass the accused products, the claims must also be construed to have that same scope when considering whether they are invalid.

The following chart is based on the Court's claim construction rulings (Dkt. Nos. 145, 153, 236), and to the extent not contradictory, Plaintiffs' apparent interpretations of the claims in Plaintiff's "Preliminary" Final Infringement Contentions. BofA notes that in many instances, Plaintiff's "Preliminary" Final Infringement Contentions ignore and/or contradict the Court's claim construction rulings. BofA does not accept the assumptions concerning the scope and meaning implicit in Plaintiff's Preliminary Final Infringement Contentions, to the extent those assumptions are discernible, and reserves the right to challenge Plaintiff's proposed (or implied) constructions. Furthermore, Plaintiffs' "Preliminary" Final Contentions fail to comply with the level of specificity required by SPR

2.1.2 and 2.1.3, and include improper reservation of rights to supplement with "additional ways in which BoA Accused Products infringe." Any alleged deficient discovery is of Plaintiffs own making. Most if not all of the deficiencies in Plaintiffs' Preliminary Final Contentions have been documented by Defendants through multiple letters since at least May 18, 2021 and up to Bank of America's most recent letter of August 19, 2023. Plaintiffs' shifting infringement contentions, deficiencies in its infringement contentions, and changing claim scope have prejudiced BofA from being able to reasonably prepare its defenses, including its invalidity contentions. BofA also reserves the right to revise and supplement these charts if and when Plaintiff is permitted to provide revised or supplemental Infringement Contentions. Citations given in the chart below are merely representative of the respective elements to provide notice to Plaintiffs' of BofA's invalidity theories and are not meant to be exhaustive.

1. Claim 1

Element	’529 Claim Recitation	Exemplary Citations in Reference
[1Pre]	A system comprising:	CoolTown discloses a system. The Cooltown User Experience (HP000088 – 093) at 1: 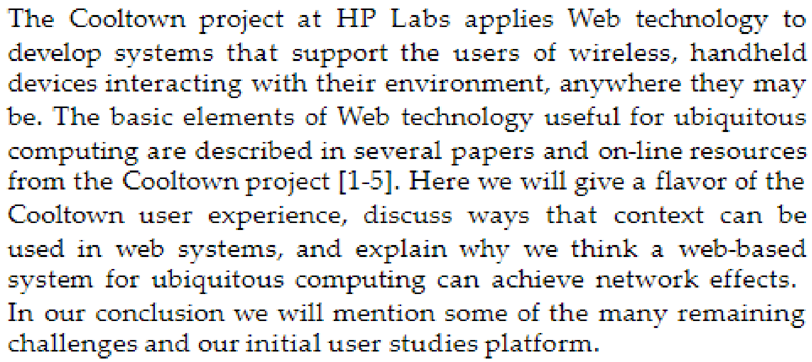 The Cooltown project at HP Labs applies Web technology to develop systems that support the users of wireless, handheld devices interacting with their environment, anywhere they may be. The basic elements of Web technology useful for ubiquitous computing are described in several papers and on-line resources from the Cooltown project [1-5]. Here we will give a flavor of the Cooltown user experience, discuss ways that context can be used in web systems, and explain why we think a web-based system for ubiquitous computing can achieve network effects. In our conclusion we will mention some of the many remaining challenges and our initial user studies platform. A Web-Based Nomadic Computing System (HP000074 – 87) at 1 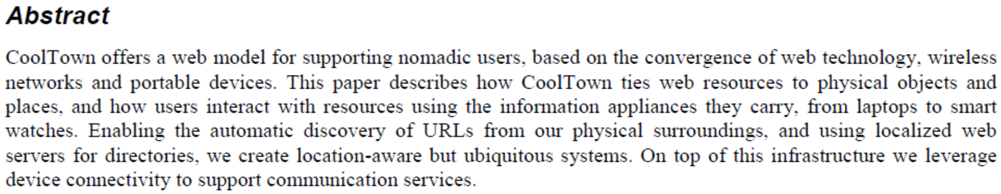 **Abstract** CoolTown offers a web model for supporting nomadic users, based on the convergence of web technology, wireless networks and portable devices. This paper describes how CoolTown ties web resources to physical objects and places, and how users interact with resources using the information appliances they carry, from laptops to smart watches. Enabling the automatic discovery of URLs from our physical surroundings, and using localized web servers for directories, we create location-aware but ubiquitous systems. On top of this infrastructure we leverage device connectivity to support communication services. Coolbase[1] Overview at 1

---

[1] *See* https://web.archive.org/web/20011030090650/http://cooltown.hp.com/dev/coolbase-overview.asp

Element	'529 Claim Recitation	Exemplary Citations in Reference
		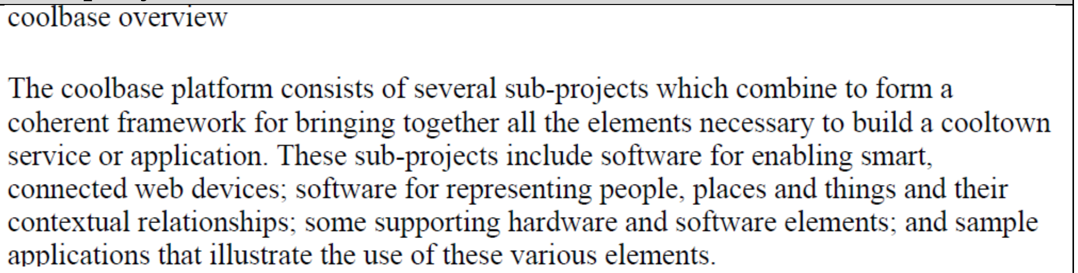 coolbase overview The coolbase platform consists of several sub-projects which combine to form a coherent framework for bringing together all the elements necessary to build a cooltown service or application. These sub-projects include software for enabling smart, connected web devices; software for representing people, places and things and their contextual relationships; some supporting hardware and software elements; and sample applications that illustrate the use of these various elements.
[1A]	a camera that captures an image;	CoolTown discloses a camera that captures an image. The Cooltown: People, Places, Things Web Presence for the Real World (HP000025 – 34) at 20:

Element	'529 Claim Recitation	Exemplary Citations in Reference
		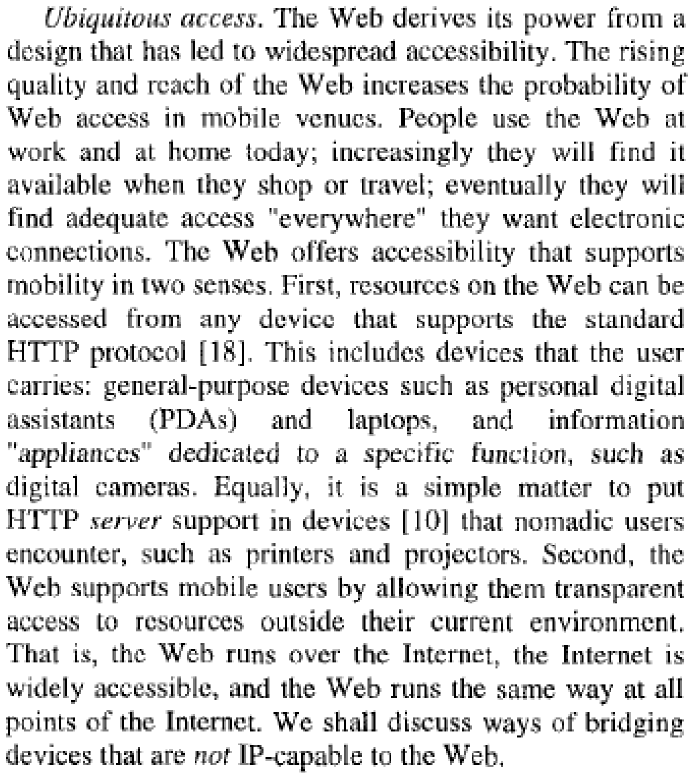 *Ubiquitous access.* The Web derives its power from a design that has led to widespread accessibility. The rising quality and reach of the Web increases the probability of Web access in mobile venues. People use the Web at work and at home today; increasingly they will find it available when they shop or travel; eventually they will find adequate access "everywhere" they want electronic connections. The Web offers accessibility that supports mobility in two senses. First, resources on the Web can be accessed from any device that supports the standard HTTP protocol [18]. This includes devices that the user carries: general-purpose devices such as personal digital assistants (PDAs) and laptops, and information "appliances" dedicated to a specific function, such as digital cameras. Equally, it is a simple matter to put HTTP *server* support in devices [10] that nomadic users encounter, such as printers and projectors. Second, the Web supports mobile users by allowing them transparent access to resources outside their current environment. That is, the Web runs over the Internet, the Internet is widely accessible, and the Web runs the same way at all points of the Internet. We shall discuss ways of bridging devices that are *not* IP-capable to the Web. The Cooltown: People, Places, Things Web Presence for the Real World (HP000025 – 34) at 24:

Element	'529 Claim Recitation	Exemplary Citations in Reference
		3.4. Infrastructure for things; ChaiServer and eSquirt Given a URL from one of the lower-layer services, we 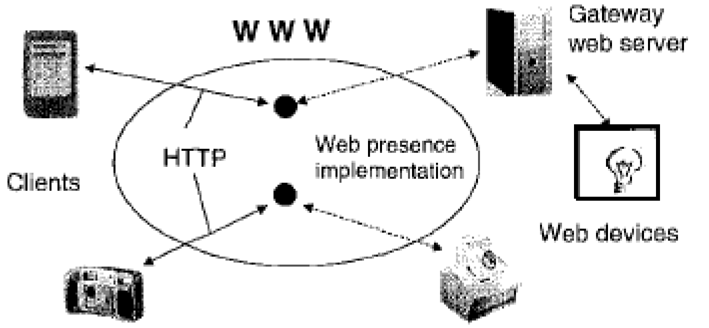 **Figure 5. Devices (printer, light), with and without embedded web servers.** Figure 5 shows interactions between two clients and their respective web devices. It shows a digital camera that interacts with a printer via HTTP operations on the printer's point of web presence. The URL for the printer resolves to its network address. The camera first interrogates the printer to find out the formats it accepts. The printer can supply this as an HTML page or as an XML document that describes it. Then the camera sends the image as, for example, JPEG encoded data. Figure 5 also shows lights connected to a home PC by a controller unit. The home PC acts as a gateway to the lighting unit. A remote user issues HTTP operations via a web browser on a PDA; the programs on the PC that handle HTTP operations directed to the light's URL manipulate the lighting unit accordingly.

Element	’529 Claim Recitation	Exemplary Citations in Reference
		The Cooltown User Experience (HP000088 – 093) at 1: 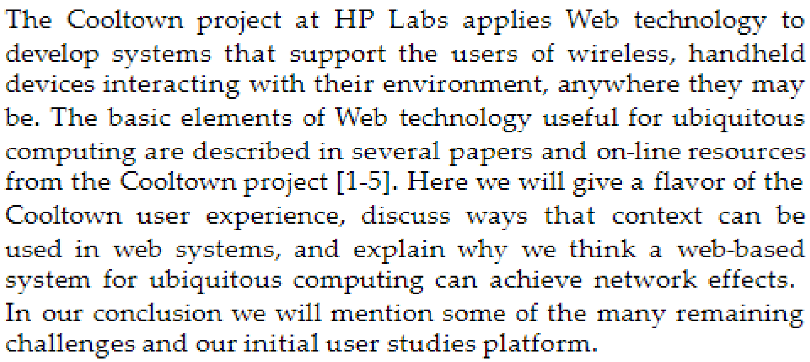 The Cooltown project at HP Labs applies Web technology to develop systems that support the users of wireless, handheld devices interacting with their environment, anywhere they may be. The basic elements of Web technology useful for ubiquitous computing are described in several papers and on-line resources from the Cooltown project [1-5]. Here we will give a flavor of the Cooltown user experience, discuss ways that context can be used in web systems, and explain why we think a web-based system for ubiquitous computing can achieve network effects. In our conclusion we will mention some of the many remaining challenges and our initial user studies platform. A Web-Based Nomadic Computing System (HP000074 – 87) at 2 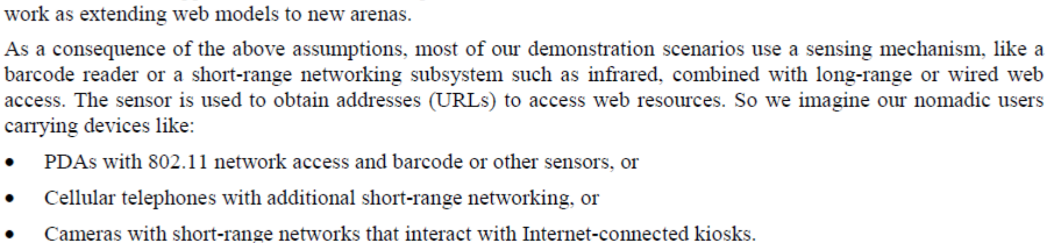 The CoolTown work is not on the devices or wireless communication layers, but rather on the layer of infrastructure that is needed to support nomadic users with portable devices. We work at the "web server" level and we see our work as extending web models to new arenas. As a consequence of the above assumptions, most of our demonstration scenarios use a sensing mechanism, like a barcode reader or a short-range networking subsystem such as infrared, combined with long-range or wired web access. The sensor is used to obtain addresses (URLs) to access web resources. So we imagine our nomadic users carrying devices like: • PDAs with 802.11 network access and barcode or other sensors, or • Cellular telephones with additional short-range networking, or • Cameras with short-range networks that interact with Internet-connected kiosks. A Web-Based Nomadic Computing System (HP000074 – 87) at 3

Element	'529 Claim Recitation	Exemplary Citations in Reference
		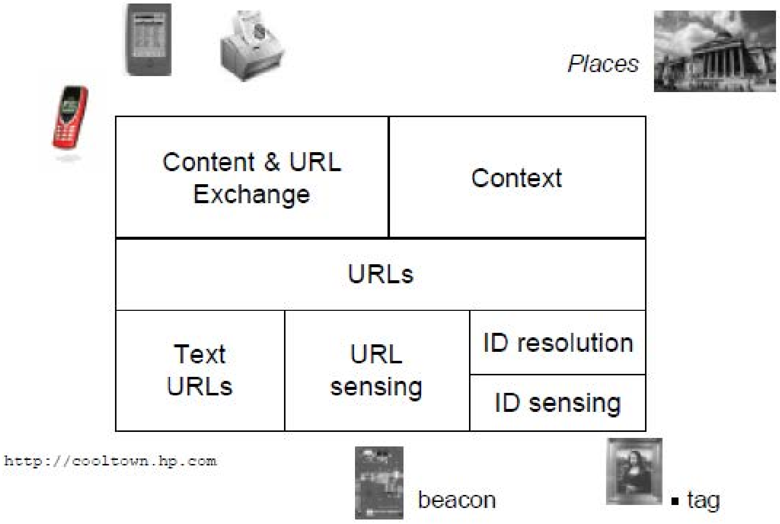 Figure 2. Component relationships. Sources of URLs, including keyboards, beacons, and resolved tags point to context or content. A Web-Based Nomadic Computing System (HP000074 – 87) at 4 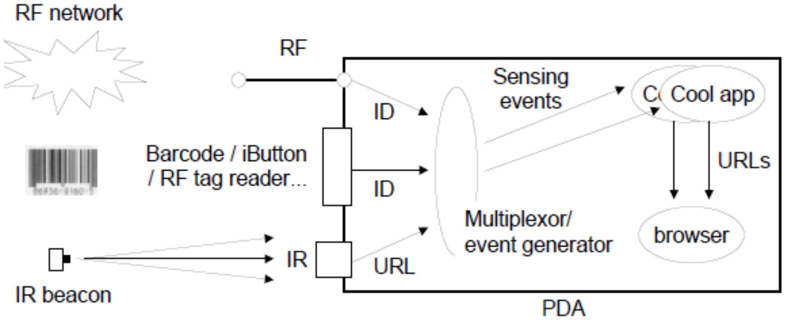 Figure 3. A sketch of a PDA equipped for sensing URLs and identifiers.

Element	'529 Claim Recitation	Exemplary Citations in Reference
		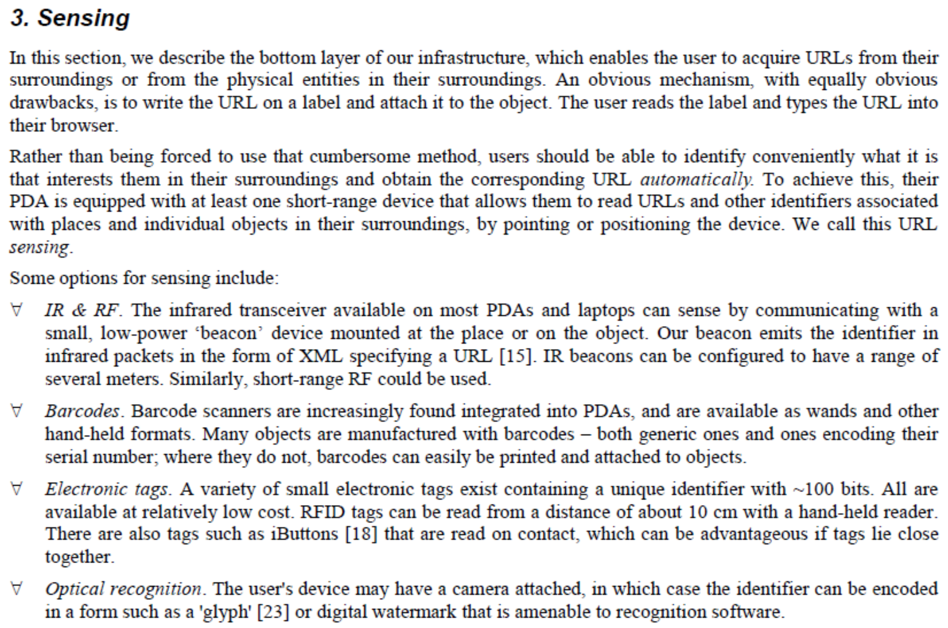 **3. Sensing** In this section, we describe the bottom layer of our infrastructure, which enables the user to acquire URLs from their surroundings or from the physical entities in their surroundings. An obvious mechanism, with equally obvious drawbacks, is to write the URL on a label and attach it to the object. The user reads the label and types the URL into their browser. Rather than being forced to use that cumbersome method, users should be able to identify conveniently what it is that interests them in their surroundings and obtain the corresponding URL *automatically*. To achieve this, their PDA is equipped with at least one short-range device that allows them to read URLs and other identifiers associated with places and individual objects in their surroundings, by pointing or positioning the device. We call this URL *sensing*. Some options for sensing include: ∀ *IR & RF*. The infrared transceiver available on most PDAs and laptops can sense by communicating with a small, low-power 'beacon' device mounted at the place or on the object. Our beacon emits the identifier in infrared packets in the form of XML specifying a URL [15]. IR beacons can be configured to have a range of several meters. Similarly, short-range RF could be used. ∀ *Barcodes*. Barcode scanners are increasingly found integrated into PDAs, and are available as wands and other hand-held formats. Many objects are manufactured with barcodes – both generic ones and ones encoding their serial number; where they do not, barcodes can easily be printed and attached to objects. ∀ *Electronic tags*. A variety of small electronic tags exist containing a unique identifier with ~100 bits. All are available at relatively low cost. RFID tags can be read from a distance of about 10 cm with a hand-held reader. There are also tags such as iButtons [18] that are read on contact, which can be advantageous if tags lie close together. ∀ *Optical recognition*. The user's device may have a camera attached, in which case the identifier can be encoded in a form such as a 'glyph' [23] or digital watermark that is amenable to recognition software. To the extent it is found that CoolTown does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that CoolTown does not anticipate this claim, CoolTown renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine CoolTown with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

Element	'529 Claim Recitation	Exemplary Citations in Reference
[1B]	a network-enabled device that conducts a data processing operation on at least a portion of the image to produce data, and sends the data to a service;	CoolTown discloses a network-enabled device that conducts a data processing operation on at least a portion of the image to produce data, and sends the data to a service. The Cooltown: People, Places, Things Web Presence for the Real World (HP000025 – 34) at 20: 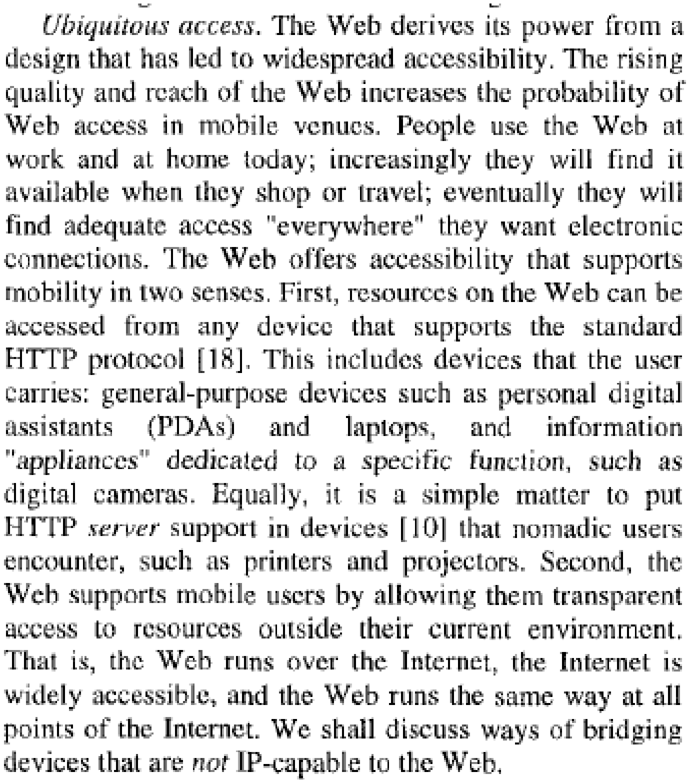 *Ubiquitous access*. The Web derives its power from a design that has led to widespread accessibility. The rising quality and reach of the Web increases the probability of Web access in mobile venues. People use the Web at work and at home today; increasingly they will find it available when they shop or travel; eventually they will find adequate access "everywhere" they want electronic connections. The Web offers accessibility that supports mobility in two senses. First, resources on the Web can be accessed from any device that supports the standard HTTP protocol [18]. This includes devices that the user carries: general-purpose devices such as personal digital assistants (PDAs) and laptops, and information "appliances" dedicated to a specific function, such as digital cameras. Equally, it is a simple matter to put HTTP *server* support in devices [10] that nomadic users encounter, such as printers and projectors. Second, the Web supports mobile users by allowing them transparent access to resources outside their current environment. That is, the Web runs over the Internet, the Internet is widely accessible, and the Web runs the same way at all points of the Internet. We shall discuss ways of bridging devices that are *not* IP-capable to the Web. The Cooltown: People, Places, Things Web Presence for the Real World (HP000025 – 34) at 24:

Element	'529 Claim Recitation	Exemplary Citations in Reference
		3.4. Infrastructure for things; ChaiServer and eSquirt Given a URL from one of the lower-layer services, we 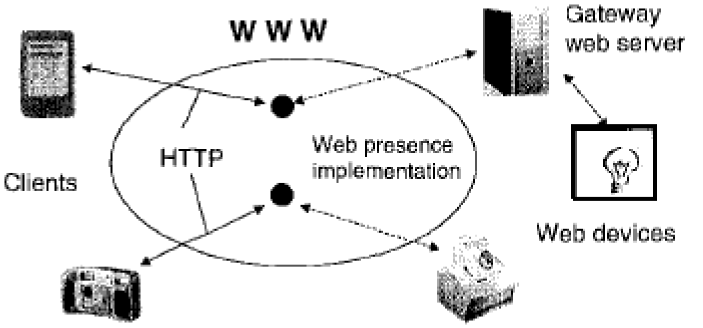 **Figure 5. Devices (printer, light), with and without embedded web servers.** Figure 5 shows interactions between two clients and their respective web devices. It shows a digital camera that interacts with a printer via HTTP operations on the

Element	'529 Claim Recitation	Exemplary Citations in Reference
		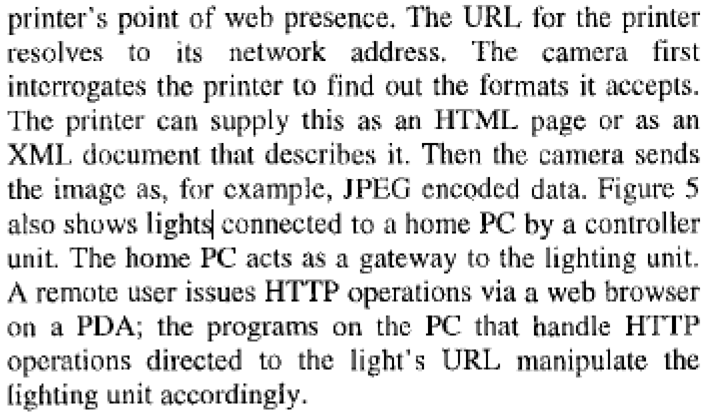 printer's point of web presence. The URL for the printer resolves to its network address. The camera first interrogates the printer to find out the formats it accepts. The printer can supply this as an HTML page or as an XML document that describes it. Then the camera sends the image as, for example, JPEG encoded data. Figure 5 also shows lights connected to a home PC by a controller unit. The home PC acts as a gateway to the lighting unit. A remote user issues HTTP operations via a web browser on a PDA; the programs on the PC that handle HTTP operations directed to the light's URL manipulate the lighting unit accordingly. The Cooltown User Experience (HP000088 – 093) at 1: 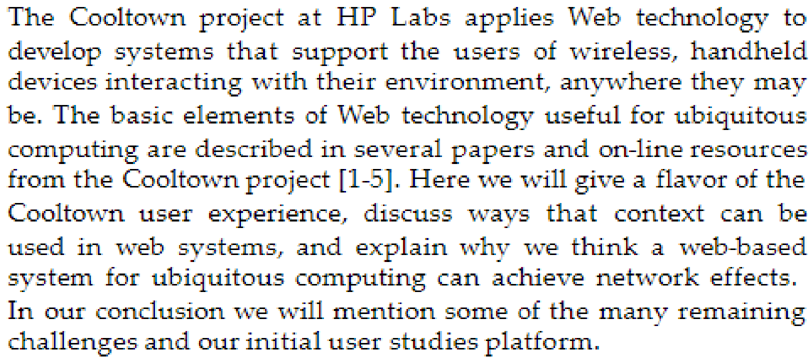 The Cooltown project at HP Labs applies Web technology to develop systems that support the users of wireless, handheld devices interacting with their environment, anywhere they may be. The basic elements of Web technology useful for ubiquitous computing are described in several papers and on-line resources from the Cooltown project [1-5]. Here we will give a flavor of the Cooltown user experience, discuss ways that context can be used in web systems, and explain why we think a web-based system for ubiquitous computing can achieve network effects. In our conclusion we will mention some of the many remaining challenges and our initial user studies platform. A Web-Based Nomadic Computing System (HP000074 – 87) at 1

Element	'529 Claim Recitation	Exemplary Citations in Reference
		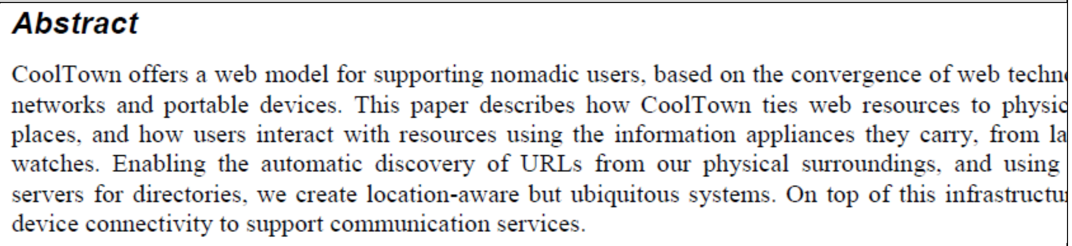 ***Abstract*** CoolTown offers a web model for supporting nomadic users, based on the convergence of web techn networks and portable devices. This paper describes how CoolTown ties web resources to physic places, and how users interact with resources using the information appliances they carry, from la watches. Enabling the automatic discovery of URLs from our physical surroundings, and using servers for directories, we create location-aware but ubiquitous systems. On top of this infrastructu device connectivity to support communication services. A Web-Based Nomadic Computing System (HP000074 – 87) at 2 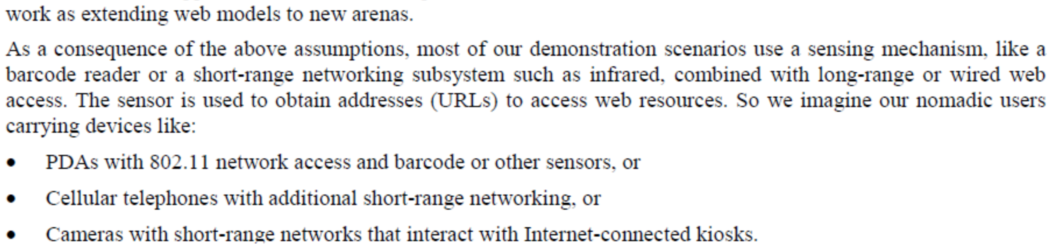 The CoolTown work is not on the devices or wireless communication layers, but rather on the layer of infrastructure that is needed to support nomadic users with portable devices. We work at the "web server" level and we see our work as extending web models to new arenas. As a consequence of the above assumptions, most of our demonstration scenarios use a sensing mechanism, like a barcode reader or a short-range networking subsystem such as infrared, combined with long-range or wired web access. The sensor is used to obtain addresses (URLs) to access web resources. So we imagine our nomadic users carrying devices like: • PDAs with 802.11 network access and barcode or other sensors, or • Cellular telephones with additional short-range networking, or • Cameras with short-range networks that interact with Internet-connected kiosks. A Web-Based Nomadic Computing System (HP000074 – 87) at 3 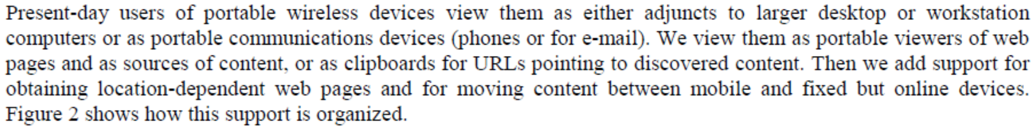 Present-day users of portable wireless devices view them as either adjuncts to larger desktop or workstation computers or as portable communications devices (phones or for e-mail). We view them as portable viewers of web pages and as sources of content, or as clipboards for URLs pointing to discovered content. Then we add support for obtaining location-dependent web pages and for moving content between mobile and fixed but online devices. Figure 2 shows how this support is organized. A Web-Based Nomadic Computing System (HP000074 – 87) at 3

Element	'529 Claim Recitation	Exemplary Citations in Reference
		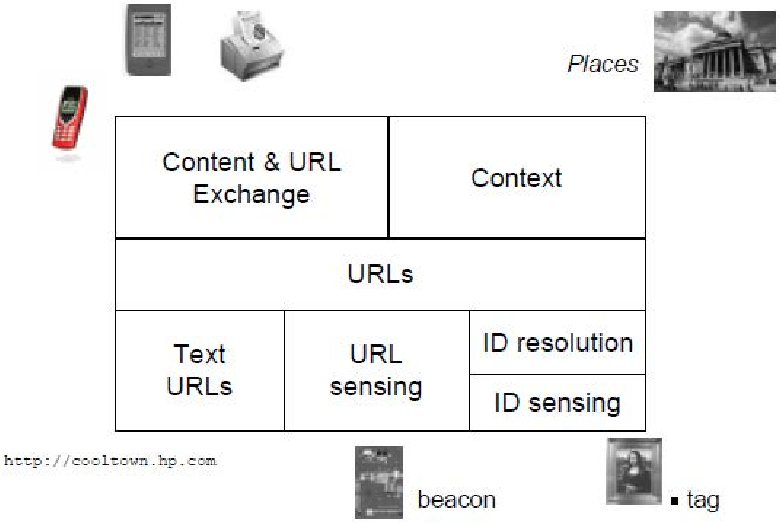 Figure 2. Component relationships. Sources of URLs, including keyboards, beacons, and resolved tags point to context or content. A Web-Based Nomadic Computing System (HP000074 – 87) at 4 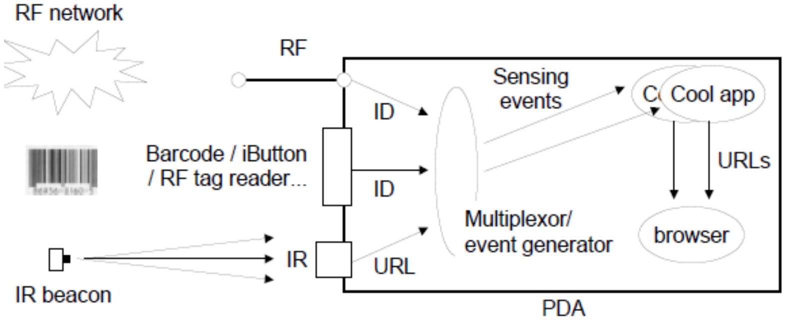 Figure 3. A sketch of a PDA equipped for sensing URLs and identifiers.

Element	'529 Claim Recitation	Exemplary Citations in Reference
		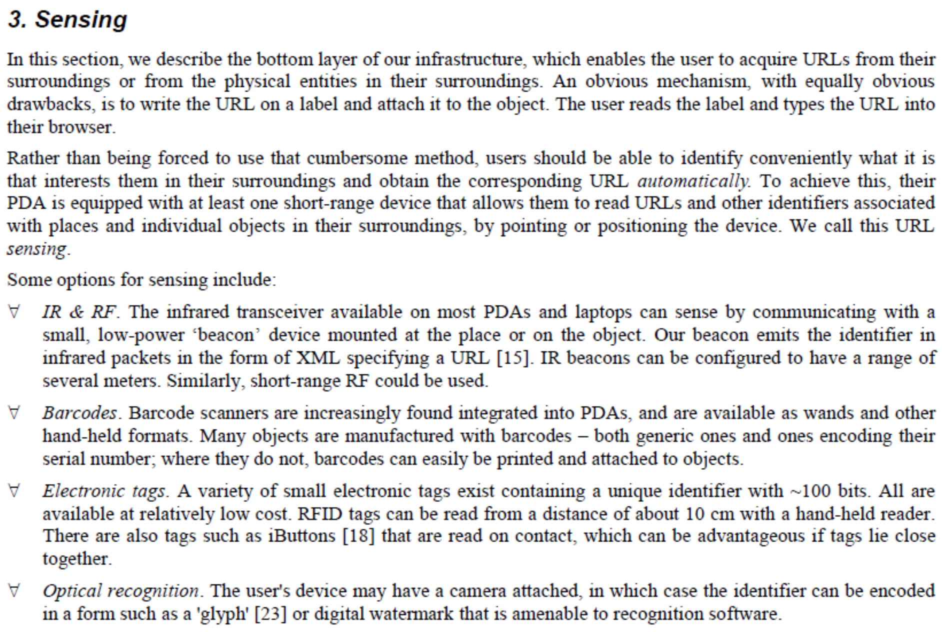 ***3. Sensing*** In this section, we describe the bottom layer of our infrastructure, which enables the user to acquire URLs from their surroundings or from the physical entities in their surroundings. An obvious mechanism, with equally obvious drawbacks, is to write the URL on a label and attach it to the object. The user reads the label and types the URL into their browser. Rather than being forced to use that cumbersome method, users should be able to identify conveniently what it is that interests them in their surroundings and obtain the corresponding URL *automatically*. To achieve this, their PDA is equipped with at least one short-range device that allows them to read URLs and other identifiers associated with places and individual objects in their surroundings, by pointing or positioning the device. We call this URL *sensing*. Some options for sensing include: ∀ *IR & RF*. The infrared transceiver available on most PDAs and laptops can sense by communicating with a small, low-power 'beacon' device mounted at the place or on the object. Our beacon emits the identifier in infrared packets in the form of XML specifying a URL [15]. IR beacons can be configured to have a range of several meters. Similarly, short-range RF could be used. ∀ *Barcodes*. Barcode scanners are increasingly found integrated into PDAs, and are available as wands and other hand-held formats. Many objects are manufactured with barcodes – both generic ones and ones encoding their serial number; where they do not, barcodes can easily be printed and attached to objects. ∀ *Electronic tags*. A variety of small electronic tags exist containing a unique identifier with ~100 bits. All are available at relatively low cost. RFID tags can be read from a distance of about 10 cm with a hand-held reader. There are also tags such as iButtons [18] that are read on contact, which can be advantageous if tags lie close together. ∀ *Optical recognition*. The user's device may have a camera attached, in which case the identifier can be encoded in a form such as a 'glyph' [23] or digital watermark that is amenable to recognition software. A Web-Based Nomadic Computing System (HP000074 – 87) at 4 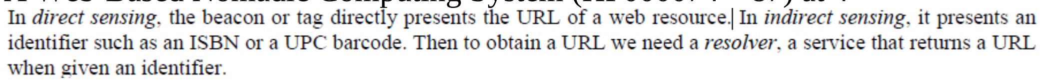 In *direct sensing*, the beacon or tag directly presents the URL of a web resource. In *indirect sensing*, it presents an identifier such as an ISBN or a UPC barcode. Then to obtain a URL we need a *resolver*, a service that returns a URL when given an identifier. A Web-Based Nomadic Computing System (HP000074 – 87) at 8 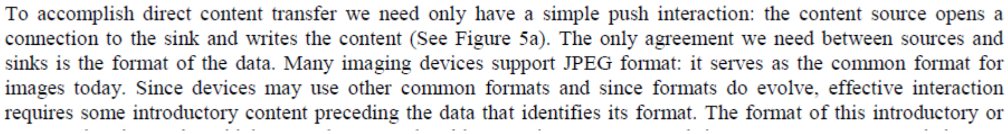 To accomplish direct content transfer we need only have a simple push interaction: the content source opens a connection to the sink and writes the content (See Figure 5a). The only agreement we need between sources and sinks is the format of the data. Many imaging devices support JPEG format: it serves as the common format for images today. Since devices may use other common formats and since formats do evolve, effective interaction requires some introductory content preceding the data that identifies its format. The format of this introductory or The Cooltown: People, Places, Things Web Presence for the Real World (HP000025 – 34) at 20:

Element	'529 Claim Recitation	Exemplary Citations in Reference
		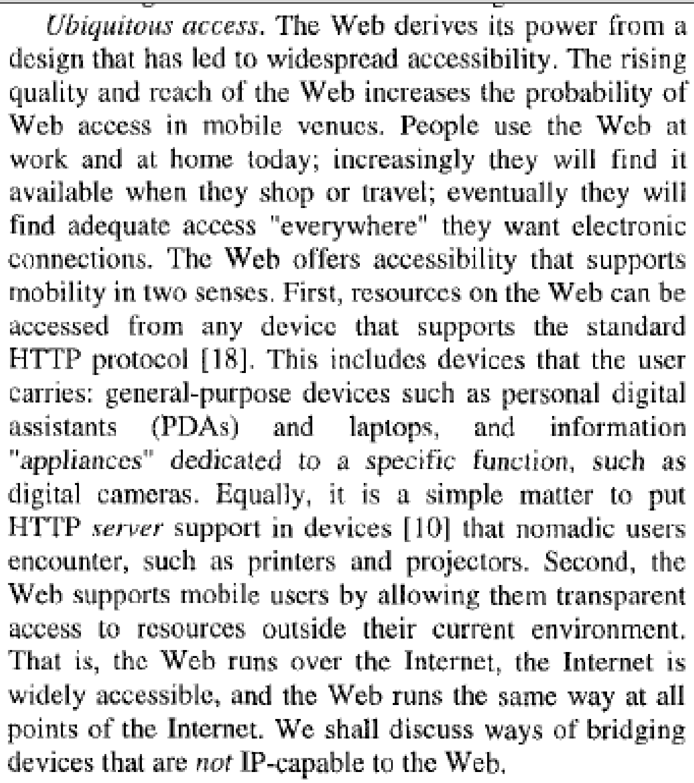 *Ubiquitous access*. The Web derives its power from a design that has led to widespread accessibility. The rising quality and reach of the Web increases the probability of Web access in mobile venues. People use the Web at work and at home today; increasingly they will find it available when they shop or travel; eventually they will find adequate access "everywhere" they want electronic connections. The Web offers accessibility that supports mobility in two senses. First, resources on the Web can be accessed from any device that supports the standard HTTP protocol [18]. This includes devices that the user carries: general-purpose devices such as personal digital assistants (PDAs) and laptops, and information "appliances" dedicated to a specific function, such as digital cameras. Equally, it is a simple matter to put HTTP *server* support in devices [10] that nomadic users encounter, such as printers and projectors. Second, the Web supports mobile users by allowing them transparent access to resources outside their current environment. That is, the Web runs over the Internet, the Internet is widely accessible, and the Web runs the same way at all points of the Internet. We shall discuss ways of bridging devices that are *not* IP-capable to the Web. To the extent it is found that CoolTown does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that CoolTown does not anticipate this claim, CoolTown renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine CoolTown with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success

Element	**'529 Claim Recitation**	**Exemplary Citations in Reference**
		for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1C]	the service programmed to receive the data;	CoolTown discloses the service programmed to receive the data. A Web-Based Nomadic Computing System (HP000074 – 87) at 4 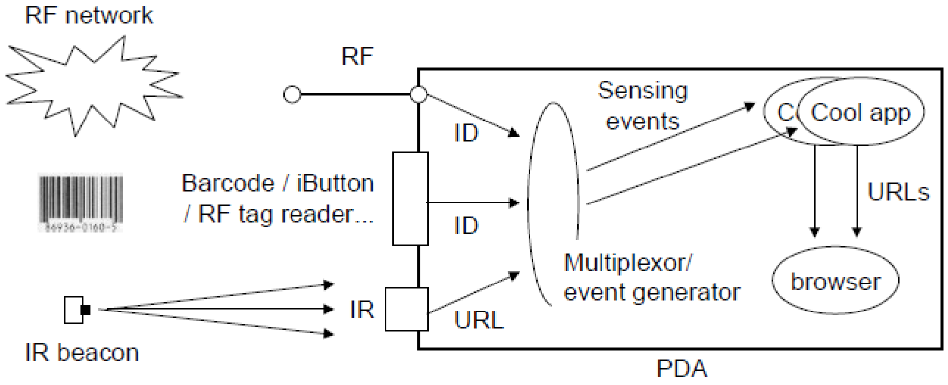 Figure 3. A sketch of a PDA equipped for sensing URLs and identifiers.

Element	'529 Claim Recitation	Exemplary Citations in Reference
		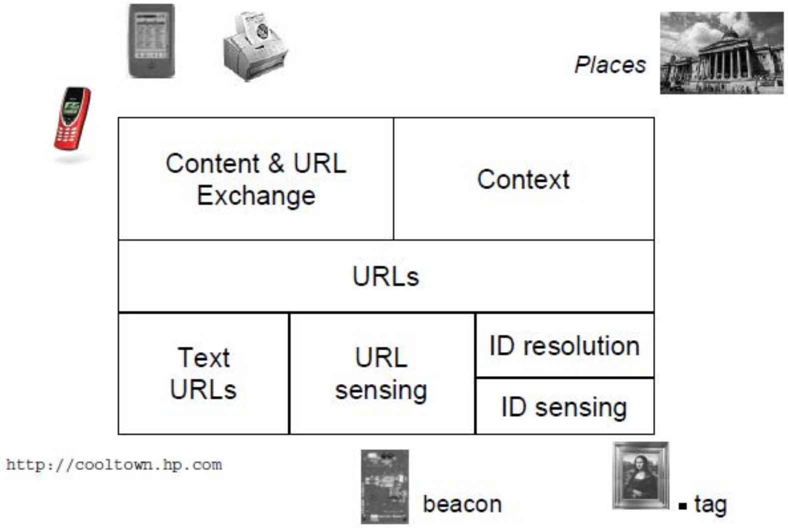 Figure 2. Component relationships. Sources of URLs, including keyboards, beacons, and resolved tags point to context or content. A Web-Based Nomadic Computing System (HP000074 – 87) at 4 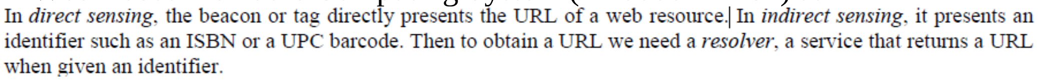 In *direct sensing*, the beacon or tag directly presents the URL of a web resource. In *indirect sensing*, it presents an identifier such as an ISBN or a UPC barcode. Then to obtain a URL we need a *resolver*, a service that returns a URL when given an identifier. A Web-Based Nomadic Computing System (HP000074 – 87) at 5 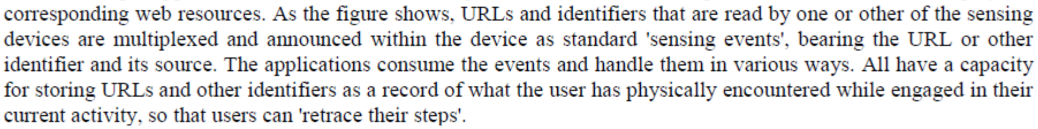 Figure 3 sketches a PDA with several sensing mechanisms (IR, barcode and RF) similar to prototypes we have developed in the CoolTown project. The device runs a web browser, and one or more CoolTown applications. The latter are simple applications that obtain URLs and present them as links, using the standard browser to display the corresponding web resources. As the figure shows, URLs and identifiers that are read by one or other of the sensing devices are multiplexed and announced within the device as standard 'sensing events', bearing the URL or other identifier and its source. The applications consume the events and handle them in various ways. All have a capacity for storing URLs and other identifiers as a record of what the user has physically encountered while engaged in their current activity, so that users can 'retrace their steps'.

Element	'529 Claim Recitation	Exemplary Citations in Reference
		To the extent it is found that CoolTown does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that CoolTown does not anticipate this claim, CoolTown renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine CoolTown with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1D]	identify an object within the image:	CoolTown discloses identify an object within the image. A Web-Based Nomadic Computing System (HP000074 – 87) at 3 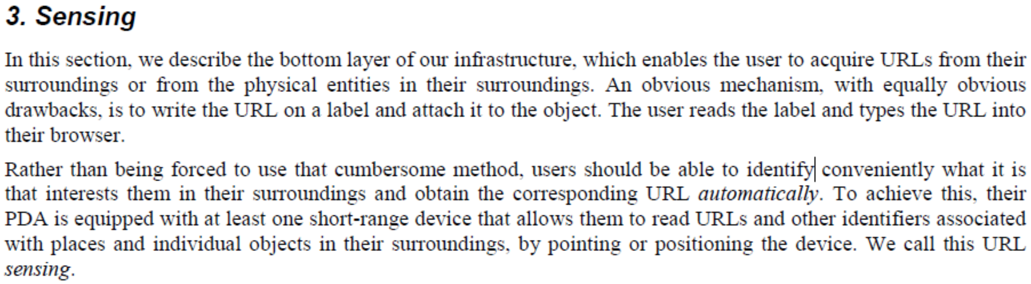 **3. Sensing** In this section, we describe the bottom layer of our infrastructure, which enables the user to acquire URLs from their surroundings or from the physical entities in their surroundings. An obvious mechanism, with equally obvious drawbacks, is to write the URL on a label and attach it to the object. The user reads the label and types the URL into their browser. Rather than being forced to use that cumbersome method, users should be able to identify conveniently what it is that interests them in their surroundings and obtain the corresponding URL *automatically*. To achieve this, their PDA is equipped with at least one short-range device that allows them to read URLs and other identifiers associated with places and individual objects in their surroundings, by pointing or positioning the device. We call this URL *sensing*. A Web-Based Nomadic Computing System (HP000074 – 87) at 4

Element	'529 Claim Recitation	Exemplary Citations in Reference
		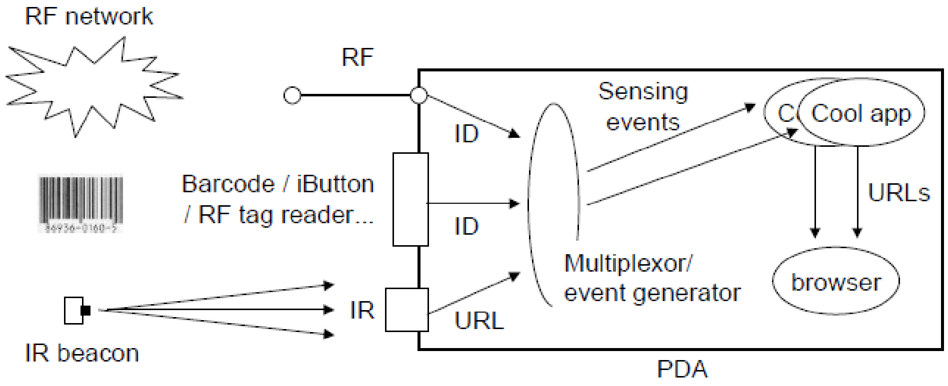 Figure 3. A sketch of a PDA equipped for sensing URLs and identifiers. Cooltown Developer's Network, People, Place, and Things: web presence for the real world (HP000025–34) (HP000025–34) at 23 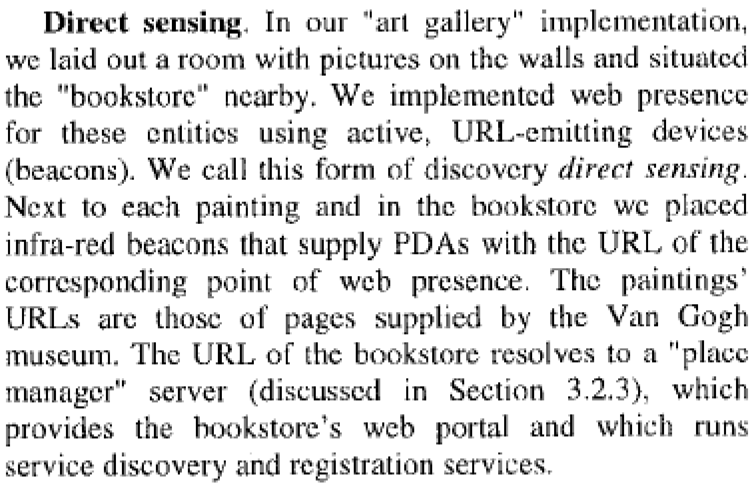 **Direct sensing**. In our "art gallery" implementation, we laid out a room with pictures on the walls and situated the "bookstore" nearby. We implemented web presence for these entities using active, URL-emitting devices (beacons). We call this form of discovery *direct sensing*. Next to each painting and in the bookstore we placed infra-red beacons that supply PDAs with the URL of the corresponding point of web presence. The paintings' URLs are those of pages supplied by the Van Gogh museum. The URL of the bookstore resolves to a "place manager" server (discussed in Section 3.2.3), which provides the bookstore's web portal and which runs service discovery and registration services.

Element	'529 Claim Recitation	Exemplary Citations in Reference
		Indirect sensing. In the case of indirect sensing, the user's client device obtains a value that it has to look up to obtain a URL. There are many types of value that can be 'sensed' as the user travels through the physical world. For example, the device could obtain latitude and longitude coordinates from a GPS sensor, which could then be translated into a zip code and hence into the URL of a place's web portal provided by Yahoo, for example. The device could even perform image recognition against a database of tourist sites to provide the URL of a corresponding web page. We have been experimenting with indirect sensing where the sensed value is obtained from a *tag*. There are electronic tags like RFID tags [33] and iButtons [22] – small devices that supply a unique identification string to a sensor placed near them or in contact with them. 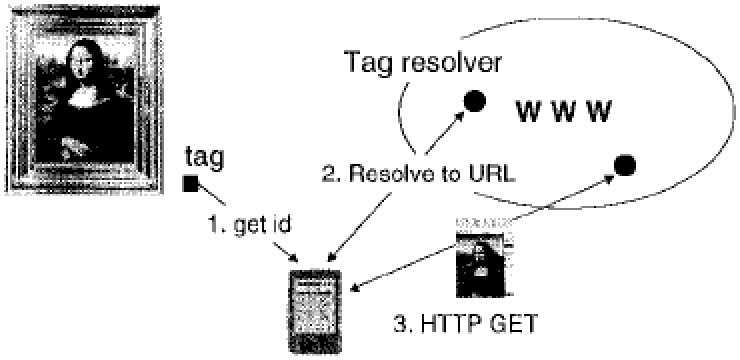 **Figure 4. Indirect sensing: the web presence for a painting.**

Element	'529 Claim Recitation	Exemplary Citations in Reference
		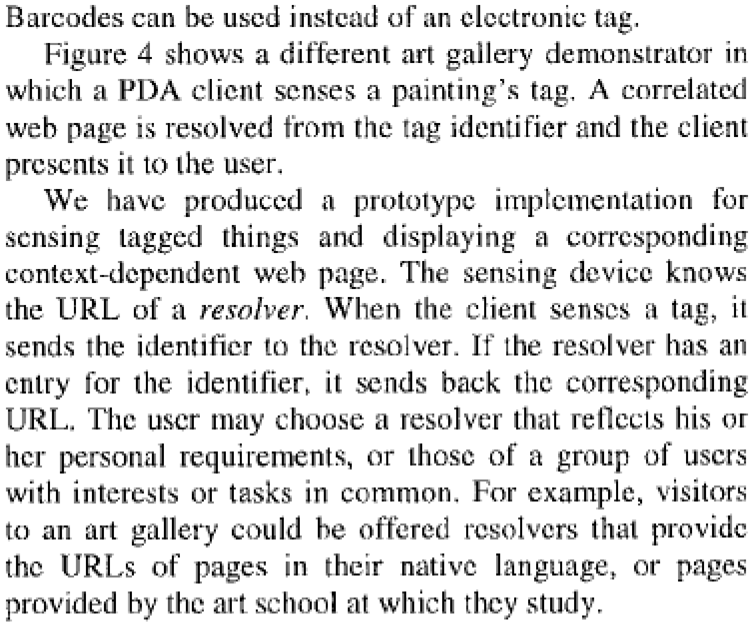Barcodes can be used instead of an electronic tag. Figure 4 shows a different art gallery demonstrator in which a PDA client senses a painting's tag. A correlated web page is resolved from the tag identifier and the client presents it to the user. We have produced a prototype implementation for sensing tagged things and displaying a corresponding context-dependent web page. The sensing device knows the URL of a *resolver*. When the client senses a tag, it sends the identifier to the resolver. If the resolver has an entry for the identifier, it sends back the corresponding URL. The user may choose a resolver that reflects his or her personal requirements, or those of a group of users with interests or tasks in common. For example, visitors to an art gallery could be offered resolvers that provide the URLs of pages in their native language, or pages provided by the art school at which they study. To the extent it is found that CoolTown does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that CoolTown does not anticipate this claim, CoolTown renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine CoolTown with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

Element	'529 Claim Recitation	Exemplary Citations in Reference
[1E]	distinguish an object present in the image from others using a database that stores data characteristics of target objects;	CoolTown discloses distinguish an object present in the image from others using a database that stores data characteristics of target objects. A Web-Based Nomadic Computing System (HP000074 – 87) at Fig. 2 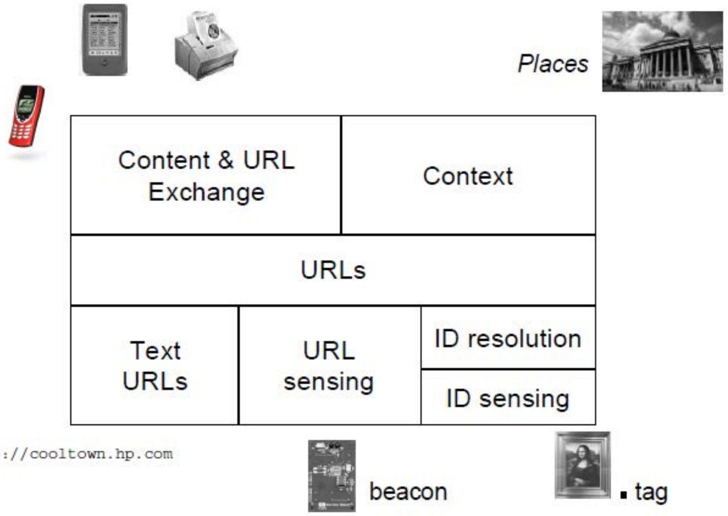 Figure 2. Component relationships. Sources of URLs, including keyboards, beacons, and resolved tags point to context or content. A Web-Based Nomadic Computing System (HP000074 – 87) at 4

Element	'529 Claim Recitation	Exemplary Citations in Reference
		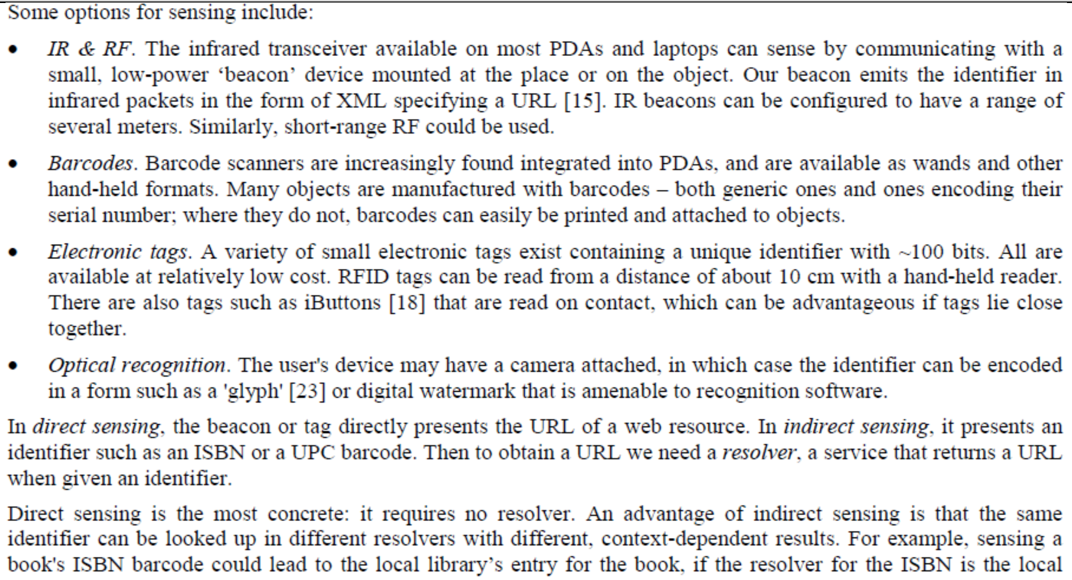 Some options for sensing include: • *IR & RF*. The infrared transceiver available on most PDAs and laptops can sense by communicating with a small, low-power 'beacon' device mounted at the place or on the object. Our beacon emits the identifier in infrared packets in the form of XML specifying a URL [15]. IR beacons can be configured to have a range of several meters. Similarly, short-range RF could be used. • *Barcodes*. Barcode scanners are increasingly found integrated into PDAs, and are available as wands and other hand-held formats. Many objects are manufactured with barcodes – both generic ones and ones encoding their serial number; where they do not, barcodes can easily be printed and attached to objects. • *Electronic tags*. A variety of small electronic tags exist containing a unique identifier with ~100 bits. All are available at relatively low cost. RFID tags can be read from a distance of about 10 cm with a hand-held reader. There are also tags such as iButtons [18] that are read on contact, which can be advantageous if tags lie close together. • *Optical recognition*. The user's device may have a camera attached, in which case the identifier can be encoded in a form such as a 'glyph' [23] or digital watermark that is amenable to recognition software. In *direct sensing*, the beacon or tag directly presents the URL of a web resource. In *indirect sensing*, it presents an identifier such as an ISBN or a UPC barcode. Then to obtain a URL we need a *resolver*, a service that returns a URL when given an identifier. Direct sensing is the most concrete: it requires no resolver. An advantage of indirect sensing is that the same identifier can be looked up in different resolvers with different, context-dependent results. For example, sensing a book's ISBN barcode could lead to the local library's entry for the book, if the resolver for the ISBN is the local A Web-Based Nomadic Computing System (HP000074 – 87) at 5 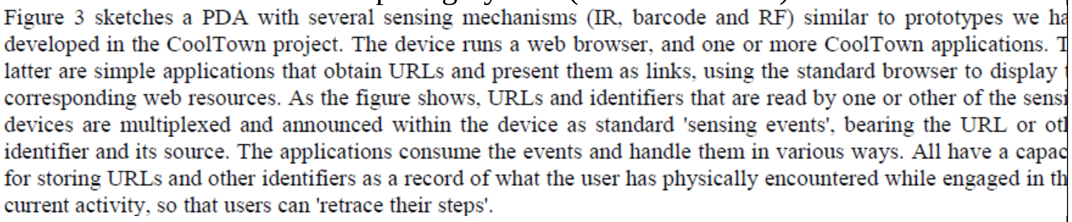 Figure 3 sketches a PDA with several sensing mechanisms (IR, barcode and RF) similar to prototypes we ha developed in the CoolTown project. The device runs a web browser, and one or more CoolTown applications. T latter are simple applications that obtain URLs and present them as links, using the standard browser to display corresponding web resources. As the figure shows, URLs and identifiers that are read by one or other of the sensi devices are multiplexed and announced within the device as standard 'sensing events', bearing the URL or ot identifier and its source. The applications consume the events and handle them in various ways. All have a capac for storing URLs and other identifiers as a record of what the user has physically encountered while engaged in th current activity, so that users can 'retrace their steps'. A Web-Based Nomadic Computing System (HP000074 – 87) at 6

Element	'529 Claim Recitation	Exemplary Citations in Reference
		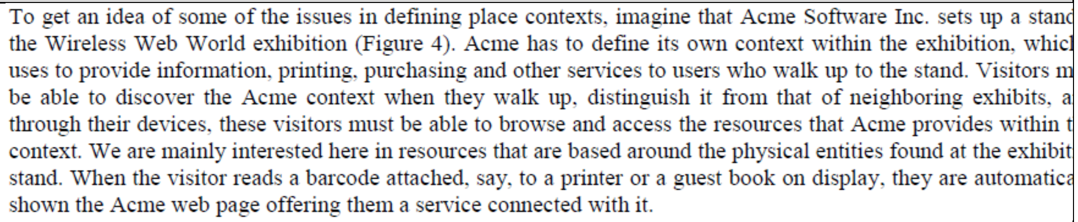 To get an idea of some of the issues in defining place contexts, imagine that Acme Software Inc. sets up a stan the Wireless Web World exhibition (Figure 4). Acme has to define its own context within the exhibition, whic uses to provide information, printing, purchasing and other services to users who walk up to the stand. Visitors m be able to discover the Acme context when they walk up, distinguish it from that of neighboring exhibits, a through their devices, these visitors must be able to browse and access the resources that Acme provides within t context. We are mainly interested here in resources that are based around the physical entities found at the exhibit stand. When the visitor reads a barcode attached, say, to a printer or a guest book on display, they are automatica shown the Acme web page offering them a service connected with it. A Web-Based Nomadic Computing System (HP000074 – 87) at 6 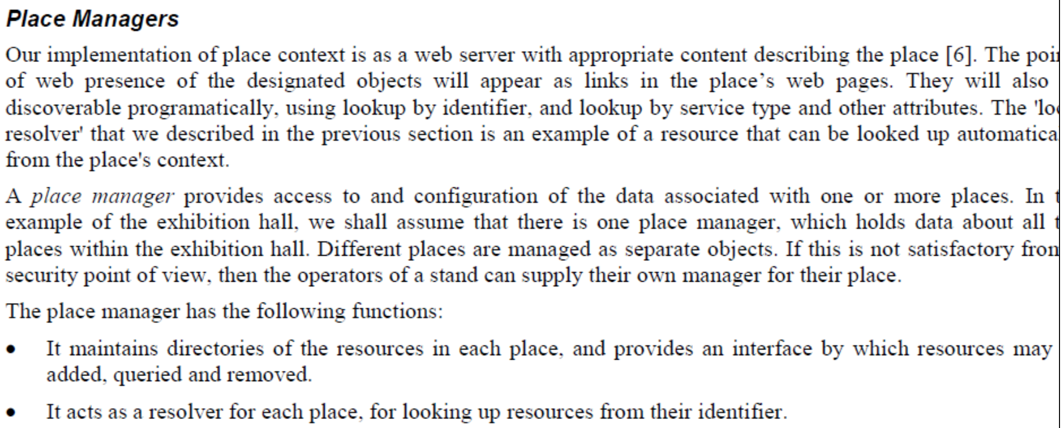 ***Place Managers*** Our implementation of place context is as a web server with appropriate content describing the place [6]. The poi of web presence of the designated objects will appear as links in the place's web pages. They will also discoverable programatically, using lookup by identifier, and lookup by service type and other attributes. The 'lo resolver' that we described in the previous section is an example of a resource that can be looked up automatica from the place's context. A *place manager* provides access to and configuration of the data associated with one or more places. In t example of the exhibition hall, we shall assume that there is one place manager, which holds data about all t places within the exhibition hall. Different places are managed as separate objects. If this is not satisfactory fron security point of view, then the operators of a stand can supply their own manager for their place. The place manager has the following functions: • It maintains directories of the resources in each place, and provides an interface by which resources may added, queried and removed. • It acts as a resolver for each place, for looking up resources from their identifier. Cooltown Developer's Network, People, Place, and Things: web presence for the real world (HP000025–34) (HP000025–34) at 8 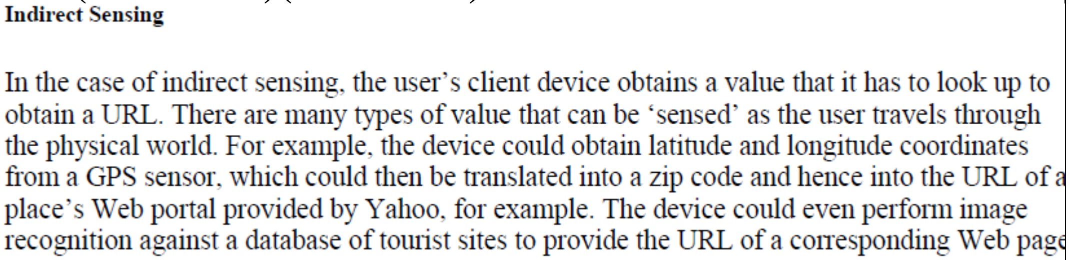 **Indirect Sensing** In the case of indirect sensing, the user's client device obtains a value that it has to look up to obtain a URL. There are many types of value that can be 'sensed' as the user travels through the physical world. For example, the device could obtain latitude and longitude coordinates from a GPS sensor, which could then be translated into a zip code and hence into the URL of a place's Web portal provided by Yahoo, for example. The device could even perform image recognition against a database of tourist sites to provide the URL of a corresponding Web page Cooltown Developer's Network, People, Place, and Things: web presence for the real world (HP000025–34) (HP000025–34) at 23

Element	'529 Claim Recitation	Exemplary Citations in Reference
		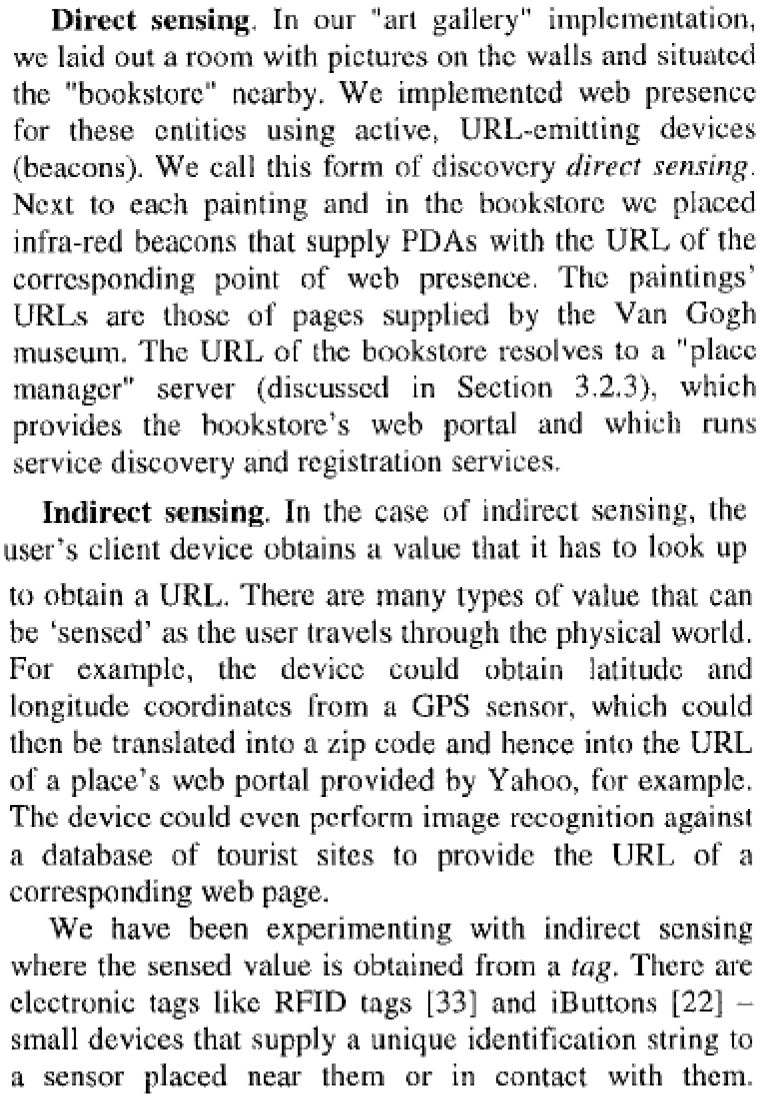 **Direct sensing**. In our "art gallery" implementation, we laid out a room with pictures on the walls and situated the "bookstore" nearby. We implemented web presence for these entities using active, URL-emitting devices (beacons). We call this form of discovery *direct sensing*. Next to each painting and in the bookstore we placed infra-red beacons that supply PDAs with the URL of the corresponding point of web presence. The paintings' URLs are those of pages supplied by the Van Gogh museum. The URL of the bookstore resolves to a "place manager" server (discussed in Section 3.2.3), which provides the bookstore's web portal and which runs service discovery and registration services. **Indirect sensing**. In the case of indirect sensing, the user's client device obtains a value that it has to look up to obtain a URL. There are many types of value that can be 'sensed' as the user travels through the physical world. For example, the device could obtain latitude and longitude coordinates from a GPS sensor, which could then be translated into a zip code and hence into the URL of a place's web portal provided by Yahoo, for example. The device could even perform image recognition against a database of tourist sites to provide the URL of a corresponding web page. We have been experimenting with indirect sensing where the sensed value is obtained from a *tag*. There are electronic tags like RFID tags [33] and iButtons [22] – small devices that supply a unique identification string to a sensor placed near them or in contact with them.

Element	'529 Claim Recitation	Exemplary Citations in Reference
		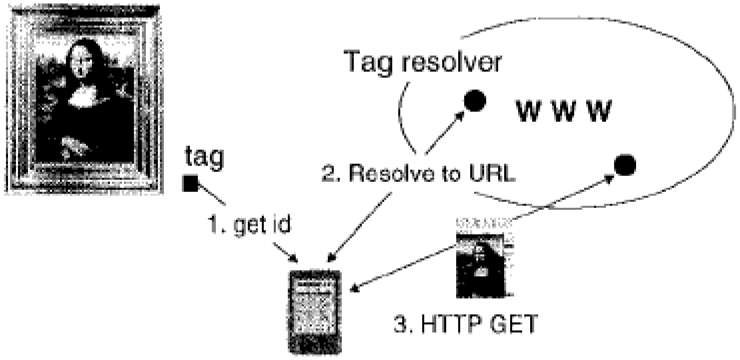 **Figure 4. Indirect sensing: the web presence for a painting.** Barcodes can be used instead of an electronic tag. Figure 4 shows a different art gallery demonstrator in which a PDA client senses a painting's tag. A correlated web page is resolved from the tag identifier and the client presents it to the user. We have produced a prototype implementation for sensing tagged things and displaying a corresponding context-dependent web page. The sensing device knows the URL of a *resolver*. When the client senses a tag, it sends the identifier to the resolver. If the resolver has an entry for the identifier, it sends back the corresponding URL. The user may choose a resolver that reflects his or her personal requirements, or those of a group of users with interests or tasks in common. For example, visitors to an art gallery could be offered resolvers that provide the URLs of pages in their native language, or pages provided by the art school at which they study.

Element	'529 Claim Recitation	Exemplary Citations in Reference
		To the extent it is found that CoolTown does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that CoolTown does not anticipate this claim, CoolTown renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine CoolTown with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1F]	associate the object with information; and	CoolTown discloses associate the object with information. A Web-Based Nomadic Computing System (HP000074 – 87) at Fig. 2

Element	'529 Claim Recitation	Exemplary Citations in Reference
		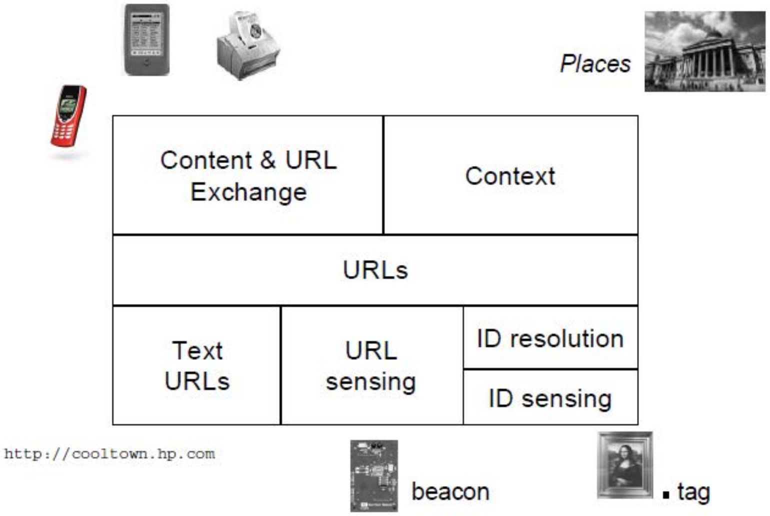 Figure 2. Component relationships. Sources of URLs, including keyboards, beacons, and resolved tags point to context or content. A Web-Based Nomadic Computing System (HP000074 – 87) at 3 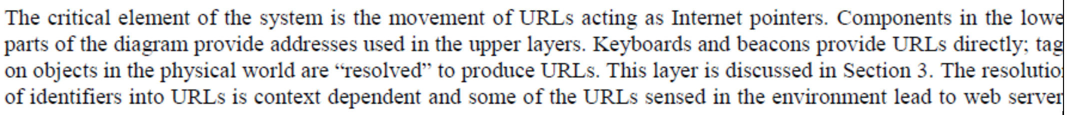 The critical element of the system is the movement of URLs acting as Internet pointers. Components in the lowe parts of the diagram provide addresses used in the upper layers. Keyboards and beacons provide URLs directly; tag on objects in the physical world are "resolved" to produce URLs. This layer is discussed in Section 3. The resolutio of identifiers into URLs is context dependent and some of the URLs sensed in the environment lead to web server A Web-Based Nomadic Computing System (HP000074 – 87) at 4

Element	'529 Claim Recitation	Exemplary Citations in Reference
		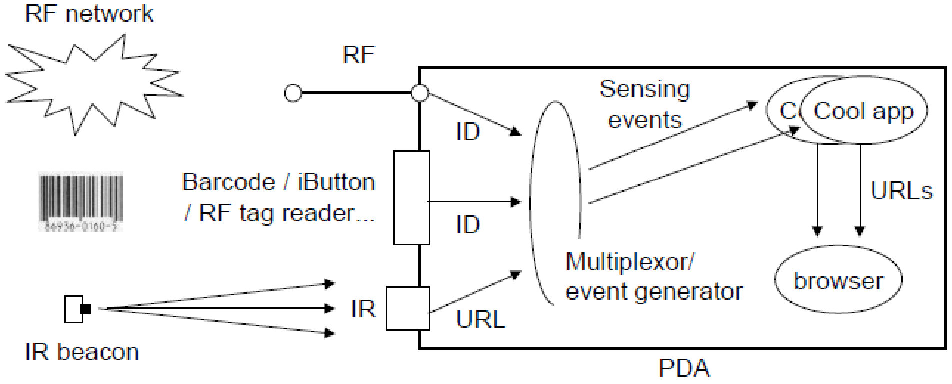 Figure 3. A sketch of a PDA equipped for sensing URLs and identifiers. A Web-Based Nomadic Computing System (HP000074 – 87) at 4–5 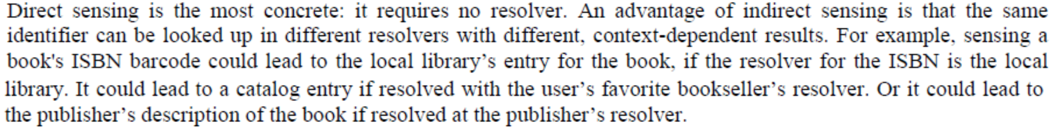 Direct sensing is the most concrete: it requires no resolver. An advantage of indirect sensing is that the same identifier can be looked up in different resolvers with different, context-dependent results. For example, sensing a book's ISBN barcode could lead to the local library's entry for the book, if the resolver for the ISBN is the local library. It could lead to a catalog entry if resolved with the user's favorite bookseller's resolver. Or it could lead to the publisher's description of the book if resolved at the publisher's resolver. A Web-Based Nomadic Computing System (HP000074 – 87) at 5 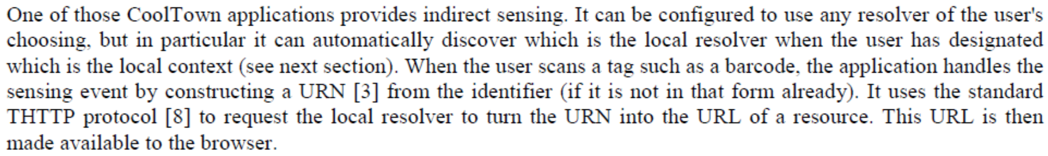 One of those CoolTown applications provides indirect sensing. It can be configured to use any resolver of the user's choosing, but in particular it can automatically discover which is the local resolver when the user has designated which is the local context (see next section). When the user scans a tag such as a barcode, the application handles the sensing event by constructing a URN [3] from the identifier (if it is not in that form already). It uses the standard THTTP protocol [8] to request the local resolver to turn the URN into the URL of a resource. This URL is then made available to the browser. A Web-Based Nomadic Computing System (HP000074 – 87) at 6

Element	'529 Claim Recitation	Exemplary Citations in Reference
		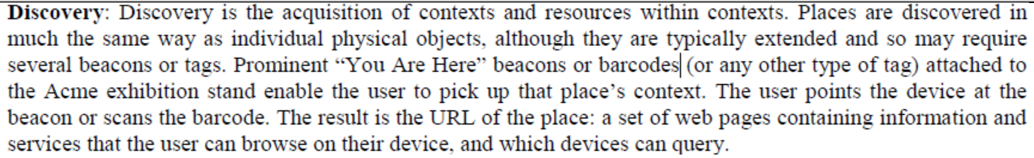 **Discovery**: Discovery is the acquisition of contexts and resources within contexts. Places are discovered in much the same way as individual physical objects, although they are typically extended and so may require several beacons or tags. Prominent "You Are Here" beacons or barcodes (or any other type of tag) attached to the Acme exhibition stand enable the user to pick up that place's context. The user points the device at the beacon or scans the barcode. The result is the URL of the place: a set of web pages containing information and services that the user can browse on their device, and which devices can query. A Web-Based Nomadic Computing System (HP000074 – 87) at 6 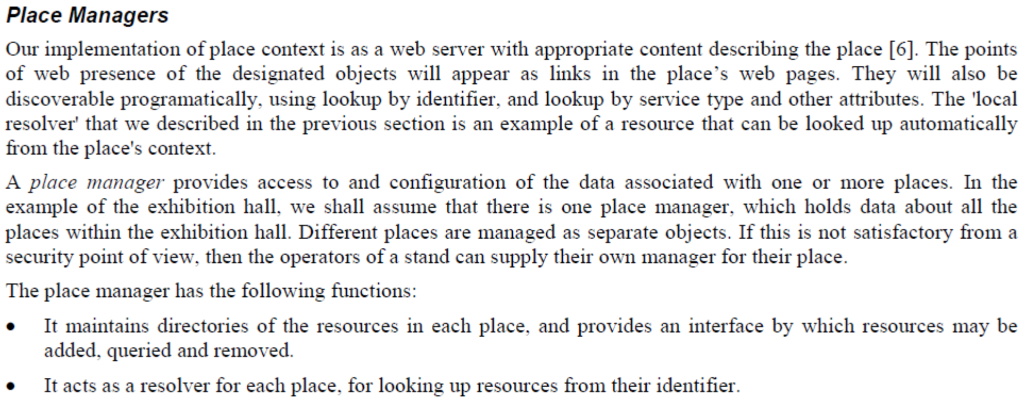 ***Place Managers*** Our implementation of place context is as a web server with appropriate content describing the place [6]. The points of web presence of the designated objects will appear as links in the place's web pages. They will also be discoverable programatically, using lookup by identifier, and lookup by service type and other attributes. The 'local resolver' that we described in the previous section is an example of a resource that can be looked up automatically from the place's context. A *place manager* provides access to and configuration of the data associated with one or more places. In the example of the exhibition hall, we shall assume that there is one place manager, which holds data about all the places within the exhibition hall. Different places are managed as separate objects. If this is not satisfactory from a security point of view, then the operators of a stand can supply their own manager for their place. The place manager has the following functions: • It maintains directories of the resources in each place, and provides an interface by which resources may be added, queried and removed. • It acts as a resolver for each place, for looking up resources from their identifier. Cooltown Developer's Network, People, Place, and Things: web presence for the real world (HP000025–34) (HP000025–34) at 20 – 21 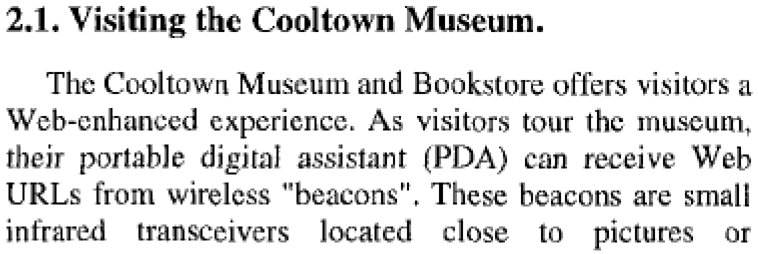 **2.1. Visiting the Cooltown Museum.** The Cooltown Museum and Bookstore offers visitors a Web-enhanced experience. As visitors tour the museum, their portable digital assistant (PDA) can receive Web URLs from wireless "beacons". These beacons are small infrared transceivers located close to pictures or

Element	'529 Claim Recitation	Exemplary Citations in Reference
		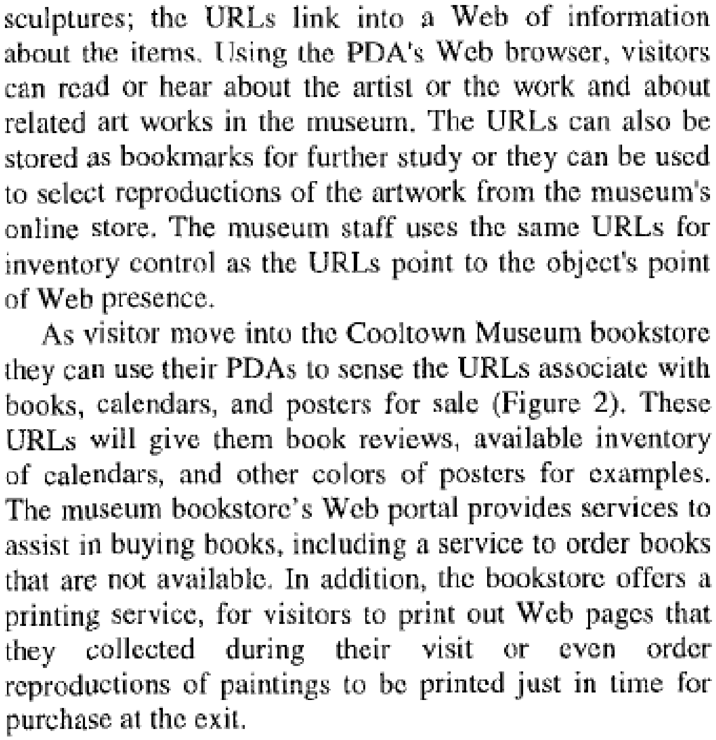 sculptures; the URLs link into a Web of information about the items. Using the PDA's Web browser, visitors can read or hear about the artist or the work and about related art works in the museum. The URLs can also be stored as bookmarks for further study or they can be used to select reproductions of the artwork from the museum's online store. The museum staff uses the same URLs for inventory control as the URLs point to the object's point of Web presence. As visitor move into the Cooltown Museum bookstore they can use their PDAs to sense the URLs associate with books, calendars, and posters for sale (Figure 2). These URLs will give them book reviews, available inventory of calendars, and other colors of posters for examples. The museum bookstore's Web portal provides services to assist in buying books, including a service to order books that are not available. In addition, the bookstore offers a printing service, for visitors to print out Web pages that they collected during their visit or even order reproductions of paintings to be printed just in time for purchase at the exit. Cooltown Developer's Network, People, Place, and Things: web presence for the real world (HP000025–34) (HP000025–34) at 23

Element	'529 Claim Recitation	Exemplary Citations in Reference
		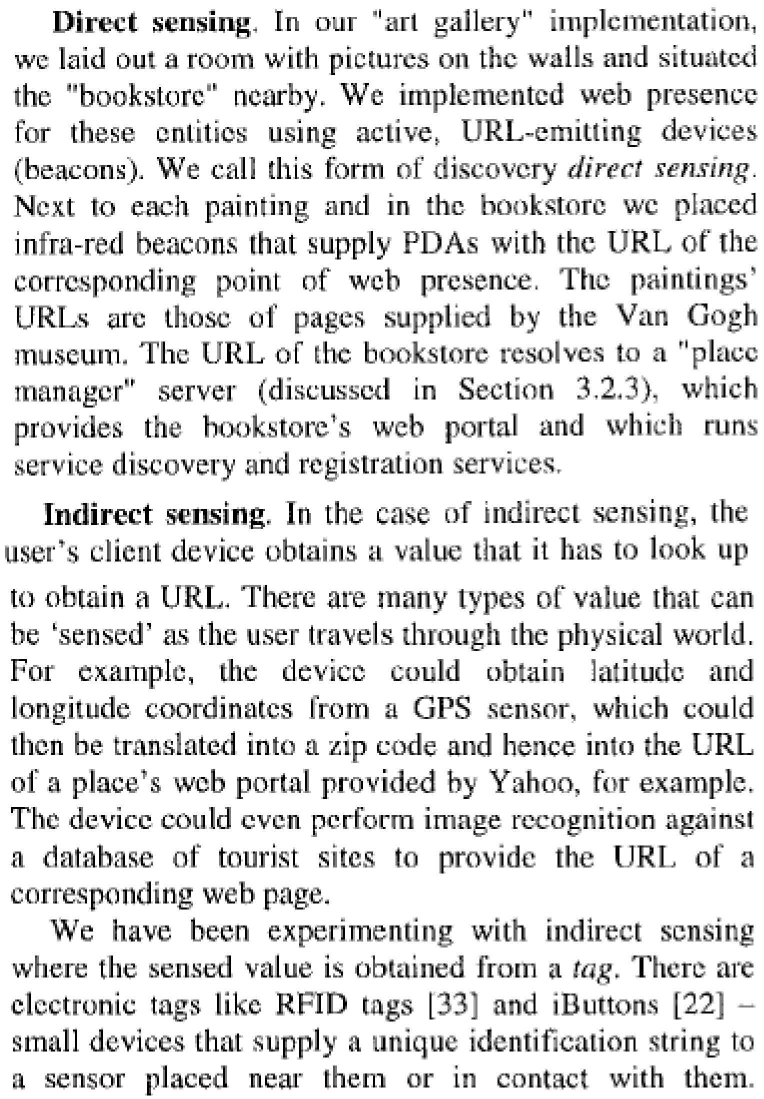 **Direct sensing**. In our "art gallery" implementation, we laid out a room with pictures on the walls and situated the "bookstore" nearby. We implemented web presence for these entities using active, URL-emitting devices (beacons). We call this form of discovery *direct sensing*. Next to each painting and in the bookstore we placed infra-red beacons that supply PDAs with the URL of the corresponding point of web presence. The paintings' URLs are those of pages supplied by the Van Gogh museum. The URL of the bookstore resolves to a "place manager" server (discussed in Section 3.2.3), which provides the bookstore's web portal and which runs service discovery and registration services. **Indirect sensing**. In the case of indirect sensing, the user's client device obtains a value that it has to look up to obtain a URL. There are many types of value that can be 'sensed' as the user travels through the physical world. For example, the device could obtain latitude and longitude coordinates from a GPS sensor, which could then be translated into a zip code and hence into the URL of a place's web portal provided by Yahoo, for example. The device could even perform image recognition against a database of tourist sites to provide the URL of a corresponding web page. We have been experimenting with indirect sensing where the sensed value is obtained from a *tag*. There are electronic tags like RFID tags [33] and iButtons [22] – small devices that supply a unique identification string to a sensor placed near them or in contact with them.

Element	'529 Claim Recitation	Exemplary Citations in Reference
		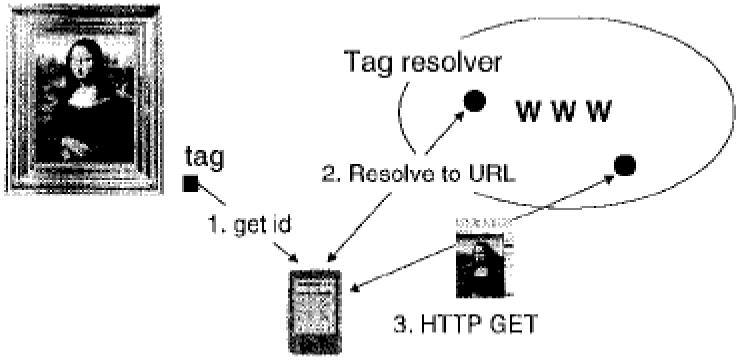 **Figure 4. Indirect sensing: the web presence for a painting.** Barcodes can be used instead of an electronic tag. Figure 4 shows a different art gallery demonstrator in which a PDA client senses a painting's tag. A correlated web page is resolved from the tag identifier and the client presents it to the user. We have produced a prototype implementation for sensing tagged things and displaying a corresponding context-dependent web page. The sensing device knows the URL of a *resolver*. When the client senses a tag, it sends the identifier to the resolver. If the resolver has an entry for the identifier, it sends back the corresponding URL. The user may choose a resolver that reflects his or her personal requirements, or those of a group of users with interests or tasks in common. For example, visitors to an art gallery could be offered resolvers that provide the URLs of pages in their native language, or pages provided by the art school at which they study.

Element	'529 Claim Recitation	Exemplary Citations in Reference
		A Web-Based Nomadic Computing System (HP000074 – 87) at 6 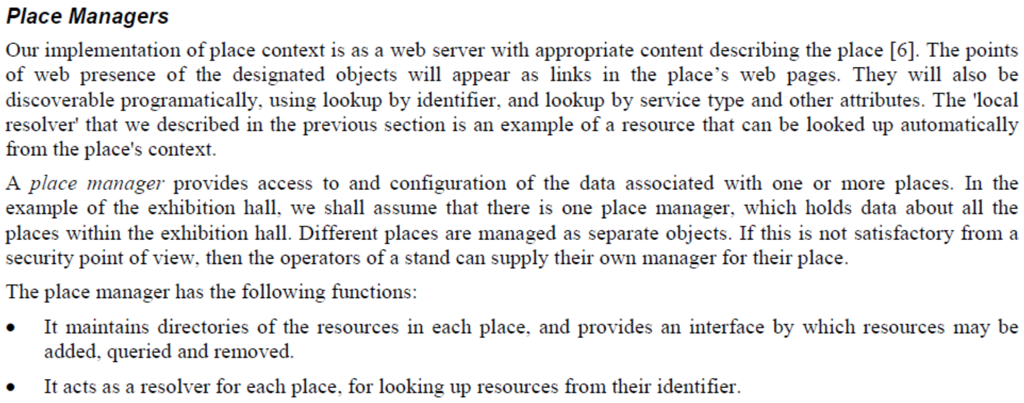 ***Place Managers*** Our implementation of place context is as a web server with appropriate content describing the place [6]. The points of web presence of the designated objects will appear as links in the place's web pages. They will also be discoverable programatically, using lookup by identifier, and lookup by service type and other attributes. The 'local resolver' that we described in the previous section is an example of a resource that can be looked up automatically from the place's context. A *place manager* provides access to and configuration of the data associated with one or more places. In the example of the exhibition hall, we shall assume that there is one place manager, which holds data about all the places within the exhibition hall. Different places are managed as separate objects. If this is not satisfactory from a security point of view, then the operators of a stand can supply their own manager for their place. The place manager has the following functions: • It maintains directories of the resources in each place, and provides an interface by which resources may be added, queried and removed. • It acts as a resolver for each place, for looking up resources from their identifier. To the extent it is found that CoolTown does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that CoolTown does not anticipate this claim, CoolTown renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine CoolTown with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1G]	return the information to the network-enabled device; and	CoolTown discloses return the information to the network-enabled device. A Web-Based Nomadic Computing System (HP000074 – 87) at Fig. 2

Element	’529 Claim Recitation	Exemplary Citations in Reference
		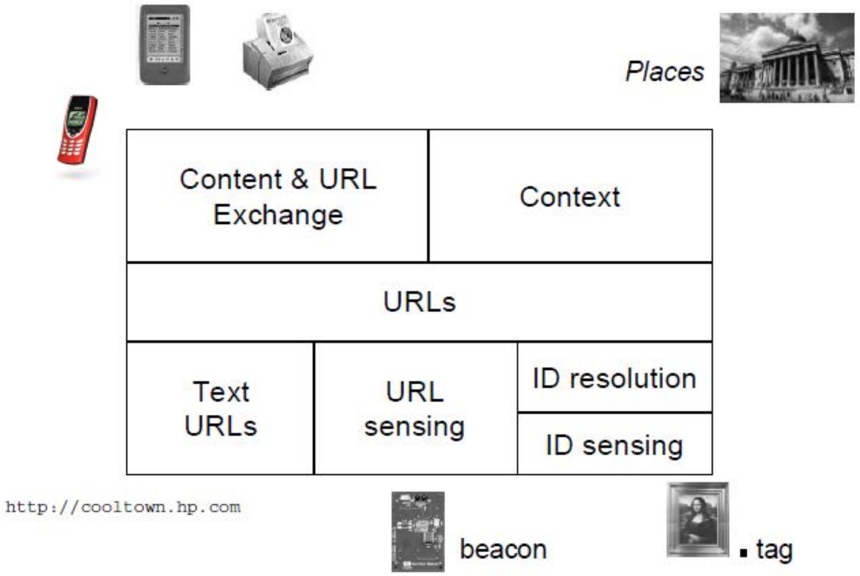 Figure 2. Component relationships. Sources of URLs, including keyboards, beacons, and resolve tags point to context or content. A Web-Based Nomadic Computing System (HP000074 – 87) at 4

Element	'529 Claim Recitation	Exemplary Citations in Reference
		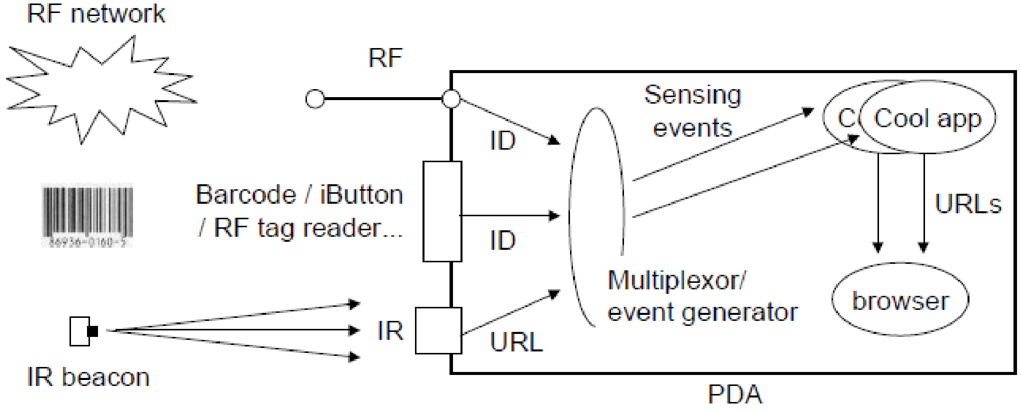 Figure 3. A sketch of a PDA equipped for sensing URLs and identifiers. A Web-Based Nomadic Computing System (HP000074 – 87) at 5 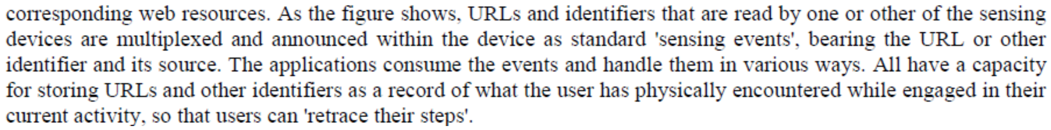 Figure 3 sketches a PDA with several sensing mechanisms (IR, barcode and RF) similar to prototypes we have developed in the CoolTown project. The device runs a web browser, and one or more CoolTown applications. The latter are simple applications that obtain URLs and present them as links, using the standard browser to display the corresponding web resources. As the figure shows, URLs and identifiers that are read by one or other of the sensing devices are multiplexed and announced within the device as standard 'sensing events', bearing the URL or other identifier and its source. The applications consume the events and handle them in various ways. All have a capacity for storing URLs and other identifiers as a record of what the user has physically encountered while engaged in their current activity, so that users can 'retrace their steps'. A Web-Based Nomadic Computing System (HP000074 – 87) at 5 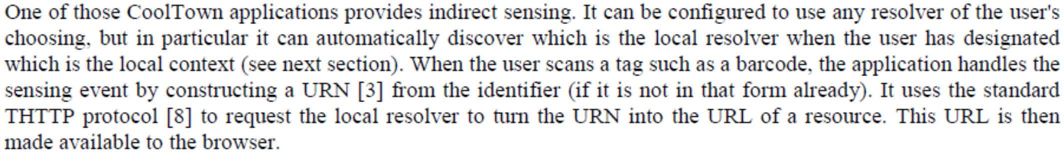 One of those CoolTown applications provides indirect sensing. It can be configured to use any resolver of the user's choosing, but in particular it can automatically discover which is the local resolver when the user has designated which is the local context (see next section). When the user scans a tag such as a barcode, the application handles the sensing event by constructing a URN [3] from the identifier (if it is not in that form already). It uses the standard THTTP protocol [8] to request the local resolver to turn the URN into the URL of a resource. This URL is then made available to the browser. A Web-Based Nomadic Computing System (HP000074 – 87) at 6

Element	'529 Claim Recitation	Exemplary Citations in Reference
		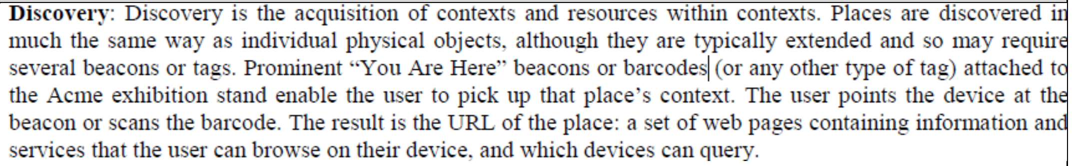 **Discovery**: Discovery is the acquisition of contexts and resources within contexts. Places are discovered in much the same way as individual physical objects, although they are typically extended and so may require several beacons or tags. Prominent "You Are Here" beacons or barcodes (or any other type of tag) attached to the Acme exhibition stand enable the user to pick up that place's context. The user points the device at the beacon or scans the barcode. The result is the URL of the place: a set of web pages containing information and services that the user can browse on their device, and which devices can query. Cooltown Developer's Network, People, Place, and Things: web presence for the real world (HP000025–34) at 5 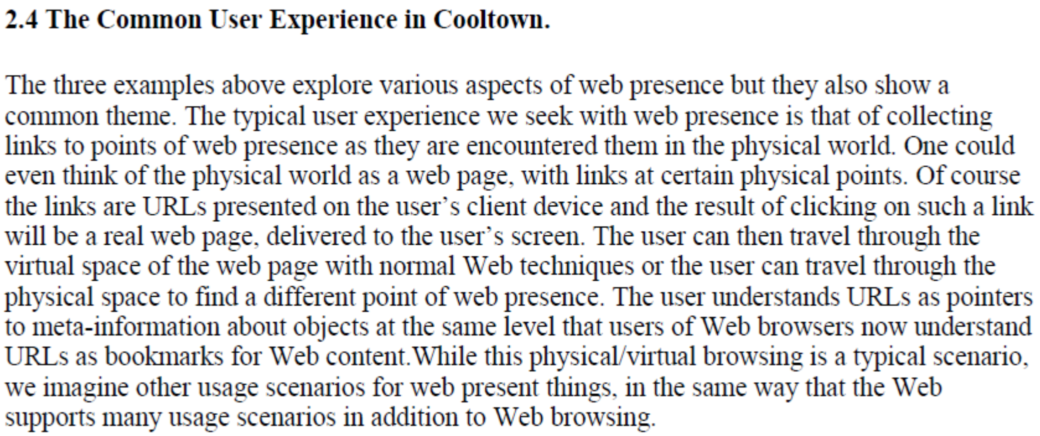 **2.4 The Common User Experience in Cooltown.** The three examples above explore various aspects of web presence but they also show a common theme. The typical user experience we seek with web presence is that of collecting links to points of web presence as they are encountered them in the physical world. One could even think of the physical world as a web page, with links at certain physical points. Of course the links are URLs presented on the user's client device and the result of clicking on such a link will be a real web page, delivered to the user's screen. The user can then travel through the virtual space of the web page with normal Web techniques or the user can travel through the physical space to find a different point of web presence. The user understands URLs as pointers to meta-information about objects at the same level that users of Web browsers now understand URLs as bookmarks for Web content.While this physical/virtual browsing is a typical scenario, we imagine other usage scenarios for web present things, in the same way that the Web supports many usage scenarios in addition to Web browsing. Cooltown Developer's Network, People, Place, and Things: web presence for the real world (HP000025–34) (HP000025–34) at 23

Element	'529 Claim Recitation	Exemplary Citations in Reference
		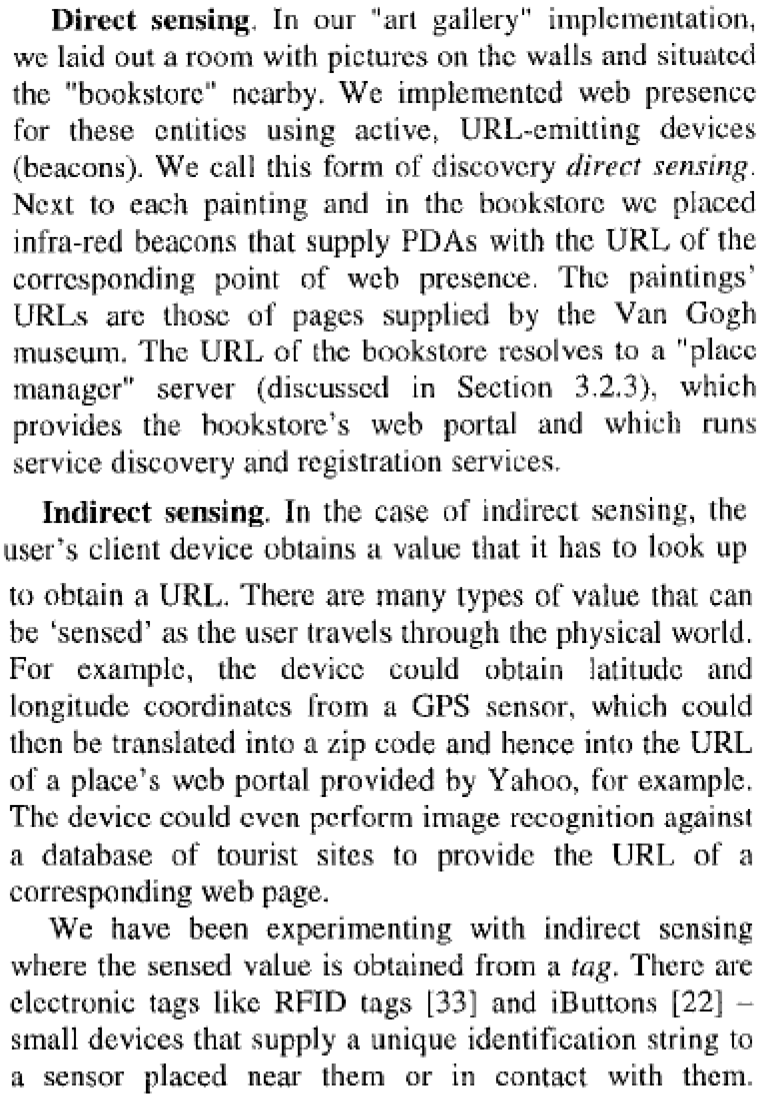 **Direct sensing**. In our "art gallery" implementation, we laid out a room with pictures on the walls and situated the "bookstore" nearby. We implemented web presence for these entities using active, URL-emitting devices (beacons). We call this form of discovery *direct sensing*. Next to each painting and in the bookstore we placed infra-red beacons that supply PDAs with the URL of the corresponding point of web presence. The paintings' URLs are those of pages supplied by the Van Gogh museum. The URL of the bookstore resolves to a "place manager" server (discussed in Section 3.2.3), which provides the bookstore's web portal and which runs service discovery and registration services. **Indirect sensing**. In the case of indirect sensing, the user's client device obtains a value that it has to look up to obtain a URL. There are many types of value that can be 'sensed' as the user travels through the physical world. For example, the device could obtain latitude and longitude coordinates from a GPS sensor, which could then be translated into a zip code and hence into the URL of a place's web portal provided by Yahoo, for example. The device could even perform image recognition against a database of tourist sites to provide the URL of a corresponding web page. We have been experimenting with indirect sensing where the sensed value is obtained from a *tag*. There are electronic tags like RFID tags [33] and iButtons [22] – small devices that supply a unique identification string to a sensor placed near them or in contact with them.

Element	'529 Claim Recitation	Exemplary Citations in Reference
		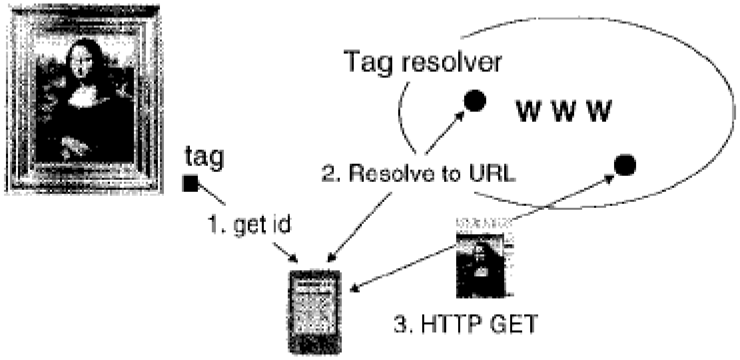 **Figure 4. Indirect sensing: the web presence for a painting.** Barcodes can be used instead of an electronic tag. Figure 4 shows a different art gallery demonstrator in which a PDA client senses a painting's tag. A correlated web page is resolved from the tag identifier and the client presents it to the user. We have produced a prototype implementation for sensing tagged things and displaying a corresponding context-dependent web page. The sensing device knows the URL of a *resolver*. When the client senses a tag, it sends the identifier to the resolver. If the resolver has an entry for the identifier, it sends back the corresponding URL. The user may choose a resolver that reflects his or her personal requirements, or those of a group of users with interests or tasks in common. For example, visitors to an art gallery could be offered resolvers that provide the URLs of pages in their native language, or pages provided by the art school at which they study.

Element	'529 Claim Recitation	Exemplary Citations in Reference
		A Web-Based Nomadic Computing System (HP000074 – 87) at 6 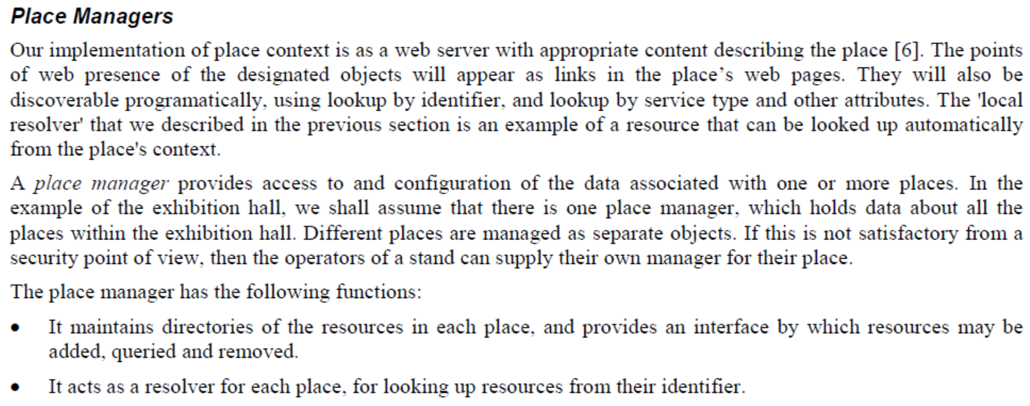 ***Place Managers*** Our implementation of place context is as a web server with appropriate content describing the place [6]. The points of web presence of the designated objects will appear as links in the place's web pages. They will also be discoverable programatically, using lookup by identifier, and lookup by service type and other attributes. The 'local resolver' that we described in the previous section is an example of a resource that can be looked up automatically from the place's context. A *place manager* provides access to and configuration of the data associated with one or more places. In the example of the exhibition hall, we shall assume that there is one place manager, which holds data about all the places within the exhibition hall. Different places are managed as separate objects. If this is not satisfactory from a security point of view, then the operators of a stand can supply their own manager for their place. The place manager has the following functions: - It maintains directories of the resources in each place, and provides an interface by which resources may be added, queried and removed. - It acts as a resolver for each place, for looking up resources from their identifier. To the extent it is found that CoolTown does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that CoolTown does not anticipate this claim, CoolTown renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine CoolTown with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1H]	the network-enabled device further programmed to present the information related to the object to	CoolTown discloses the network-enabled device further programmed to present the information related to the object to a user.

Element	’529 Claim Recitation	Exemplary Citations in Reference
	a user.	A Web-Based Nomadic Computing System (HP000074 – 87) at 4 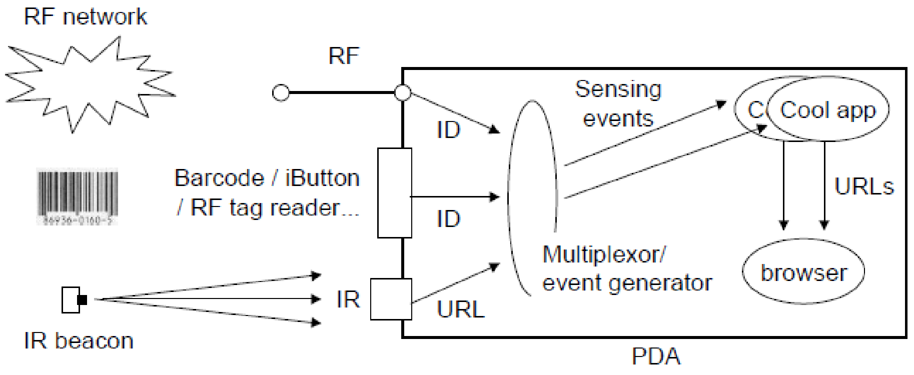 Figure 3. A sketch of a PDA equipped for sensing URLs and identifiers. A Web-Based Nomadic Computing System (HP000074 – 87) at 5 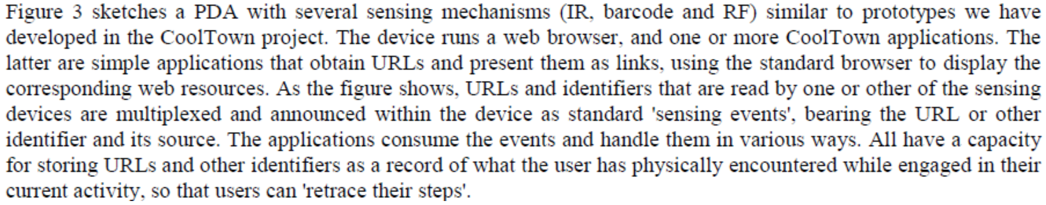 Figure 3 sketches a PDA with several sensing mechanisms (IR, barcode and RF) similar to prototypes we have developed in the CoolTown project. The device runs a web browser, and one or more CoolTown applications. The latter are simple applications that obtain URLs and present them as links, using the standard browser to display the corresponding web resources. As the figure shows, URLs and identifiers that are read by one or other of the sensing devices are multiplexed and announced within the device as standard 'sensing events', bearing the URL or other identifier and its source. The applications consume the events and handle them in various ways. All have a capacity for storing URLs and other identifiers as a record of what the user has physically encountered while engaged in their current activity, so that users can 'retrace their steps'. A Web-Based Nomadic Computing System (HP000074 – 87) at 5 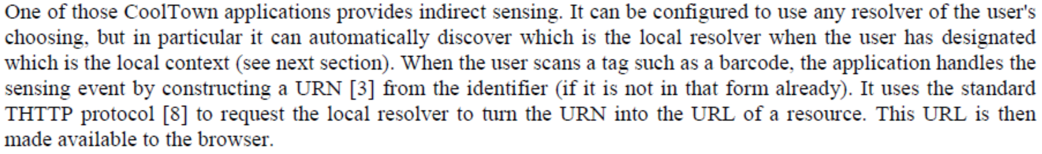 One of those CoolTown applications provides indirect sensing. It can be configured to use any resolver of the user's choosing, but in particular it can automatically discover which is the local resolver when the user has designated which is the local context (see next section). When the user scans a tag such as a barcode, the application handles the sensing event by constructing a URN [3] from the identifier (if it is not in that form already). It uses the standard THTTP protocol [8] to request the local resolver to turn the URN into the URL of a resource. This URL is then made available to the browser. A Web-Based Nomadic Computing System (HP000074 – 87) at 6

Element	'529 Claim Recitation	Exemplary Citations in Reference
		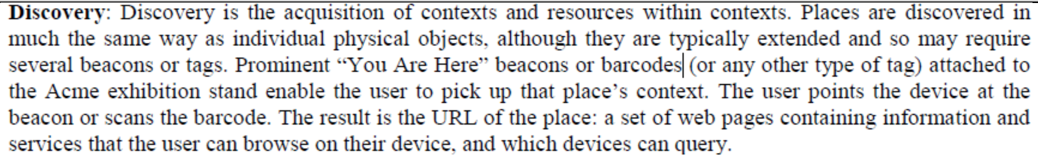 **Discovery**: Discovery is the acquisition of contexts and resources within contexts. Places are discovered in much the same way as individual physical objects, although they are typically extended and so may require several beacons or tags. Prominent "You Are Here" beacons or barcodes (or any other type of tag) attached to the Acme exhibition stand enable the user to pick up that place's context. The user points the device at the beacon or scans the barcode. The result is the URL of the place: a set of web pages containing information and services that the user can browse on their device, and which devices can query. Cooltown Developer's Network, People, Place, and Things: web presence for the real world (HP000025–34) at 5 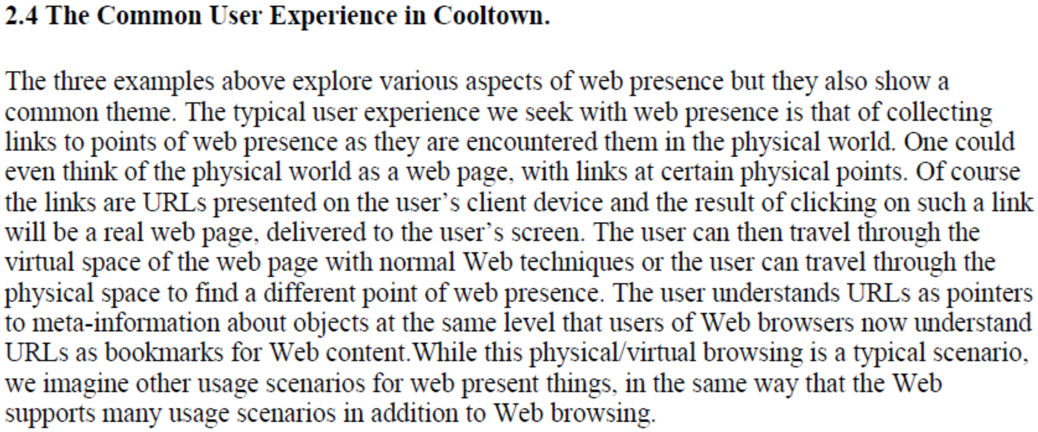 **2.4 The Common User Experience in Cooltown.** The three examples above explore various aspects of web presence but they also show a common theme. The typical user experience we seek with web presence is that of collecting links to points of web presence as they are encountered them in the physical world. One could even think of the physical world as a web page, with links at certain physical points. Of course the links are URLs presented on the user's client device and the result of clicking on such a link will be a real web page, delivered to the user's screen. The user can then travel through the virtual space of the web page with normal Web techniques or the user can travel through the physical space to find a different point of web presence. The user understands URLs as pointers to meta-information about objects at the same level that users of Web browsers now understand URLs as bookmarks for Web content.While this physical/virtual browsing is a typical scenario, we imagine other usage scenarios for web present things, in the same way that the Web supports many usage scenarios in addition to Web browsing. Cooltown Developer's Network, People, Place, and Things: web presence for the real world (HP000025–34) at 9

Nantworks, LLC v. Bank of America Corporation
Central District of California, No. 2:20-cv-7872
Invalidity Chart for U.S. Patent 7,881,529 based on
Invalidity Chart for U.S. Patent 7,881,529 based on Hewlett-Packard CoolTown System ("CoolTown")

Element	'529 Claim Recitation	Exemplary Citations in Reference
		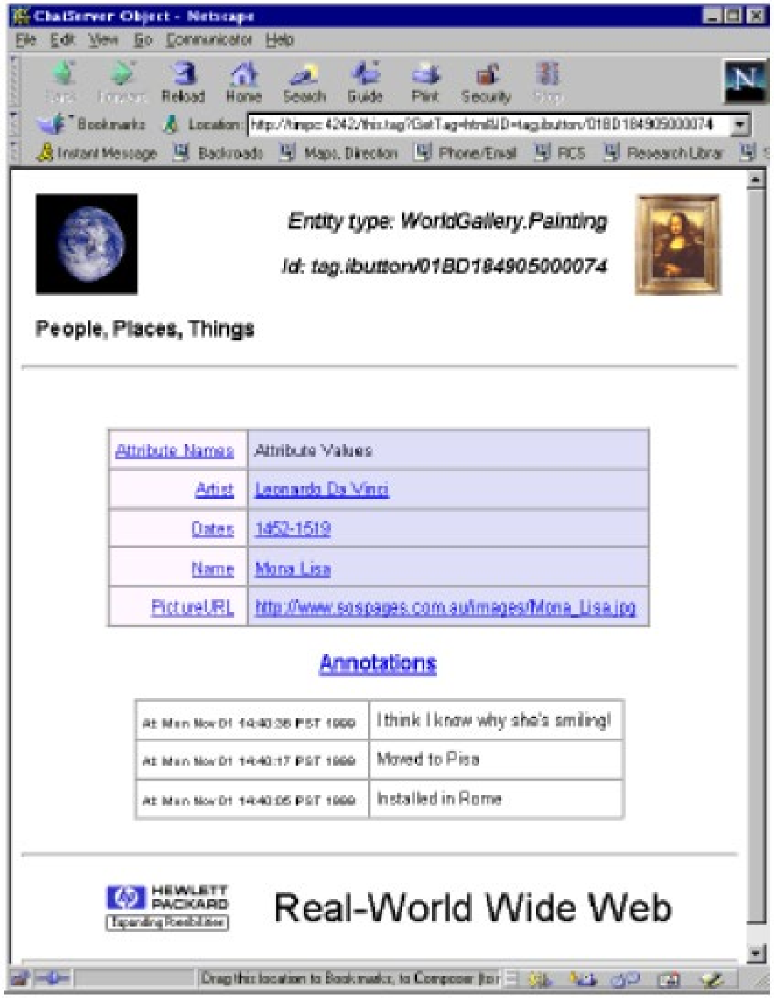 Entity type: WorldGallery.Painting Id: tag.ibutton/01BD184905000074 People, Places, Things Attribute Names Attribute Values Artist Leonardo Da Vinci Dates 1452-1519 Name Mona Lisa PictureURL http://www.sospages.com.au/images/Mona_Lisa.jpg Annotations At Mon Nov 01 14:40:36 PST 1999 I think I know why she's smiling! At Mon Nov 01 14:40:17 PST 1999 Moved to Pisa At Mon Nov 01 14:40:05 PST 1999 Installed in Rome HEWLETT PACKARD Expanding Possibilities Real-World Wide Web Figure 5. A Web presence for *Mona Lisa*.

Element	'529 Claim Recitation	Exemplary Citations in Reference
		Cooltown Developer's Network, People, Place, and Things: web presence for the real world (HP000025–34) (HP000025–34) at 23 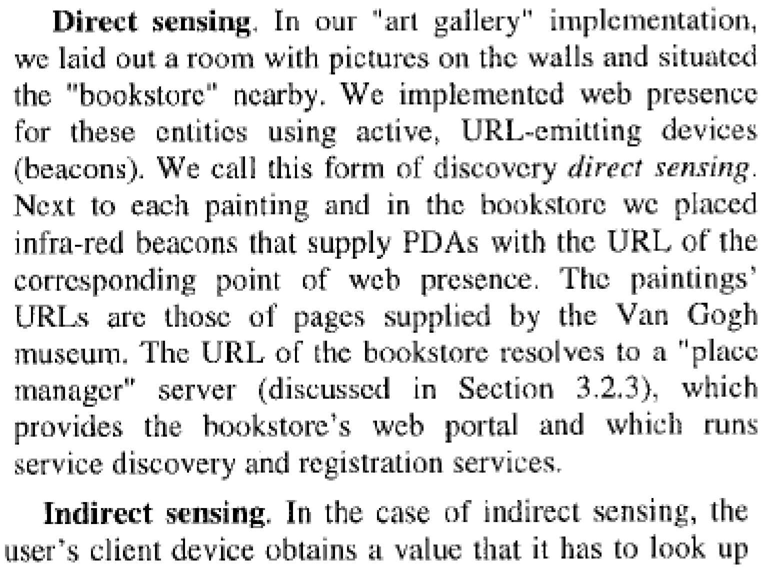 **Direct sensing**. In our "art gallery" implementation, we laid out a room with pictures on the walls and situated the "bookstore" nearby. We implemented web presence for these entities using active, URL-emitting devices (beacons). We call this form of discovery *direct sensing*. Next to each painting and in the bookstore we placed infra-red beacons that supply PDAs with the URL of the corresponding point of web presence. The paintings' URLs are those of pages supplied by the Van Gogh museum. The URL of the bookstore resolves to a "place manager" server (discussed in Section 3.2.3), which provides the bookstore's web portal and which runs service discovery and registration services. **Indirect sensing**. In the case of indirect sensing, the user's client device obtains a value that it has to look up

Element	'529 Claim Recitation	Exemplary Citations in Reference
		to obtain a URL. There are many types of value that can be 'sensed' as the user travels through the physical world. For example, the device could obtain latitude and longitude coordinates from a GPS sensor, which could then be translated into a zip code and hence into the URL of a place's web portal provided by Yahoo, for example. The device could even perform image recognition against a database of tourist sites to provide the URL of a corresponding web page. We have been experimenting with indirect sensing where the sensed value is obtained from a *tag*. There are electronic tags like RFID tags [33] and iButtons [22] – small devices that supply a unique identification string to a sensor placed near them or in contact with them. 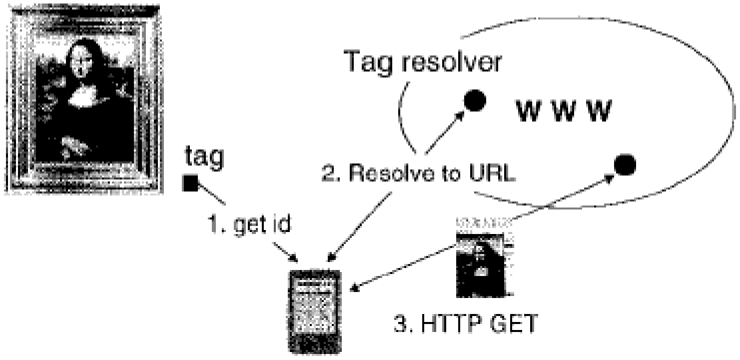 **Figure 4. Indirect sensing: the web presence for a painting.**

Element	'529 Claim Recitation	Exemplary Citations in Reference
		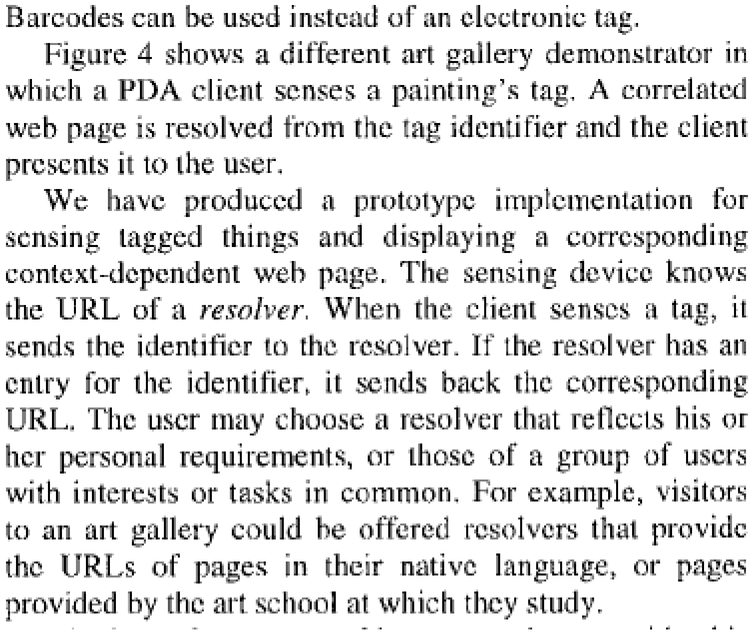 Barcodes can be used instead of an electronic tag. Figure 4 shows a different art gallery demonstrator in which a PDA client senses a painting's tag. A correlated web page is resolved from the tag identifier and the client presents it to the user. We have produced a prototype implementation for sensing tagged things and displaying a corresponding context-dependent web page. The sensing device knows the URL of a *resolver*. When the client senses a tag, it sends the identifier to the resolver. If the resolver has an entry for the identifier, it sends back the corresponding URL. The user may choose a resolver that reflects his or her personal requirements, or those of a group of users with interests or tasks in common. For example, visitors to an art gallery could be offered resolvers that provide the URLs of pages in their native language, or pages provided by the art school at which they study. To the extent it is found that CoolTown does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that CoolTown does not anticipate this claim, CoolTown renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine CoolTown with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

2. Claim 20

Claim	'529 Claim Recitation	Exemplary Citations in Reference
20	The system of claim 1, wherein the data processing operation is selected from the group consisting of contrast enhancement, noise removal, and de-blurring.	To the extent it is found that CoolTown does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that CoolTown does not anticipate this claim, CoolTown renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine CoolTown with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

Nantworks, LLC v. Bank of America Corporation
Central District of California, No. 2:20-cv-7872
Invalidity Chart for U.S. Patent 8,520,897 based on
U.S. Patent No. 6,947,417 to Rhoads ("Rhoads")

Exhibit H-3

Bank of America Corporation ("BofA" or "Defendant") provides this chart subject to all reservations, objections, statements, and disclaimers set forth herein and in BofA's Final Invalidity Contentions Cover Pleading, as well as any amendment, supplement, or modification thereof, which are incorporated herein by reference in their entirety.

On information and belief, the application that matured into Rhoads was filed on May 15, 2000, claiming the priority date of May 19, 1999. Rhoads published at least as of September 20, 2005, and is available as prior art at least under 35 U.S.C. § 102(e).

As shown in the chart below, Rhoads anticipates asserted claims 25, 30, 33–34, and 39 of the '897 Patent. To the extent Rhoads is found not to expressly disclose certain limitations in the asserted claims, such limitations are inherent. To the extent Rhoads is found not to anticipate any asserted claims or claim elements of the '897 Patent, the reference nevertheless renders those claims or claim elements obvious under 35 U.S.C. § 103, either alone or in combination with other art identified herein or in the cover pleading of BofA's Final Invalidity Contentions, for the reasons identified in the cover pleading. These Final Invalidity Contentions are not an admission by BofA that the accused products, including any current or past versions of the accused products, are covered by, or infringe the asserted claims, but are based instead on the recognition that if the claims are interpreted to be broad enough to encompass the accused products, the claims must also be construed to have that same scope when considering whether they are invalid.

The following chart is based on the Court's claim construction rulings (Dkt. Nos. 145, 153, 236), and to the extent not contradictory, Plaintiffs' apparent interpretations of the claims in Plaintiff's "Preliminary" Final Infringement Contentions. BofA notes that in many instances, Plaintiff's "Preliminary" Final Infringement Contentions ignore and/or contradict the Court's claim construction rulings. BofA does not accept the assumptions concerning the scope and meaning implicit in Plaintiff's Preliminary Final Infringement Contentions, to the extent those assumptions are discernible, and reserves the right to challenge Plaintiff's proposed (or implied) constructions. Furthermore, Plaintiffs' "Preliminary" Final Contentions fail to comply with the level of specificity required by SPR 2.1.2 and 2.1.3, and include improper reservation of rights to supplement with "additional ways in which BoA Accused Products infringe." Any alleged deficient discovery is of Plaintiffs own making. Most if not all of the deficiencies in Plaintiffs' Preliminary Final Contentions have been documented by Defendants through multiple letters since at least May 18, 2021 and up to Bank of America's most recent letter of August 19, 2023. Plaintiffs' shifting infringement contentions, deficiencies in its infringement contentions, and changing claim scope have prejudiced BofA from being able to reasonably prepare its defenses, including its invalidity contentions. BofA also reserves the right to revise and supplement these charts if and when Plaintiff is permitted to provide revised or supplemental Infringement Contentions. Citations given in the chart below are merely representative of the respective elements to provide notice to Plaintiffs' of BofA's invalidity theories and are not meant to be exhaustive.

1. Claim 25

Element	'897 Claim Recitation	Exemplary Citations in Reference
[1Pre]	A method of conducting a transaction with an interactive system, the method comprising	Rhoads at Abstract • "A cell phone is equipped with a 2D optical sensor, enabling a variety of applications. For example, such a phone may also be provided with a digital watermark decoder, permitting decoding of steganographic data on imaged objects. Movement of a phone may be inferred by sensing movement of an imaged pattern across the optical sensor's field of view, allowing use of the phone as a gestural input device through which a user can signal instructions to a computer-based process. A variety of other arrangements by which electronic devices can interact with the physical world are also detailed, e.g., involving sensing and responding to digital watermarks, bar codes, RFIDs, etc." Rhoads at 5:54-56. • "Basically, the technology detailed in this disclosure may be regarded as enhanced systems by which users can interact with computer-based devices." Rhoads at 9:35–41 • "The Steganographic decoding process may entail three Steps. In the first, the object is located. In the Second, the objects orientation is discerned. In the third, the Bedoop data is extracted from the image data corresponding to the Bedoop object." Rhoads at Fig. 2 (and accompanying description)

Element	'897 Claim Recitation	Exemplary Citations in Reference
		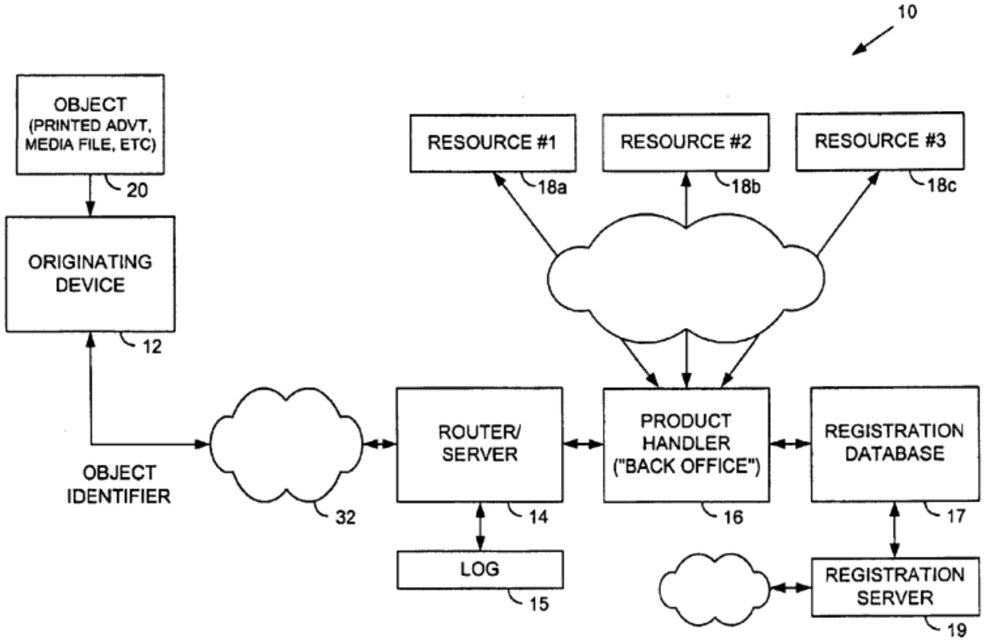 FIG. 2 Rhoads at 49:15-18 - "Referring to FIG. 2, a system 10 according to the exemplary embodiment includes an originating device 12, a router/server 14, a product handler 16, a registration database 17, and one or more remote resources 18." Rhoads at 50:21-48. - "As more particularly detailed below, the handler 16 provides a response in accordance with the particular watermark payload. The response may be provided directly by the product handler to the device 12, or the handler

Element	'897 Claim Recitation	Exemplary Citations in Reference
		may respond by communicating with a remote resource 18 (which may be, e.g., a data repository or service provider). In the former case, the handler 16 may identify a URL corresponding to the watermark (using the database 17), and return the URL to the application 28 c. Application 28 c can then pass the URL to a web browser 28 b in the device 12, and initiate a link to the internet site identified by the URL. Or the handler may have some locally stored data (e.g., audio or video, or software updates) and send it to the device 12 in response to the watermark. In the latter case, the handler 16 does not respond directly to the device 12. Instead, the handler responds by communicating with a remote resource 18. The communication can be as simple as logging receipt of the watermark message in a remote repository. Or it can be to authenticate device 12 (or a user thereof) to a remote resource in anticipation of a further transaction (e.g., the communication can form part of an on-line licensing or digital rights management transaction). Or the communication can request the remote resource to provide data or a service back to device 12 or to another destination (e.g., to initiate an FTP file transfer, or to request that a song selection identified by the watermark be downloaded to a user's personal music library, or to update software installed on device 12). In still other cases, hybrids of the two foregoing cases can be employed, e.g., handler 16 can send some data back to device 12, while also communicating with a remote resource 18. In some cases, the response returned to the device 12 by handler 6 (or a remote resource 18) can serve to trigger some further action by the device 12. For example, the response returned to device 12 can include a WindowsMedia audio file, together with a request that the device 12 launch the WindowsMedia player installed on the device. (The launching of a browser pointed to a URL is another example of such triggering.)"
[1A]	providing access to a device having a display;	Rhoads discloses providing access to a device having a display.

Element	'897 Claim Recitation	Exemplary Citations in Reference
		Rhoads at 46:47-53 • "The provision of image Sensors in cell phones enables other functionality. One is the capture of still or video imagery. Such image data can be compressed (typically by lossy processes such as MPEG, JPEG, or the like, implemented with dedicated hardware CODECs) and transmitted with the audio data. The Screens on Such phones can likewise be used for display of incoming image or Video data." To the extent it is found that Carlton does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Carlton does not anticipate this claim, Carlton renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Carlton with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

Element	'897 Claim Recitation	Exemplary Citations in Reference
[1B]	displaying a displayed image on the display of the device;	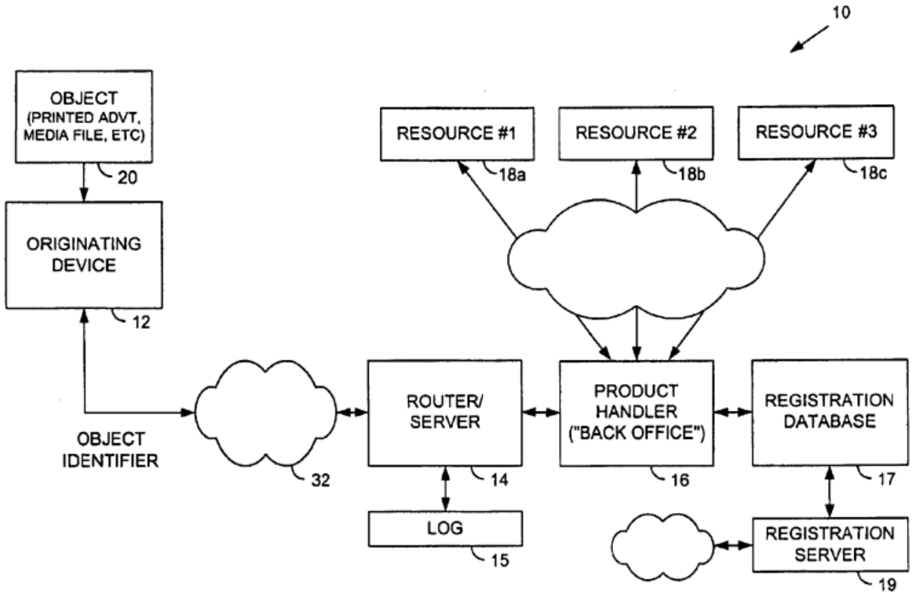 **FIG. 2** Rhoads at 49:24-27 - "Device 12 interacts with an object 20. The object can be electronic or not. Electronic objects can include computer files, representations of audio, video, or still imagery (e.g., files or in streaming form), etc." Rhoads at Fig. 3

Element	'897 Claim Recitation	Exemplary Citations in Reference
		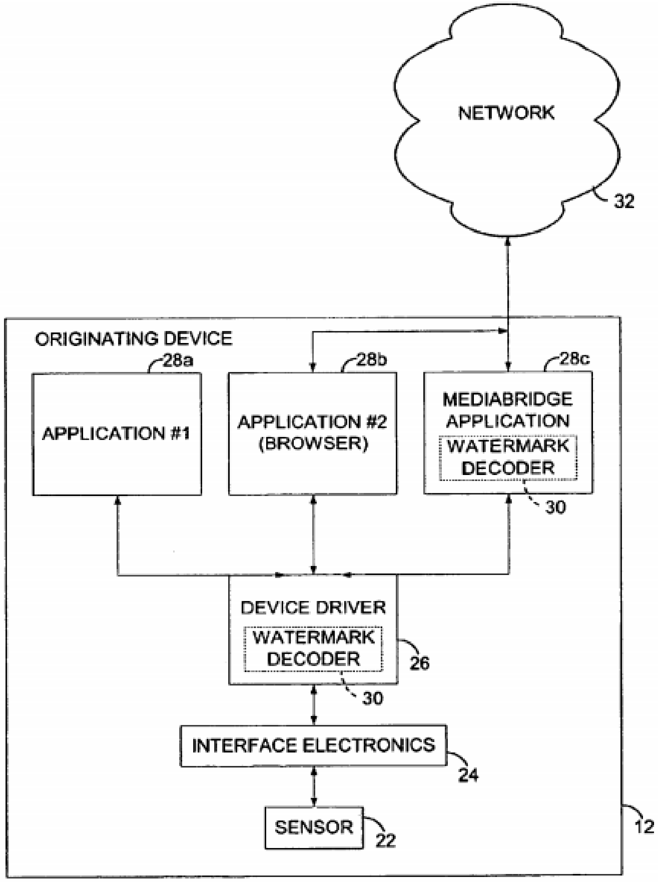 **FIG. 3** Rhoads at 7:24-42 • "Device 12 interacts with an object 20. The object can be electronic or not. Electronic objects can include computer files, representations of

Element	'897 Claim Recitation	Exemplary Citations in Reference
		audio, video, or still imagery (e.g., files or in streaming form), etc. Non-electronic objects can include physical objects such as newspapers, magazine pages, posters, product packaging, event tickets, credit cards, paper currency, etc. Non-electronic objects can also include sounds produced by loudspeakers. • When used with nonelectronic objects, device 12 (FIG. 2) typically includes some form of sensor or transducer 22 to produce electronic signals or data corresponding to the object. Examples include CCD- or CMOS-based optical sensors (either as part of still- or video cameras, flatbed scanners, mice, or otherwise), microphones, barcode scanners, RF ID sensors, mag stripe readers, etc. In such cases, the sensor 22 may be coupled to associated interface electronics 24, which in turn may be coupled to device driver software 26, which in turn may be coupled to one or more application programs 28." Rhoads at 7:13-29 • "Referring to FIG. 11, a basic embodiment 110 of the present technology includes an optical sensor 112, a computer 114, and a network connection 116 to the internet 118. The illustrated optical sensor 112 is a digital camera having a resolution of 320 by 200 pixels (color or black and white) that stares out, grabbing frames of image data five times per second and storing same in one or more frame buffers. These frames of image data are analyzed by a computer 114 for the presence of Bedoop data. (Essentially, Bedoop data is any form of plural-bit data encoding recognized by the system 110—data which, in many embodiments, initiates some action.) Once detected, the system responds in accordance with the detected Bedoop data (e.g., by initiating some local action, or by communication with a remote computer, such as over the internet, via an online service such as AOL, or using point-to-point dial-up communications, as with a bulletin board system)."

Element	'897 Claim Recitation	Exemplary Citations in Reference
		Rhoads at 9:35–10:10 - "The Steganographic decoding process may entail three Steps. In the first, the object is located. In the Second, the objects orientation is discerned. In the third, the Bedoop data is extracted from the image data corresponding to the Bedoop object. The first step, object location, can be assisted by various clues. One is the placement of the object; typically the center of the image field will be a point on the object. The surrounding data can then be analyzed to try and discern the object's boundaries. Another location technique is slight movement. Although the user will typically try to hold the object still, there will usually be some jitter of the Bedoop object within the image frame (e.g., a few pixels back and forth). Background visual clutter, in contrast, will typically be stationary. Such movement may thus be sensed and used to identify the Bedoop object from within the image data. Still another object-location clue is object shape. Many Bedoop objects are rectangular in shape (or trapezoidal as viewed by the camera). Straight edge boundaries can thus be used to define an area of likely Bedoop data. Color is a further object identification clue that may be useful in some contexts. Yet another object location clue is spatial frequency. In imaging systems with well defined focal zones, undesired visual clutter may be at focal distances that results in blurring. The Bedoop object, in contrast, will be in focus and may be characterized by fine detail. Analyzing the image data for the high frequencies associated with fine detail can be used to distinguish the intended object from others.

Element	'897 Claim Recitation	Exemplary Citations in Reference
		Characteristic markings on the object (as discussed below in connection with determining object orientation) can also be sensed and used in locating the object. Once the Bedoop object has been located within the image data, masking can be applied (if desired) to eliminate image data not corresponding to the intended object." Rhoads at 46:47-53 - "The provision of image Sensors in cell phones enables other functionality. One is the capture of still or video imagery. Such image data can be compressed (typically by lossy processes such as MPEG, JPEG, or the like, implemented with dedicated hardware CODECs) and transmitted with the audio data. The Screens on Such phones can likewise be used for display of incoming image or Video data." To the extent it is found that Carlton does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Carlton does not anticipate this claim, Carlton renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Carlton with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1C]	presenting the displayed image proximate to an optical sensor; and	Rhoads discloses presenting the displayed image proximate to an optical sensor.

Element	'897 Claim Recitation	Exemplary Citations in Reference
		Rhoads at Fig. 2 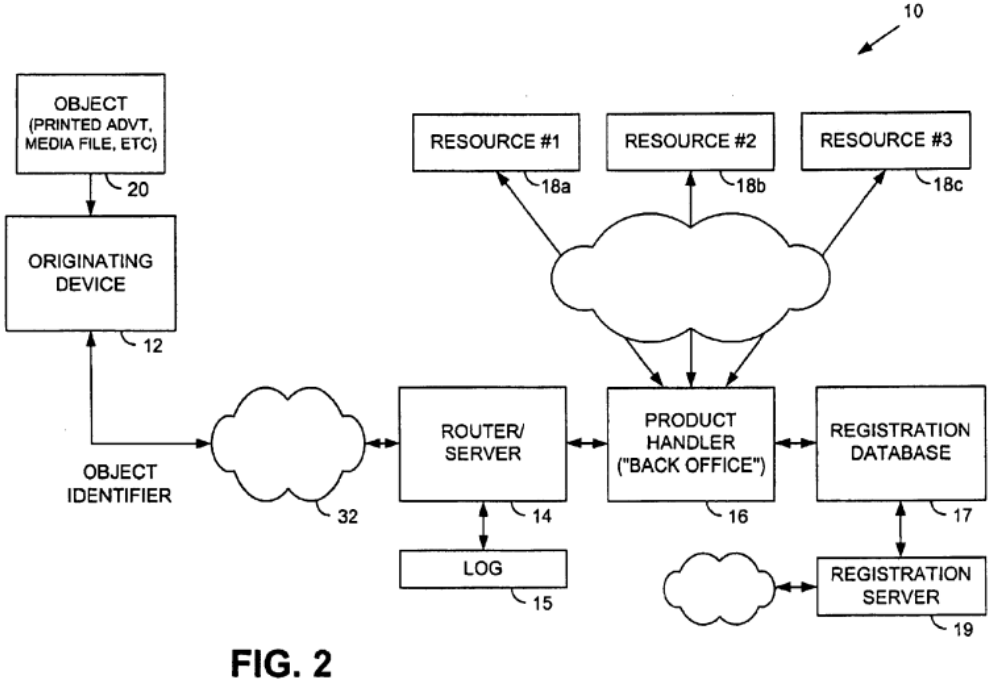 FIG. 2 Rhoads at 3:41–51 - "According to another aspect, the invention includes a promotional method comprising: (a) presenting an object within the field of view of an optical sensor device, the object being selected from the list consisting of a retail product, packaging for a retail product, or printed advertising; (b) acquiring optical data corresponding to the object; (c) decoding plural-bit digital data from the optical data; (d) submitting at least some of said decoded data to a remote computer; and (e) determining at the remote computer whether a prize should be awarded in response to submission of said decoded data."

Element	'897 Claim Recitation	Exemplary Citations in Reference
		Rhoads at 7:24-42 • "Device 12 interacts with an object 20. The object can be electronic or not. Electronic objects can include computer files, representations of audio, video, or still imagery (e.g., files or in streaming form), etc. Non-electronic objects can include physical objects such as newspapers, magazine pages, posters, product packaging, event tickets, credit cards, paper currency, etc. Non-electronic objects can also include sounds produced by loudspeakers. When used with nonelectronic objects, device 12 (FIG. 2) typically includes some form of sensor or transducer 22 to produce electronic signals or data corresponding to the object. Examples include CCD- or CMOS-based optical sensors (either as part of still- or video cameras, flatbed scanners, mice, or otherwise), microphones, barcode scanners, RF ID sensors, mag stripe readers, etc. In such cases, the sensor 22 may be coupled to associated interface electronics 24, which in turn may be coupled to device driver software 26, which in turn may be coupled to one or more application programs 28." Rhoads at Fig. 11

Element	'897 Claim Recitation	Exemplary Citations in Reference
		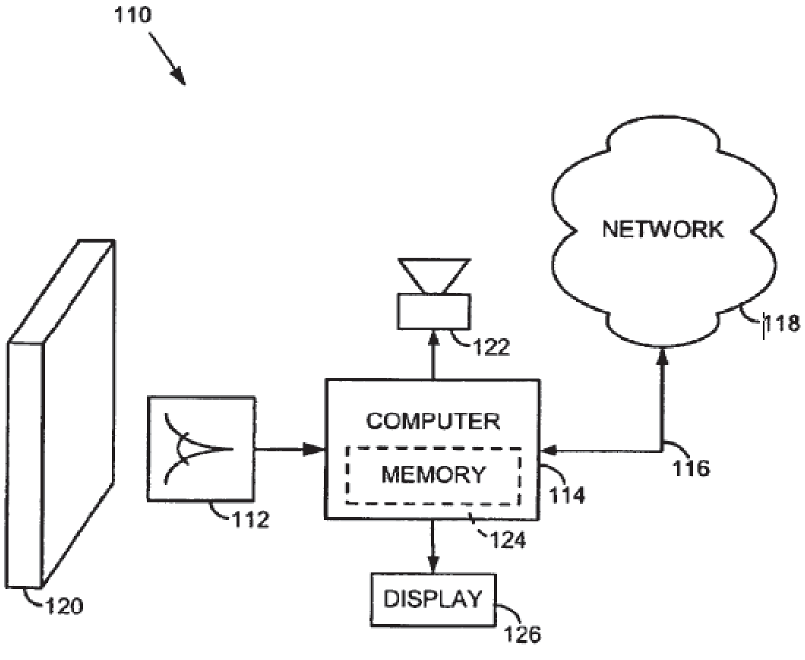 FIG. 11 Rhoads at 7:13-29 - "Referring to FIG. 11, a basic embodiment 110 of the present technology includes an optical sensor 112, a computer 114, and a network connection 116 to the internet 118. The illustrated optical sensor 112 is a digital camera having a resolution of 320 by 200 pixels (color or black and white) that stares out, grabbing frames of image data five times per second and storing same in one or more frame buffers. These frames of image data are

Element	'897 Claim Recitation	Exemplary Citations in Reference
		analyzed by a computer 114 for the presence of Bedoop data. (Essentially, Bedoop data is any form of plural-bit data encoding recognized by the system 110—data which, in many embodiments, initiates some action.) Once detected, the system responds in accordance with the detected Bedoop data (e.g., by initiating some local action, or by communication with a remote computer, such as over the internet, via an online service such as AOL, or using point-to-point dial-up communications, as with a bulletin board system)." Rhoads 8:66 – 9:12 • "In still other embodiments, the camera (or other sensor) can be equipped with one or more auxiliary fixed-focus lenses that can be selectively used, depending on the particular application. Some such embodiments have a first fixed focused lens that always overlies the sensor, with which one or more auxiliary lenses can be optically cascaded (e.g., by hinge or slide arrangements). Such arrangements are desirable, e.g., when a camera is not a dedicated Bedoop sensor but also performs other imaging tasks. When the camera is to be used for Bedoop, the auxiliary lens is positioned (e.g., flipped into) place, changing the focal length of the first lens (which may by unsuitably long for Bedoop purposes, such as infinity) to an appropriate Bedoop imaging range (such as one foot)." Rhoads at 17:46–54 • "The building can be provided with an image sensor (such as a video camera or the like), an associated Bedoop. detection system, and the proximity detector. When a user wearing the badge approaches, the proximity detector signals the camera to capture image data. The Bedoop detection system identifies the badge photograph (e.g., by clues as are described in the prior applications, or without such aids), captures optical

Element	'897 Claim Recitation	Exemplary Citations in Reference
		data, and decodes same to extract the steganographically-embedded data hidden therein." Rhoads at 46:47-53 • "The provision of image Sensors in cell phones enables other functionality. One is the capture of still or video imagery. Such image data can be compressed (typically by lossy processes such as MPEG, JPEG, or the like, implemented with dedicated hardware CODECs) and transmitted with the audio data. The Screens on Such phones can likewise be used for display of incoming image or Video data." Rhoads at 50:21-48 • "As more particularly detailed below, the handler 16 provides a response in accordance with the particular watermark payload. The response may be provided directly by the product handler to the device 12, or the handler may respond by communicating with a remote resource 18 (which may be, e.g., a data repository or service provider). In the former case, the handler 16 may identify a URL corresponding to the watermark (using the database 17), and return the URL to the application 28 c. Application 28 c can then pass the URL to a web browser 28 b in the device 12, and initiate a link to the internet site identified by the URL. Or the handler may have some locally stored data (e.g., audio or video, or software updates) and send it to the device 12 in response to the watermark. In the latter case, the handler 16 does not respond directly to the device 12. Instead, the handler responds by communicating with a remote resource 18. The communication can be as simple as logging receipt of the watermark message in a remote repository. Or it can be to authenticate device 12 (or a user thereof) to a remote resource in anticipation of a further transaction (e.g., the communication can form part of an on-line licensing

Element	'897 Claim Recitation	Exemplary Citations in Reference
		or digital rights management transaction). Or the communication can request the remote resource to provide data or a service back to device 12 or to another destination (e.g., to initiate an FTP file transfer, or to request that a song selection identified by the watermark be downloaded to a user's personal music library, or to update software installed on device 12). In still other cases, hybrids of the two foregoing cases can be employed, e.g., handler 16 can send some data back to device 12, while also communicating with a remote resource 18. In some cases, the response returned to the device 12 by handler 6 (or a remote resource 18) can serve to trigger some further action by the device 12. For example, the response returned to device 12 can include a WindowsMedia audio file, together with a request that the device 12 launch the WindowsMedia player installed on the device. (The launching of a browser pointed to a URL is another example of such triggering.)" To the extent it is found that Carlton does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Carlton does not anticipate this claim, Carlton renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Carlton with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1D]	enabling, via the interactive system, identification of features from the displayed image, recognition of a target based on the features,	Rhoads discloses enabling, via the interactive system, identification of features from the displayed image, recognition of a target based on the features, association of the target with target information pertinent to the target, and performance of a transaction based on the target information.

Element	'897 Claim Recitation	Exemplary Citations in Reference
	association of the target with target information pertinent to the target, and performance of a transaction based on the target information.	Rhoads at Fig. 2 (and accompanying description) 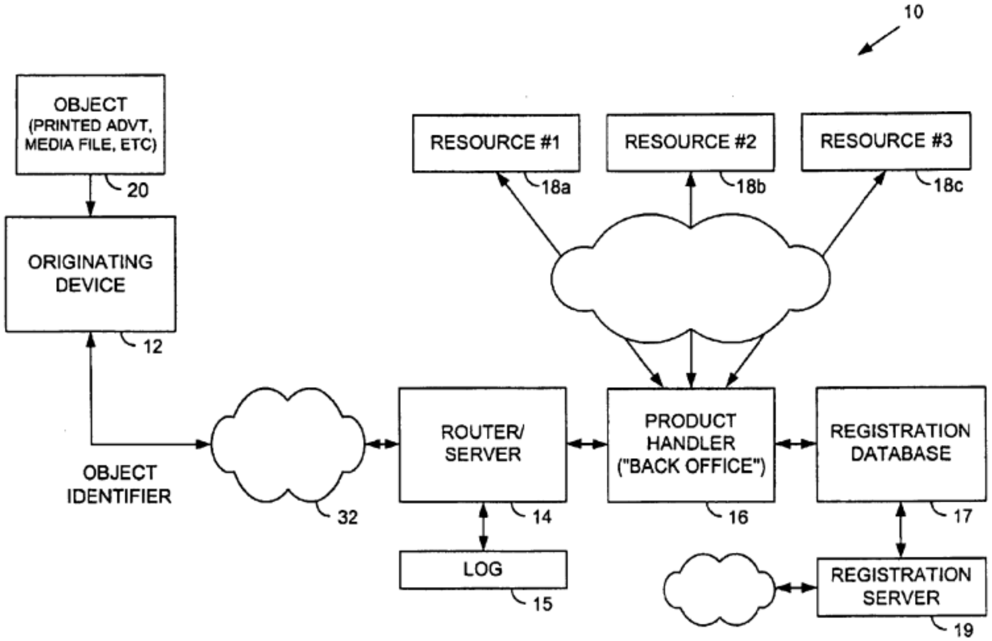 FIG. 2 Rhoads at 9:35–10:21 - "The Steganographic decoding process may entail three Steps. In the first, the object is located. In the Second, the objects orientation is discerned. In the third, the Bedoop data is extracted from the image data corresponding to the Bedoop object. The first step, object location, can be assisted by various clues. One is the placement of the object; typically the center of the image field will be a point on the object. The surrounding data can then be analyzed to try and discern the object's boundaries.

Element	'897 Claim Recitation	Exemplary Citations in Reference
		Another location technique is slight movement. Although the user will typically try to hold the object still, there will usually be some jitter of the Bedoop object within the image frame (e.g., a few pixels back and forth). Background visual clutter, in contrast, will typically be stationary. Such movement may thus be sensed and used to identify the Bedoop object from within the image data. Still another object-location clue is object shape. Many Bedoop objects are rectangular in shape (or trapezoidal as viewed by the camera). Straight edge boundaries can thus be used to define an area of likely Bedoop data. Color is a further object identification clue that may be useful in some contexts. Yet another object location clue is spatial frequency. In imaging systems with well defined focal zones, undesired visual clutter may be at focal distances that results in blurring. The Bedoop object, in contrast, will be in focus and may be characterized by fine detail. Analyzing the image data for the high frequencies associated with fine detail can be used to distinguish the intended object from others. Characteristic markings on the object (as discussed below in connection with determining object orientation) can also be sensed and used in locating the object. Once the Bedoop object has been located within the image data, masking can be applied (if desired) to eliminate image data not corresponding to the intended object. The next step in the decoding process, determining orientation of the Bedoop data, can likewise be discerned by reference to visual clues. For example, some objects include subliminal graticule data, or other

Element	'897 Claim Recitation	Exemplary Citations in Reference
		calibration data, steganographically encoded with the Bedoop data to aid in determining orientation. Others can employ overt markings, either placed for that sole purpose (e.g. reference lines or fiducials), or serving another purpose as well (e.g. lines of text), to discern orientation. Edge-detection algorithms can also be employed to deduce the orientation of the object by reference to its edges." Rhoads at 16:66 – 17:10 • "Inside the building the candidate may encounter a security guard. Seeing an unfamiliar person, the guard may visually compare the photo on the badge with the candidate's face. Additionally, the guard can present the badge to a portable Bedoop device, or to one of many Bedoop systems scattered through the building (e g., at every telephone). The Bedoop system extracts the Bedoop data from the card (i.e., from the DMV photograph), interrogates the DMV's DNS server with this Bedoop data, and receives in reply the name of the person depicted in the photograph. (If the Bedoop system is a telephone, the name may be displayed on a small LCD display commonly provided on telephones.)" Rhoads at 22:9–27 • "To aid in identification, the Bedoop objects (e.g., badges) can be colored a distinctive color, permitting the system to more easily identify candidate objects from other items within the optical capture devices' field of view. Or the object can be provided with a retro-reflective coating, and the elevator can be equipped with one or more illumination sources of known spectral or temporal quality (e.g., constant infra red, or constant illumination with a single- or multi-line spectrum, or a pulsed light source of known periodicity; LEDs or semiconductor lasers, each with an associated diffuser, can be used for each the foregoing and can be paired with the image capture devices). Other such tell-tale clues can likewise be

Element	'897 Claim Recitation	Exemplary Citations in Reference
		used to aid in object location. In all such cases, the optical capture device can sense the tell-tale clue(s) using a wide field of view sensor. The device can then be physically or electronically steered, and/or zoomed, to acquire a higher resolution image of the digitally-encoded object suitable for decoding." Rhoads at 23:33–54 • "When a recipient of a business card holds it in front of a Bedoop sensor, the operating system on the local system launches a local Bedoop application. That local Bedoop application, in turn, establishes an external internet connection to a remote business card server. The address of that server may already be known to the local Bedoop application (e.g., having been stored from previous use), or the local Bedoop system can traverse the above-described public network of DNS servers to reach the business card server. A database on the business card name server maintains a large collection of business card data, one database record per UID. When that server receives Bedoop data from a local Bedoop system, it parses out the UID and accesses the corresponding database record. This record typically includes more information than is commonly printed on conventional business cards. Sample fields from the record may include, for example, name, title, office phone, office fax, home phone, home fax, cellular phone, email address, company name, corporate web page address, personal web page address, secretary's name, spouse's name, and birthday. This record is transmitted back to the originating Bedoop system." Rhoads at 48:32–53 • "As noted, the following disclosure focuses on one particular application—a system for linking print media to electronic content It should be reiterated,

Element	'897 Claim Recitation	Exemplary Citations in Reference
		however, that the technology is not so limited, and may more generally be viewed as a system for linking any object (physical or electronic) to a corresponding networked or local resource. In accordance with an exemplary application of the below-detailed technology, digital watermarking is employed to convey a plural bit identifier within print media, such as magazine advertisements or articles, direct mail coupons or catalogs, bank- or credit-cards, and business cards. This identifier is read by software in a user's computing device and forwarded to a remote database. The remote database identifies a URL corresponding to the identifier, and provides the URL back to the user's computer—permitting a browser on the user's computer to display the URL-identified web page. That web page can provide additional information or services—more timely and/or more extensive than that provided by the print material. By such arrangement, more efficient internet navigation and access is provided to consumers, and more effective means for linking readers to e-commerce points of sale is provided to advertisers." Rhoads at 49:3-14: • "Detection is the complementary operation to embedding, i.e., discerning a digital identifier from an object. Response refers to the action taken based on the discerned identifier. The middle two steps—embedding and detection—can employ any of myriad well-known technologies, including 1D and 2D barcodes, magnetic ink character recognition (MICR), optical character recognition (OCR), optical mark recognition (OMR), radio frequency identification (RF/ID), data glyphs, organic transistors, magnetic stripe, metadata, file header information, UV/IR identifiers, and other machine-readable indicia and techniques for associating plural-bit digital data with an electronic or physical object. The detailed embodiment employs watermarking

Element	'897 Claim Recitation	Exemplary Citations in Reference
		technology, although this is illustrative only." Rhoads at 50:21-48 • "As more particularly detailed below, the **handler 16** provides a response in accordance with the particular watermark payload. The response may be provided directly by the product handler to the device 12, or the handler may respond by communicating with a remote resource 18 (which may be, e.g., a data repository or service provider). In the former case, the handler 16 may identify a URL corresponding to the watermark (using the database 17), and return the URL to the application 28 c. Application 28 c can then pass the URL to a web browser 28 b in the device 12, and initiate a link to the internet site identified by the URL. Or the handler may have some locally stored data (e.g., audio or video, or software updates) and send it to the device 12 in response to the watermark. In the latter case, the handler 16 does not respond directly to the device 12. Instead, the handler responds by communicating with a remote resource 18. The communication can be as simple as logging receipt of the watermark message in a remote repository. Or it can be to authenticate device 12 (or a user thereof) to a remote resource in anticipation of a further transaction (e.g., the communication can form part of an on-line licensing or digital rights management transaction). Or the communication can request the remote resource to provide data or a service back to device 12 or to another destination (e.g., to initiate an FTP file transfer, or to request that a song selection identified by the watermark be downloaded to a user's personal music library, or to update software installed on device 12). In still other cases, hybrids of the two foregoing cases can be employed, e.g., handler 16 can send some data back to device 12, while also communicating with a remote resource 18.

Element	'897 Claim Recitation	Exemplary Citations in Reference
		In some cases, the response returned to the device 12 by handler 6 (or a remote resource 18) can serve to trigger some further action by the device 12. For example, the response returned to device 12 can include a WindowsMedia audio file, together with a request that the device 12 launch the WindowsMedia player installed on the device. (The launching of a browser pointed to a URL is another example of such triggering.)" Rhoads at 51:18-37 • "When shown a watermarked image, the application 28 c analyzes the image and extracts the embedded watermark payload (more particularly detailed below) from the image. The application sends some or all of this information in a message format to the router 14. The router 14 decodes the received message, looking for vendor and product information. Based on this information, it passes the message to a corresponding product handler 16. The product handler receives the message and attempts to match the detected watermark serial number to a registered watermark serial number earlier stored in the database 17. If a match is found, the product handler performs the desired action. As noted, typical actions include returning a URL for web redirection, serving up an HTML page for initial user navigation, initiating software downloads, etc. If a match is not found, the product handler returns an error code and message to the application 28 c. If a match is found, but the corresponding action is unavailable, incomplete, inactive or invalid, the product handler returns an error code and message to the calling application." Rhoads at 84:11 – 17 • "B55. A method comprising:

Element	'897 Claim Recitation	Exemplary Citations in Reference
		presenting an object to an optical sensor, the optical sensor producing image data, the object being steganographically encoded with plural bits of digital data; decoding said digital data from the image data; and in response to said digital data, displaying an image of a like object." To the extent it is found that Carlton does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Carlton does not anticipate this claim, Carlton renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Carlton with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

2. Claim 30

Claim	'897 Claim Recitation	Exemplary Citations in Reference
30	The method of claim 25, wherein the step of displaying the displayed image includes displaying a time varying image.	Rhoads discloses the step of displaying the displayed image includes displaying a time varying image. Rhoads at 7:24-42 • "Device 12 interacts with an object 20. The object can be electronic or not. Electronic objects can include computer files, representations of audio, video, or still imagery (e.g., files or in streaming form), etc. Non-

		electronic objects can include physical objects such as newspapers, magazine pages, posters, product packaging, event tickets, credit cards, paper currency, etc. Non-electronic objects can also include sounds produced by loudspeakers. When used with nonelectronic objects, device 12 (FIG. 2) typically includes some form of sensor or transducer 22 to produce electronic signals or data corresponding to the object. Examples include CCD- or CMOS-based optical sensors (either as part of still- or video cameras, flatbed scanners, mice, or otherwise), microphones, barcode scanners, RF ID sensors, mag stripe readers, etc. In such cases, the sensor 22 may be coupled to associated interface electronics 24, which in turn may be coupled to device driver software 26, which in turn may be coupled to one or more application programs 28." Rhoads at 17:46–54 • "The building can be provided with an image sensor (such as a video camera or the like), an associated Bedoop. detection system, and the proximity detector. When a user wearing the badge approaches, the proximity detector signals the camera to capture image data. The Bedoop detection system identifies the badge photograph (e.g., by clues as are described in the prior applications, or without such aids), captures optical data, and decodes same to extract the steganographically-embedded data hidden therein." Rhoads at 46:47-53 • "The provision of image Sensors in cell phones enables other functionality. One is the capture of still or video imagery. Such image data can be compressed (typically by lossy processes such as MPEG, JPEG, or the like, implemented with dedicated hardware CODECs) and transmitted with the audio data. The Screens on Such phones can likewise be used for display of incoming image or Video data."

		To the extent it is found that Carlton does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Carlton does not anticipate this claim, Carlton renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Carlton with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

3. Claim 33

Claim	'897 Claim Recitation	Exemplary Citations in Reference
33	The method of claim 25, wherein the step of enabling performance of the transaction includes initiating a verification.	Rhoads discloses the step of enabling performance of the transaction includes initiating a verification. Rhoads at 16:66 – 17:16 • "Inside the building the candidate may encounter a security guard. Seeing an unfamiliar person, the guard may visually compare the photo on the badge with the candidate's face. Additionally, the guard can present the badge to a portable Bedoop device, or to one of many Bedoop systems scattered through the building (e g., at every telephone). The Bedoop system extracts the Bedoop data from the card (i.e., from the DMV photograph), interrogates the DMV's DNS server with this Bedoop data, and receives in reply the name of the person depicted in the photograph. (If the Bedoop system is a telephone, the name may be displayed on a small LCD display commonly provided on telephones.)

		The guard checks the name returned by the Bedoop System with the name printed on the badge. On Seeing that the printed and Bedoop-decoded names match (and optionally checking the door log to see that a person of that name was authorized to enter and did So), the Security guard can let the candidate pass." Rhoads at 17:46–54 • "The building can be provided with an image sensor (such as a video camera or the like), an associated Bedoop. detection system, and the proximity detector. When a user wearing the badge approaches, the proximity detector signals the camera to capture image data. The Bedoop detection system identifies the badge photograph (e.g., by clues as are described in the prior applications, or without such aids), captures optical data, and decodes same to extract the steganographically-embedded data hidden therein." Rhoads at 50:21-48. • "As more particularly detailed below, the handler 16 provides a response in accordance with the particular watermark payload. The response may be provided directly by the product handler to the device 12, or the handler may respond by communicating with a remote resource 18 (which may be, e.g., a data repository or service provider). In the former case, the handler 16 may identify a URL corresponding to the watermark (using the database 17), and return the URL to the application 28 c. Application 28 c can then pass the URL to a web browser 28 b in the device 12, and initiate a link to the internet site identified by the URL. Or the handler may have some locally stored data (e.g., audio or video, or software updates) and send it to the device 12 in response to the watermark.

		In the latter case, the handler 16 does not respond directly to the device 12. Instead, the handler responds by communicating with a remote resource 18. The communication can be as simple as logging receipt of the watermark message in a remote repository. Or it can be to authenticate device 12 (or a user thereof) to a remote resource in anticipation of a further transaction (e.g., the communication can form part of an on-line licensing or digital rights management transaction). Or the communication can request the remote resource to provide data or a service back to device 12 or to another destination (e.g., to initiate an FTP file transfer, or to request that a song selection identified by the watermark be downloaded to a user's personal music library, or to update software installed on device 12). In still other cases, hybrids of the two foregoing cases can be employed, e.g., handler 16 can send some data back to device 12, while also communicating with a remote resource 18. In some cases, the response returned to the device 12 by handler 6 (or a remote resource 18) can serve to trigger some further action by the device 12. For example, the response returned to device 12 can include a WindowsMedia audio file, together with a request that the device 12 launch the WindowsMedia player installed on the device. (The launching of a browser pointed to a URL is another example of such triggering.)" To the extent it is found that Carlton does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Carlton does not anticipate this claim, Carlton renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Carlton with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these

		references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

4. Claim 34

Claim	'897 Claim Recitation	Exemplary Citations in Reference
34	The method of claim 25 wherein the step of enabling performance of the transaction includes initiating an authorization.	Rhoads discloses the step of enabling performance of the transaction includes initiating an authorization. Rhoads at 16:66 – 17:16 • "Inside the building the candidate may encounter a security guard. Seeing an unfamiliar person, the guard may visually compare the photo on the badge with the candidate's face. Additionally, the guard can present the badge to a portable Bedoop device, or to one of many Bedoop systems scattered through the building (e g., at every telephone). The Bedoop system extracts the Bedoop data from the card (i.e., from the DMV photograph), interrogates the DMV's DNS server with this Bedoop data, and receives in reply the name of the person depicted in the photograph. (If the Bedoop system is a telephone, the name may be displayed on a small LCD display commonly provided on telephones.) The guard checks the name returned by the Bedoop System with the name printed on the badge. On Seeing that the printed and Bedoop-decoded names match (and optionally checking the door log to see that a person of that name was authorized to enter and did So), the Security guard can let the candidate pass." Rhoads at 19:6–8 • "In another example, the UID field serves an authentication purpose, e.g., to verify that the printed medium actually was printed at a particular place, or by a particular user or at a particular time."

		Rhoads at 35:15–29 • "A shopper's credit card can be Bedoop-encoded so as to lead Bedoop systems of particular stores (i.e., stores pre-authorized by the shopper) to a profile on the shopper (e.g., containing size information, repeat purchase information, return history, style/color preferences, etc.). Credit Card Purchases When a consumer visits a commercial web site and wishes to purchase a displayed product, the transaction can be speeded simply by presenting a Bedoop-encoded credit card to a Bedoop sensor on the user's computer. The Bedoop data on the card leads to a database entry containing the credit card number and expiration date. The Bedoop application then sends this information (optionally after encrypting same) to the web site with instructions to purchase the depicted product." Rhoads at 46:46-53 • "Another function enabled by image sensors in cell phones is user-verification, e.g., by retinal scanning or other optically-sensed biometrics, before the phone will permit a call to be placed. A great number of such biometric verification techniques are known." Rhoads at 50:21-48 • "As more particularly detailed below, the handler 16 provides a response in accordance with the particular watermark payload. The response may be provided directly by the product handler to the device 12, or the handler may respond by communicating with a remote resource 18 (which may be, e.g., a data repository or service provider). In the former case, the handler 16 may identify a URL corresponding to the watermark (using the database 17), and return the URL to the application 28 c. Application 28 c can then pass the URL to a web

		browser 28 b in the device 12, and initiate a link to the internet site identified by the URL. Or the handler may have some locally stored data (e.g., audio or video, or software updates) and send it to the device 12 in response to the watermark. In the latter case, the handler 16 does not respond directly to the device 12. Instead, the handler responds by communicating with a remote resource 18. The communication can be as simple as logging receipt of the watermark message in a remote repository. Or it can be to authenticate device 12 (or a user thereof) to a remote resource in anticipation of a further transaction (e.g., the communication can form part of an on-line licensing or digital rights management transaction). Or the communication can request the remote resource to provide data or a service back to device 12 or to another destination (e.g., to initiate an FTP file transfer, or to request that a song selection identified by the watermark be downloaded to a user's personal music library, or to update software installed on device 12). In still other cases, hybrids of the two foregoing cases can be employed, e.g., handler 16 can send some data back to device 12, while also communicating with a remote resource 18. In some cases, the response returned to the device 12 by handler 6 (or a remote resource 18) can serve to trigger some further action by the device 12. For example, the response returned to device 12 can include a WindowsMedia audio file, together with a request that the device 12 launch the WindowsMedia player installed on the device. (The launching of a browser pointed to a URL is another example of such triggering.)" To the extent it is found that Carlton does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Carlton does not anticipate this claim, Carlton renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final

		Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Carlton with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

5. Claim 38

Claim	'897 Claim Recitation	Exemplary Citations in Reference
38	The method of claim 25, wherein the step of enabling performance of the interaction includes initiating a transaction with an account.	Rhoads discloses the step of enabling performance of the interaction includes initiating a transaction with an account. Rhoads at 50:24-48 • "In the latter case, the handler 16 does not respond directly to the device 12. Instead, the handler responds by communicating with a remote resource 18. The communication can be as simple as logging receipt of the watermark message in a remote repository. Or it can be to authenticate device 12 (or a user thereof) to a remote resource in anticipation of a further transaction (e.g., the communication can form part of an on-line licensing or digital rights management transaction). Or the communication can request the remote resource to provide data or a service back to device 12 or to another destination (e.g., to initiate an FTP file transfer, or to request that a song selection identified by the watermark be downloaded to a user's personal music library, or to update software installed on device 12). In still other cases, hybrids of the two foregoing cases can be employed, e.g., handler 16 can send some data back to device 12, while also communicating with a remote resource 18. In some cases, the response returned to the device 12 by handler 6 (or a

		remote resource 18) can serve to trigger some further action by the device 12. For example, the response returned to device 12 can include a WindowsMedia audio file, together with a request that the device 12 launch the WindowsMedia player installed on the device. (The launching of a browser pointed to a URL is another example of such triggering.)" Rhoads at 30:4-7 • "Bank cards (debit, credit, etc.) can similarly be encoded with Bedoop data to permit the holder to access bank records corresponding to the bank card account. (Entry of a PIN code may be required to assure privacy.)" Rhoads at 48:37-42 • "In accordance with an exemplary application of the below-detailed technology, digital watermarking is employed to convey a plural bit identifier within print media, such as magazine advertisements or articles, direct mail coupons or catalogs, bank- or credit-cards, and business cards." To the extent it is found that Carlton does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Carlton does not anticipate this claim, Carlton renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Carlton with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to

		address similar known problems, according to known methods which would yield predictable results.

6. Claim 39

Claim	'897 Claim Recitation	Exemplary Citations in Reference
39	The method of claim 38, wherein the account includes at least one of the following: an on-line account, an account linked with the portable device, an account linked to a user of the portable device, a bank account, and a credit card account.	Rhoads discloses the account includes at least one of the following: an on-line account, an account linked with the portable device, an account linked to a user of the portable device, a bank account, and a credit card account. Rhoads at 35:15–29 • "A shopper's credit card can be Bedoop-encoded so as to lead Bedoop systems of particular stores (i.e., stores pre-authorized by the shopper) to a profile on the shopper (e.g., containing size information, repeat purchase information, return history, style/color preferences, etc.). Rhoads at 50:24-48 • "In the latter case, the handler 16 does not respond directly to the device 12. Instead, the handler responds by communicating with a remote resource 18. The communication can be as simple as logging receipt of the watermark message in a remote repository. Or it can be to authenticate device 12 (or a user thereof) to a remote resource in anticipation of a further transaction (e.g., the communication can form part of an on-line licensing or digital rights management transaction). Or the communication can request the remote resource to provide data or a service back to device 12 or to another destination (e.g., to initiate an FTP file transfer, or to request that a song selection identified by the watermark be downloaded to a user's personal music library, or to update software installed on device 12). In still other cases, hybrids of the two foregoing cases can be employed, e.g., handler 16 can send some data back to device 12, while also

		communicating with a remote resource 18. In some cases, the response returned to the device 12 by handler 6 (or a remote resource 18) can serve to trigger some further action by the device 12. For example, the response returned to device 12 can include a WindowsMedia audio file, together with a request that the device 12 launch the WindowsMedia player installed on the device. (The launching of a browser pointed to a URL is another example of such triggering.)" Rhoads at 30:4-7 • "Bank cards (debit, credit, etc.) can similarly be encoded with Bedoop data to permit the holder to access bank records corresponding to the bank card account. (Entry of a PIN code may be required to assure privacy.)" Rhoads at 48:37-42 • "In accordance with an exemplary application of the below-detailed technology, digital watermarking is employed to convey a plural bit identifier within print media, such as magazine advertisements or articles, direct mail coupons or catalogs, bank- or credit-cards, and business cards." 48:37-42. To the extent it is found that Carlton does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Carlton does not anticipate this claim, Carlton renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Carlton with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to

		address similar known problems, according to known methods which would yield predictable results.

Nantworks, LLC v. Bank of America Corporation
Central District of California, No. 2:20-cv-7872
Invalidity Chart for U.S. Patent 8,520,897 based on
U.S. Patent No. 6,036,086 to Sizer ("Sizer")

Exhibit H-5

Bank of America Corporation ("BofA" or "Defendant") provides this chart subject to all reservations, objections, statements, and disclaimers set forth herein and in BofA's Final Invalidity Contentions Cover Pleading, as well as any amendment, supplement, or modification thereof, which are incorporated herein by reference in their entirety.

On information and belief, the application that matured into Sizer was filed on March 28, 1997. Sizer published at least as of March 14, 2000, and is available as prior art at least under 35 U.S.C. § 102(a), (b) and (e).

As shown in the chart below, Sizer anticipates asserted claims 25, 30, 33, 34, and39 of the '897 Patent. To the extent Sizer is found not to expressly disclose certain limitations in the asserted claims, such limitations are inherent. To the extent Sizer is found not to anticipate any asserted claims or claim elements of the '897 Patent, the reference nevertheless renders those claims or claim elements obvious under 35 U.S.C. § 103, either alone or in combination with other art identified herein or in the cover pleading of BofA's Final Invalidity Contentions, for the reasons identified in the cover pleading. These Final Invalidity Contentions are not an admission by BofA that the accused products, including any current or past versions of the accused products, are covered by, or infringe the asserted claims, but are based instead on the recognition that if the claims are interpreted to be broad enough to encompass the accused products, the claims must also be construed to have that same scope when considering whether they are invalid.

The following chart is based on the Court's claim construction rulings (Dkt. Nos. 145, 153, 236), and to the extent not contradictory, Plaintiffs' apparent interpretations of the claims in Plaintiffs' "Preliminary" Final Infringement Contentions. BofA notes that in many instances, Plaintiffs' "Preliminary" Final Infringement Contentions ignore and/or contradict the Court's claim construction rulings. BofA does not accept the assumptions concerning the scope and meaning implicit in Plaintiffs' Preliminary Final Infringement Contentions, to the extent those assumptions are discernible, and reserves the right to challenge Plaintiffs' proposed (or implied) constructions. Furthermore, Plaintiffs' "Preliminary" Final Contentions fail to comply with the level of specificity required by SPR 2.1.2 and 2.1.3, and include improper reservation of rights to supplement with "additional ways in which BoA Accused Products infringe." Any alleged deficient discovery is of Plaintiffs own making. Most if not all of the deficiencies in Plaintiffs' Preliminary Final Contentions have been documented by Defendants through multiple letters since at least May 18, 2021 and up to Bank of America's most recent letter of August 19, 2023. Plaintiffs' shifting infringement contentions, deficiencies in its infringement contentions, and changing claim scope have prejudiced BofA from being able to reasonably prepare its defenses, including its invalidity contentions. BofA also reserves the right to revise and supplement these charts if and when Plaintiffs are permitted to provide revised or supplemental Infringement Contentions. Citations given in the chart below are merely representative of the respective elements to provide notice to Plaintiffs of BofA's invalidity theories and are not meant to be exhaustive.

1. Claim 25

Element	'897 Claim Recitation	Exemplary Citations in Reference
[1Pre]	A method of conducting a transaction with an interactive system, the method comprising	Sizer discloses a method of conducting a transaction with an interactive system. Sizer at Abstract: • In accordance with the present invention, an apparatus initiates a transaction and includes a capture device for capturing transaction data from marks contained on an object. The marks have a code corresponding to transaction data for initiating a transaction. The capture device includes a scanner, operable by a user, for reading marks contained on the object. A controller interprets the marks and retrieves the transaction data embedded in the marks. The capture device originates a telephone call and transfers at least a portion of the transaction data to a desired destination for initiating a transaction. Sizer at 12:62-13:12 • In accordance with the present invention, a capture device now can initiate a transaction by capturing transaction data from marks contained on an object, such as a magazine advertisement, where the marks have embedded therein a code corresponding to transaction data for initiating a transaction. The capture device includes a scanner, operable by a user, for reading marks contained on the object and a controller for interpreting the marks and retrieving transaction data embedded in the marks. Through associated apparatus as described above, a telephone call can be originated by information contained in the transaction data in at least a portion of the transaction data transferred to a desired destination for initiating a transaction. As a result, it is now possible now for magazines, newspapers, printed media and many other objects to contain marks relating to an advertisement or solicitation which can be easily scanned and a transaction

Element	'897 Claim Recitation	Exemplary Citations in Reference
		completed to the vendor of the advertisement or solicitation. This saves much time and allows easier access to services and products.
[1A]	providing access to a device having a display;	Sizer discloses providing access to a device having a display. Sizer at FIG. 1 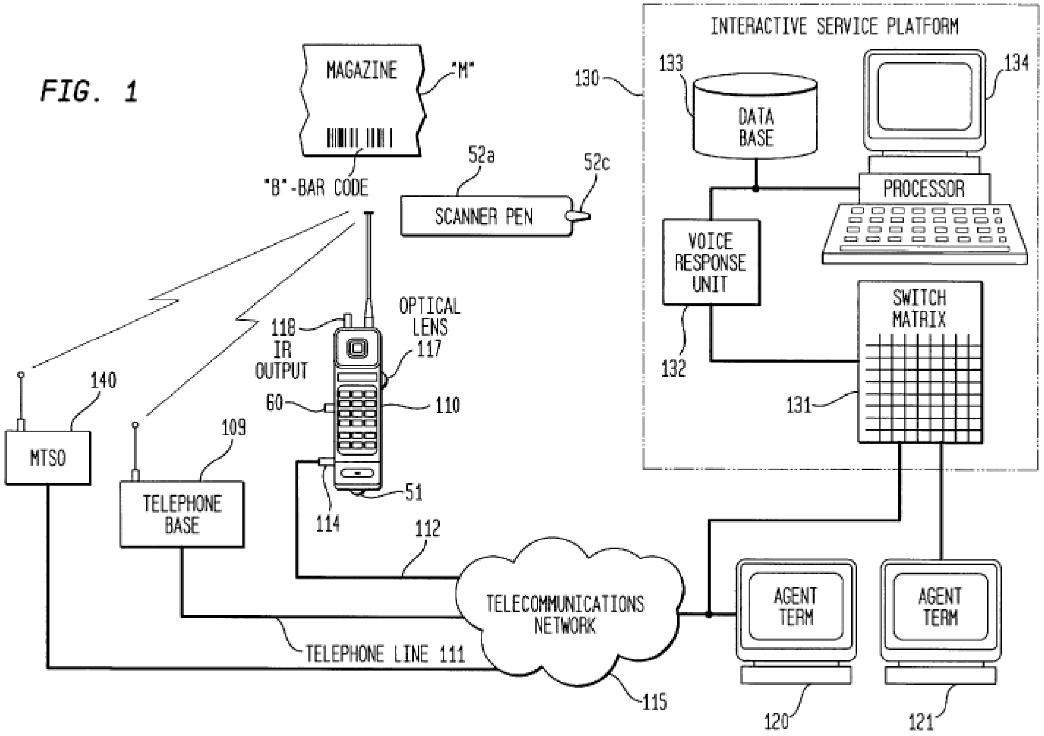 FIG. 1 Sizer at 5:3-20 - "Briefly, it was found using a spectrum analyzer on a typical video signal, that there are comparatively large frequency components at the line rate and at the frame rate and its harmonics, but that between these frequencies,

Element	'897 Claim Recitation	Exemplary Citations in Reference
		there are other frequency bands in which little information is carried. One such frequency band is between 15 and 30 kHz. By adding a low level carrier signal or tone at a frequency in this band, say 25 kHz, the video image is not degraded, but a properly tuned decoder can receive and decode the encoded information. In this way, digital information can be subliminally inserted in a video signal by adding to the video signal an amplitude shift keyed (ASK) or frequency shift keyed (FSK) carrier signal, and the digital information can later be recovered using ASK or FSK decoding. Encoded data can also be inserted in a television signal and recovered by a received that responds to the picture displayed on the television, in the manner described in U.S. Pat. No. 4,807,031 cited above." Sizer at 5:23-26 • "If any non-discernible encoded data is part of a video signal, such sensing takes the form of light from display area being collected by an optical lens 117." Sizer at 6:42-47 • "If the encoded non-perceptible data is inserted in the video portion of a television signal, the images displayed on a visual display area a monitor or television are captured by optical lens 117 and an associated photo-detector 203, which is arranged to supply an electrical signal to a filter/receiver 215 representing the image." Sizer at 9:33-53 • "For background purposes, a televised advertisement for a product could include additional ordering or "coupon" information embedded in a subliminal visual pattern within the video image displayed on a television. The coupon information describes the item(s) offered for sale, the regular

Element	'897 Claim Recitation	Exemplary Citations in Reference
		price, the coupon value, and contains additional information needed for ordering. This embedded information could be transmitted between a television and capture device at a rate of about 60 bits per second, that is sufficient such that the necessary information can be transmitted in a very brief period of time. In the application just described, the embedded product information, dialing number, coupon information and necessary checkbits together comprise approximately 2,000 bits of information, so that between three and four seconds would be required for transmission. Once the coupon information is stored in capture device 110, the user of the capture device can supplement the coupon information with user entered data which can be entered using keypad 247. This supplemental information may include credit card or other billing information pertaining to the user, as an example." Sizer at 10:27-37 • In the active mode, a user operates capture device 110 in real time, typically while scanning a product or bar code, or in more complicated devices viewing a video display or listening to an audio presentation. Scanning a product or bar code, or in more complicated devices the user activates the SAVE button 241. When button 241 is activated, a positive result occurs in step 503, thereby capturing the encoded data, displaying all or a portion of the data on display 225, and storing same in memory 221, all in step 507. Until the SAVE button is activated, capture device 110 may be in a "wait state" 505." To the extent it is found that Sizer does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Sizer does not anticipate this claim, Sizer renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have

Element	'897 Claim Recitation	Exemplary Citations in Reference
		been motivated to combine Sizer with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1B]	displaying a displayed image on the display of the device;	Sizer discloses displaying a displayed image on the display of the device. Sizer at FIG. 1 FIG. 1 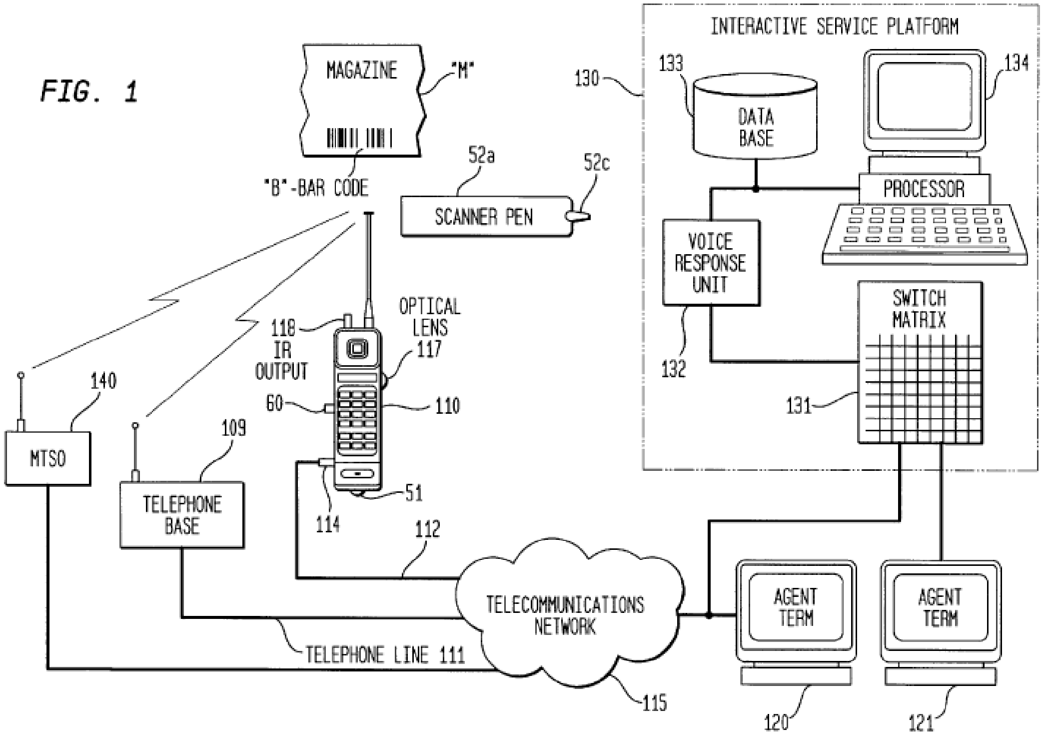

Element	'897 Claim Recitation	Exemplary Citations in Reference
		Sizer at 2:60-62 • At least a portion of the transaction data can be displayed on a LCD or similar type of display. Sizer at 5:3-20 • "Briefly, it was found using a spectrum analyzer on a typical video signal, that there are comparatively large frequency components at the line rate and at the frame rate and its harmonics, but that between these frequencies, there are other frequency bands in which little information is carried. One such frequency band is between 15 and 30 kHz. By adding a low level carrier signal or tone at a frequency in this band, say 25 kHz, the video image is not degraded, but a properly tuned decoder can receive and decode the encoded information. In this way, digital information can be subliminally inserted in a video signal by adding to the video signal an amplitude shift keyed (ASK) or frequency shift keyed (FSK) carrier signal, and the digital information can later be recovered using ASK or FSK decoding. Encoded data can also be inserted in a television signal and recovered by a received that responds to the picture displayed on the television, in the manner described in U.S. Pat. No. 4,807,031 cited above." Sizer at 5:23-26 • "If any non-discernible encoded data is part of a video signal, such sensing takes the form of light from display area being collected by an optical lens 117." Sizer at 5:36-48 • In addition to receiving non-discernible encoded data from the scanner pen, the capture device 110 has the capability of (a) storing the encoded data, (b) transmitting (or otherwise outputting) the encoded data (or

Element	'897 Claim Recitation	Exemplary Citations in Reference
		information retrieved using such data) to a remote location or device, and (c) displaying the encoded data to a user. The data may include information for initiating a transaction or originating a telephone call, as well as additional information (e.g., price and ordering information) associated with the scanned material. The remote location can be agent terminals 120 or 121, interactive service platform 130, or a point of sale system, and the communication between the capture device 110 and the remote location may be one-way or two-way. Sizer at 6:42-47 • "If the encoded non-perceptible data is inserted in the video portion of a television signal, the images displayed on a visual display area a monitor or television are captured by optical lens 117 and an associated photo-detector 203, which is arranged to supply an electrical signal to a filter/receiver 215 representing the image." Sizer at 7:12-16 • The captured data may also be displayed on a display 225, such as a liquid crystal display, so that a user will be aware of exactly what information was captured. Sizer at 9:14-19 • During the SEND DATA function, the information being output may also be displayed on display 225. This allows a person using the capture device to identify the information that was captured, and use the information in an actual transaction. Sizer at 9:33-53 • "For background purposes, a televised advertisement for a product could include additional ordering or "coupon" information embedded in a

Element	'897 Claim Recitation	Exemplary Citations in Reference
		subliminal visual pattern within the video image displayed on a television. The coupon information describes the item(s) offered for sale, the regular price, the coupon value, and contains additional information needed for ordering. This embedded information could be transmitted between a television and capture device at a rate of about 60 bits per second, that is sufficient such that the necessary information can be transmitted in a very brief period of time. In the application just described, the embedded product information, dialing number, coupon information and necessary checkbits together comprise approximately 2,000 bits of information, so that between three and four seconds would be required for transmission. Once the coupon information is stored in capture device 110, the user of the capture device can supplement the coupon information with user entered data which can be entered using keypad 247. This supplemental information may include credit card or other billing information pertaining to the user, as an example." Sizer at 10:27-37 • In the active mode, a user operates capture device 110 in real time, typically while scanning a product or bar code, or in more complicated devices viewing a video display or listening to an audio presentation. Scanning a product or bar code, or in more complicated devices the user activates the SAVE button 241. When button 241 is activated, a positive result occurs in step 503, thereby capturing the encoded data, displaying all or a portion of the data on display 225, and storing same in memory 221, all in step 507. Until the SAVE button is activated, capture device 110 may be in a "wait state" 505." To the extent it is found that Sizer does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Sizer does not anticipate this claim, Sizer renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination

Element	'897 Claim Recitation	Exemplary Citations in Reference
		with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Sizer with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1C]	presenting the displayed image proximate to an optical sensor; and	Sizer discloses presenting the displayed image proximate to an optical sensor. Sizer at FIG. 1

Element	'897 Claim Recitation	Exemplary Citations in Reference
		FIG. 1 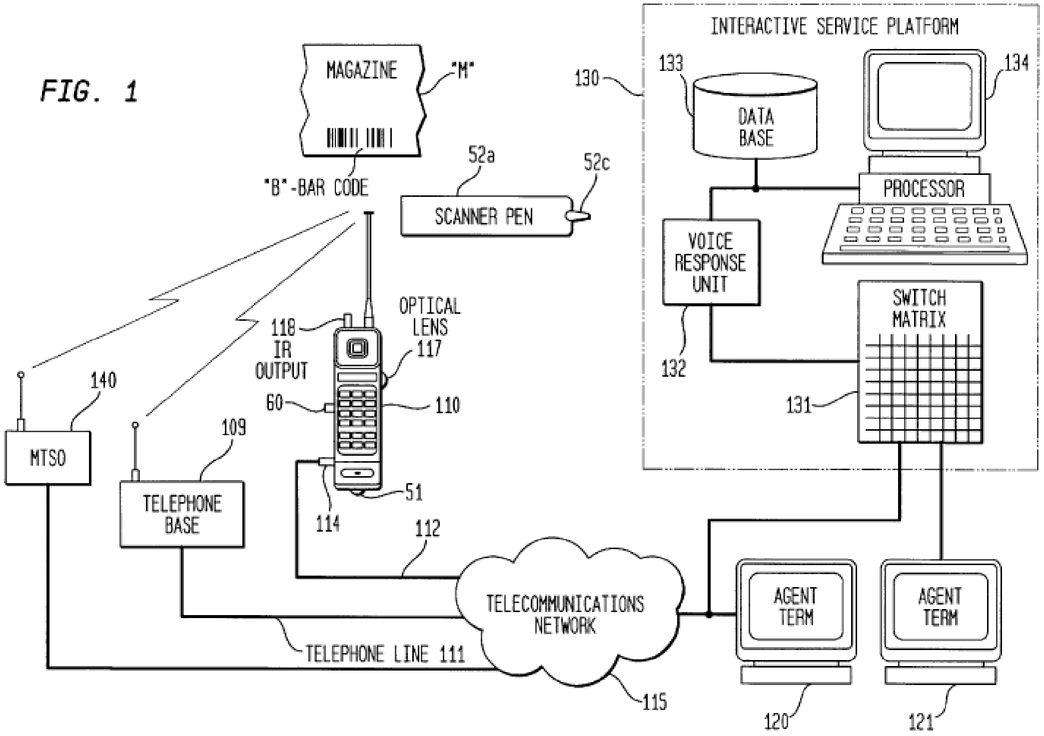 Sizer at 5:3-20 - "Briefly, it was found using a spectrum analyzer on a typical video signal, that there are comparatively large frequency components at the line rate and at the frame rate and its harmonics, but that between these frequencies, there are other frequency bands in which little information is carried. One such frequency band is between 15 and 30 kHz. By adding a low level carrier signal or tone at a frequency in this band, say 25 kHz, the video image is not degraded, but a properly tuned decoder can receive and decode the encoded information. In this way, digital information can be

Element	'897 Claim Recitation	Exemplary Citations in Reference
		subliminally inserted in a video signal by adding to the video signal an amplitude shift keyed (ASK) or frequency shift keyed (FSK) carrier signal, and the digital information can later be recovered using ASK or FSK decoding. Encoded data can also be inserted in a television signal and recovered by a received that responds to the picture displayed on the television, in the manner described in U.S. Pat. No. 4,807,031 cited above." Sizer at 5:23-26 • "If any non-discernible encoded data is part of a video signal, such sensing takes the form of light from display area being collected by an optical lens 117." Sizer at 6:41-62 • "If the encoded non-perceptible data is inserted in the video portion of a television signal, the images displayed on a visual display area a monitor or television are captured by optical lens 117 and an associated photo-detector 203, which is arranged to supply an electrical signal to a filter/receiver 215 representing the image. On the other hand, if the encoded non-perceptible data is inserted in the audio portion of a television signal, e.g. the output from a speaker is captured by a microphone 211 and an associated amplifier 213, which is arranged to supply an electrical signal to filter/receiver 215 representing the sound energy. In either event, the encoded non-perceptible data is decoded in the filter/receiver 215 in a manner consistent with the manner in which the original data was encoded. Thus, filter/receiver 215 can be arranged to perform the same functionality as the elements illustrated in FIG. 4 of the above referenced Broughton et al. patent, or, alternatively, the operation of filter/receiver 215 can be as described in conjunction with FIGS. 4 or 7 in the copending application of T. Sizer."

Element	'897 Claim Recitation	Exemplary Citations in Reference
		Sizer at 9:33-53 • "For background purposes, a televised advertisement for a product could include additional ordering or "coupon" information embedded in a subliminal visual pattern within the video image displayed on a television. The coupon information describes the item(s) offered for sale, the regular price, the coupon value, and contains additional information needed for ordering. This embedded information could be transmitted between a television and capture device at a rate of about 60 bits per second, that is sufficient such that the necessary information can be transmitted in a very brief period of time. In the application just described, the embedded product information, dialing number, coupon information and necessary checkbits together comprise approximately 2,000 bits of information, so that between three and four seconds would be required for transmission. Once the coupon information is stored in capture device 110, the user of the capture device can supplement the coupon information with user entered data which can be entered using keypad 247. This supplemental information may include credit card or other billing information pertaining to the user, as an example." Sizer at 10:27-37 • In the active mode, a user operates capture device 110 in real time, typically while scanning a product or bar code, or in more complicated devices viewing a video display or listening to an audio presentation. Scanning a product or bar code, or in more complicated devices the user activates the SAVE button 241. When button 241 is activated, a positive result occurs in step 503, thereby capturing the encoded data, displaying all or a portion of the data on display 225, and storing same in memory 221, all in step 507. Until the SAVE button is activated, capture device 110 may be in a "wait state" 505."

Element	'897 Claim Recitation	Exemplary Citations in Reference
		To the extent it is found that Sizer does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Sizer does not anticipate this claim, Sizer renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Sizer with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1D]	enabling, via the interactive system, identification of features from the displayed image, recognition of a target based on the features, association of the target with target information pertinent to the target, and performance of a transaction based on the target information.	Sizer discloses enabling, via the interactive system, identification of features from the displayed image, recognition of a target based on the features, association of the target with target information pertinent to the target, and performance of a transaction based on the target information. Sizer at FIG. 2

Element	'897 Claim Recitation	Exemplary Citations in Reference
		FIG. 2 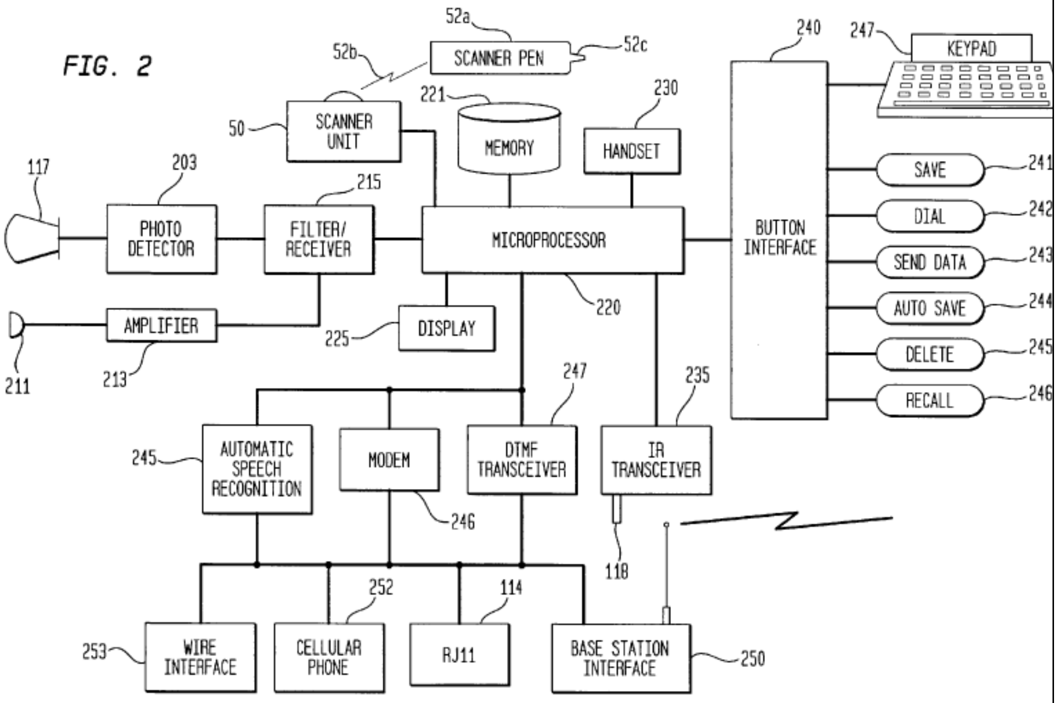 Sizer at 2:36-42 • It would be advantageous if select pieces of information concerning the advertisement or solicitation could be retrieved to initiate a transaction by dialing a telephone number concerning the vendor of that particular advertisement or solicitation and then initiating a transaction automatically. Sizer at 4:45-64 • It should be understood that the data could also be received from a light emitting device such as a television or audio signal from a speaker. In

Element	'897 Claim Recitation	Exemplary Citations in Reference
		that instance, in a television signal would be received by a television via over-the-air transmission, as from transmission tower via a transmission link, or alternatively supplied by a cable TV connection, or from a VCR/video-tape. The video portion of the television signal could be displayed on the picture tube or other visual display area of a monitor, computer or a television, while the audio portion of the television signal is played from speaker. Associated with the program that is seen and possibly heard by a viewer/listener could be non-discernible encoded data that is transmitted as part of the video and possibly audio signal. For example, the program may be an advertisement for a service or product, and the indiscernible data may create a subliminal visual pattern on the visual display area, that when properly received and decoded, contains information needed to order those services or products, including prices, delivery intervals, shipping details, coupon offers, and so-on. Sizer at 6:18-38 • "If an outgoing telephone call in initiated by capture device 110, the call can illustratively be routed through telecommunications network 115 directly to an agent terminal 120, or to an interactive service platform 130, which is connected to another agent terminal 121. When a connection is completed to platform 130, the call may be routed by a switch matrix 131 to a voice response unit 132 that operates under the control of a processor 134. Information included in the captured encoded can be provided to the attendants, or processor 134 can perform a look-up and retrieval operation in database 131, using encoded data as a key, in order to effectuate a desired transaction. For certain transactions, information connected with a transaction may also be transmitted from service platform 130 to capture device 110, either for display to a user, storage for later use, or output to a point of sale system or device. This information may, for example, be confirmation information, indicating that a transaction has been effected,

Element	'897 Claim Recitation	Exemplary Citations in Reference
		or details regarding a just completed transaction, such as ticket information relating to a travel reservation." Sizer at 6:41-62 • "If the encoded non-perceptible data is inserted in the video portion of a television signal, the images displayed on a visual display area a monitor or television are captured by optical lens 117 and an associated photo-detector 203, which is arranged to supply an electrical signal to a filter/receiver 215 representing the image. On the other hand, if the encoded non-perceptible data is inserted in the audio portion of a television signal, e.g. the output from a speaker is captured by a microphone 211 and an associated amplifier 213, which is arranged to supply an electrical signal to filter/receiver 215 representing the sound energy. In either event, the encoded non-perceptible data is decoded in the filter/receiver 215 in a manner consistent with the manner in which the original data was encoded. Thus, filter/receiver 215 can be arranged to perform the same functionality as the elements illustrated in FIG. 4 of the above referenced Broughton et al. patent, or, alternatively, the operation of filter/receiver 215 can be as described in conjunction with FIGS. 4 or 7 in the copending application of T. Sizer." Sizer at 10:45-49 • "A call is then originated launch, such by transmitting the dialing information from memory 221 through base station interface 250 to base station 109 and thence to telephone line 111 and telecommunications network 115 to interactive service platform 130." Sizer at 12:41-43 • "The call may be placed to an interactive service platform, and additional captured data, such as coupon or similar information, can then be used to

Element	'897 Claim Recitation	Exemplary Citations in Reference
		effectuate a transaction." Sizer at 12:62-13:12 In accordance with the present invention, a capture device now can initiate a transaction by capturing transaction data from marks contained on an object, such as a magazine advertisement, where the marks have embedded therein a code corresponding to transaction data for initiating a transaction. The capture device includes a scanner, operable by a user, for reading marks contained on the object and a controller for interpreting the marks and retrieving transaction data embedded in the marks. Through associated apparatus as described above, a telephone call can be originated by information contained in the transaction data in at least a portion of the transaction data transferred to a desired destination for initiating a transaction. As a result, it is now possible now for magazines, newspapers, printed media and many other objects to contain marks relating to an advertisement or solicitation which can be easily scanned and a transaction completed to the vendor of the advertisement or solicitation. This saves much time and allows easier access to services and products. To the extent it is found that Sizer does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Sizer does not anticipate this claim, Sizer renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Sizer with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

2. Claim 30

Claim	'897 Claim Recitation	Exemplary Citations in Reference
30	The method of claim 25, wherein the step of displaying the displayed image includes displaying a time varying image.	Sizer discloses the step of displaying the displayed image includes displaying a time varying image. Sizer at 5:3-20 • "Briefly, it was found using a spectrum analyzer on a typical video signal, that there are comparatively large frequency components at the line rate and at the frame rate and its harmonics, but that between these frequencies, there are other frequency bands in which little information is carried. One such frequency band is between 15 and 30 kHz. By adding a low level carrier signal or tone at a frequency in this band, say 25 kHz, the video image is not degraded, but a properly tuned decoder can receive and decode the encoded information. In this way, digital information can be subliminally inserted in a video signal by adding to the video signal an amplitude shift keyed (ASK) or frequency shift keyed (FSK) carrier signal, and the digital information can later be recovered using ASK or FSK decoding. Encoded data can also be inserted in a television signal and recovered by a received that responds to the picture displayed on the television, in the manner described in U.S. Pat. No. 4,807,031 cited above." Sizer at 5:23-26 • "If any non-discernible encoded data is part of a video signal, such sensing takes the form of light from display area being collected by an optical lens 117." Sizer at 6:42-47 • "If the encoded non-perceptible data is inserted in the video portion of a television signal, the images displayed on a visual display area a monitor or television are captured by optical lens 117 and an associated photo-

		detector 203, which is arranged to supply an electrical signal to a filter/receiver 215 representing the image." Sizer at 9:33-53 • "For background purposes, a televised advertisement for a product could include additional ordering or "coupon" information embedded in a subliminal visual pattern within the video image displayed on a television. The coupon information describes the item(s) offered for sale, the regular price, the coupon value, and contains additional information needed for ordering. This embedded information could be transmitted between a television and capture device at a rate of about 60 bits per second, that is sufficient such that the necessary information can be transmitted in a very brief period of time. In the application just described, the embedded product information, dialing number, coupon information and necessary checkbits together comprise approximately 2,000 bits of information, so that between three and four seconds would be required for transmission. Once the coupon information is stored in capture device 110, the user of the capture device can supplement the coupon information with user entered data which can be entered using keypad 247. This supplemental information may include credit card or other billing information pertaining to the user, as an example." Sizer at 10:27-37 • In the active mode, a user operates capture device 110 in real time, typically while scanning a product or bar code, or in more complicated devices viewing a video display or listening to an audio presentation. Scanning a product or bar code, or in more complicated devices the user activates the SAVE button 241. When button 241 is activated, a positive result occurs in step 503, thereby capturing the encoded data, displaying all or a portion of the data on display 225, and storing same in memory 221, all in step 507. Until the SAVE button is activated, capture device

		110 may be in a "wait state" 505." To the extent it is found that Sizer does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Sizer does not anticipate this claim, Sizer renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Sizer with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

3. Claim 33

Claim	'897 Claim Recitation	Exemplary Citations in Reference
33	The method of claim 25, wherein the step of enabling performance of the transaction includes initiating a verification.	Sizer discloses the step of enabling performance of the transaction includes initiating a verification. To the extent it is found that Sizer does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Sizer does not anticipate this claim, Sizer renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Sizer with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success

		for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

4. Claim 34

Claim	'897 Claim Recitation	Exemplary Citations in Reference
34	The method of claim 25 wherein the step of enabling performance of the transaction includes initiating an authorization.	Sizer discloses the step of enabling performance of the transaction includes initiating an authorization. Sizer at 6:18-38 • "If an outgoing telephone call in initiated by capture device 110, the call can illustratively be routed through telecommunications network 115 directly to an agent terminal 120, or to an interactive service platform 130, which is connected to another agent terminal 121. When a connection is completed to platform 130, the call may be routed by a switch matrix 131 to a voice response unit 132 that operates under the control of a processor 134. Information included in the captured encoded can be provided to the attendants, or processor 134 can perform a look-up and retrieval operation in database 131, using encoded data as a key, in order to effectuate a desired transaction. For certain transactions, information connected with a transaction may also be transmitted from service platform 130 to capture device 110, either for display to a user, storage for later use, or output to a point of sale system or device. This information may, for example, be confirmation information, indicating that a transaction has been effected, or details regarding a just completed transaction, such as ticket information relating to a travel reservation." To the extent it is found that Sizer does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Sizer does not anticipate

		this claim, Sizer renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Sizer with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

5. Claim 38

Claim	'897 Claim Recitation	Exemplary Citations in Reference
38	The method of claim 25, wherein the step of enabling performance of the interaction includes initiating a transaction with an account.	Sizer discloses the step of enabling performance of the interaction includes initiating a transaction with an account. To the extent it is found that Sizer does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Sizer does not anticipate this claim, Sizer renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Sizer with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

6. Claim 39

Claim	'897 Claim Recitation	Exemplary Citations in Reference
39	The method of claim 38, wherein the account includes at least one of the following: an on-line account, an account linked with the portable device, an account linked to a user of the portable device, a bank account, and a credit card account.	Sizer discloses the account includes at least one of the following: an on-line account, an account linked with the portable device, an account linked to a user of the portable device, a bank account, and a credit card account. Sizer at FIG. 1 FIG. 1 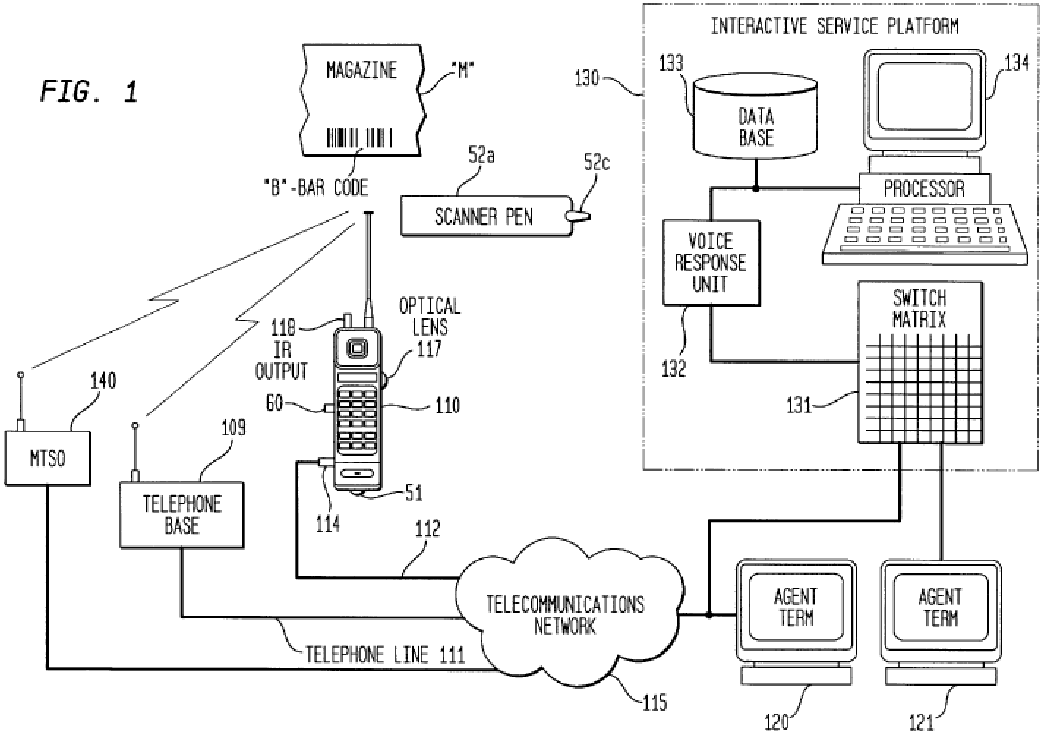 Sizer at 6:18-38

		• "If an outgoing telephone call in initiated by capture device 110, the call can illustratively be routed through telecommunications network 115 directly to an agent terminal 120, or to an interactive service platform 130, which is connected to another agent terminal 121. When a connection is completed to platform 130, the call may be routed by a switch matrix 131 to a voice response unit 132 that operates under the control of a processor 134. Information included in the captured encoded can be provided to the attendants, or processor 134 can perform a look-up and retrieval operation in database 131, using encoded data as a key, in order to effectuate a desired transaction. For certain transactions, information connected with a transaction may also be transmitted from service platform 130 to capture device 110, either for display to a user, storage for later use, or output to a point of sale system or device. This information may, for example, be confirmation information, indicating that a transaction has been effected, or details regarding a just completed transaction, such as ticket information relating to a travel reservation." Sizer at 9:33-53 • "For background purposes, a televised advertisement for a product could include additional ordering or "coupon" information embedded in a subliminal visual pattern within the video image displayed on a television. The coupon information describes the item(s) offered for sale, the regular price, the coupon value, and contains additional information needed for ordering. This embedded information could be transmitted between a television and capture device at a rate of about 60 bits per second, that is sufficient such that the necessary information can be transmitted in a very brief period of time. In the application just described, the embedded product information, dialing number, coupon information and necessary checkbits together comprise approximately 2,000 bits of information, so that between three and four seconds would be required for transmission. Once the coupon

		information is stored in capture device 110, the user of the capture device can supplement the coupon information with user entered data which can be entered using keypad 247. This supplemental information may include credit card or other billing information pertaining to the user, as an example." Sizer at 10:45-49 • "A call is then originated launch, such by transmitting the dialing information from memory 221 through base station interface 250 to base station 109 and thence to telephone line 111 and telecommunications network 115 to interactive service platform 130." Sizer at 12:41-43 • "The call may be placed to an interactive service platform, and additional captured data, such as coupon or similar information, can then be used to effectuate a transaction." To the extent it is found that Sizer does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Sizer does not anticipate this claim, Sizer renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Sizer with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

Nantworks, LLC v. Bank of America Corporation
Central District of California, No. 2:20-cv-7872
Invalidity Chart for U.S. Patent 8,520,897 based on
U.S. Patent No. 6,666,377 to Harris ("Harris")

Exhibit H-6

Bank of America Corporation ("BofA" or "Defendant") provides this chart subject to all reservations, objections, statements, and disclaimers set forth herein and in BofA's Final Invalidity Contentions Cover Pleading, as well as any amendment, supplement, or modification thereof, which are incorporated herein by reference in their entirety.

On information and belief, the application that matured into Harris was filed on July 18, 2000. Harris published at least as of December 23, 2003, and is available as prior art at least under 35 U.S.C. § 102(e).

As shown in the chart below, Harris anticipates asserted claims 25, 30,, 33, 34, and 39 of the '897 Patent. To the extent Harris is found not to expressly disclose certain limitations in the asserted claims, such limitations are inherent. To the extent Harris is found not to anticipate any asserted claims or claim elements of the '897 Patent, the reference nevertheless renders those claims or claim elements obvious under 35 U.S.C. § 103, either alone or in combination with other art identified herein or in the cover pleading of BofA's Final Invalidity Contentions, for the reasons identified in the cover pleading. These Final Invalidity Contentions are not an admission by BofA that the accused products, including any current or past versions of the accused products, are covered by, or infringe the asserted claims, but are based instead on the recognition that if the claims are interpreted to be broad enough to encompass the accused products, the claims must also be construed to have that same scope when considering whether they are invalid.

The following chart is based on the Court's claim construction rulings (Dkt. Nos. 145, 153, 236), and to the extent not contradictory, Plaintiffs' apparent interpretations of the claims in Plaintiffs' "Preliminary" Final Infringement Contentions. BofA notes that in many instances, Plaintiffs' "Preliminary" Final Infringement Contentions ignore and/or contradict the Court's claim construction rulings. BofA does not accept the assumptions concerning the scope and meaning implicit in Plaintiffs' Preliminary Final Infringement Contentions, to the extent those assumptions are discernible, and reserves the right to challenge Plaintiffs' proposed (or implied) constructions. Furthermore, Plaintiffs' "Preliminary" Final Contentions fail to comply with the level of specificity required by SPR 2.1.2 and 2.1.3, and include improper reservation of rights to supplement with "additional ways in which BoA Accused Products infringe." Any alleged deficient discovery is of Plaintiffs own making. Most if not all of the deficiencies in Plaintiffs' Preliminary Final Contentions have been documented by Defendants through multiple letters since at least May 18, 2021 and up to Bank of America's most recent letter of August 19, 2023. Plaintiffs' shifting infringement contentions, deficiencies in its infringement contentions, and changing claim scope have prejudiced BofA from being able to reasonably prepare its defenses, including its invalidity contentions. BofA also reserves the right to revise and supplement these charts if and when Plaintiffs are permitted to provide revised or supplemental Infringement Contentions. Citations given in the chart below are merely representative of the respective elements to provide notice to Plaintiffs of BofA's invalidity theories and are not meant to be exhaustive.

1. Claim 25

Element	'897 Claim Recitation	Exemplary Citations in Reference
[1Pre]	A method of conducting a transaction with an interactive system, the method comprising	Harris discloses a method: Harris at 2: 3-19 • "The bar codes can be imaged/scanned in a number of different ways. One embodiment discloses using a camera to input and decode these bar codes. The embodiments are shown in FIGS. 1A and 1B. FIG. 1A describes using a personal digital assistant 100 as the input device. One preferred input device of this type is the Palm V(TM) type hand held computer. A bar code scanner can be used, such as the commercially available Symbol Technology SPT 1700. Alternatively, a camera add-on unit can be added to the Palm V and used as described herein. The PDA includes a screen 145 and user interface 146 all within the same housing 99. FIG. 1B shows the client being a cellular telephone which also may include a screen 145 and user interface 150. The cellular telephone can be associated with either a dedicated bar code scanner or an image sensor 160 of the type used to obtain photographs for video telephony."
[1A]	providing access to a device having a display;	Harris discloses providing access to a device having a display. Harris at FIG. 1A and FIG. 1B

Element	'897 Claim Recitation	Exemplary Citations in Reference
		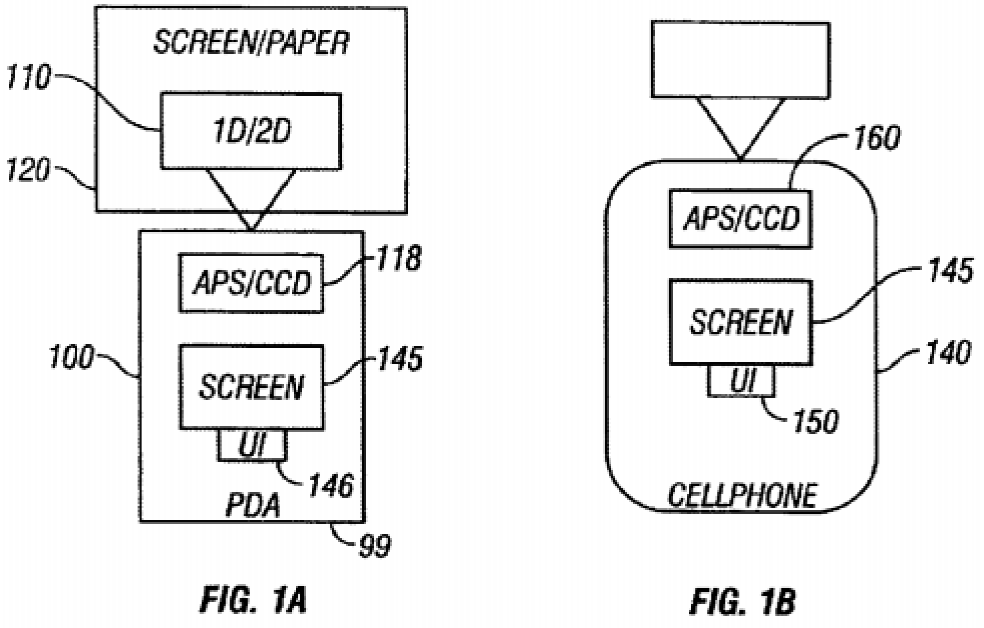 FIG. 1A FIG. 1B Harris at 2: 3-19 • "The bar codes can be imaged/scanned in a number of different ways. One embodiment discloses using a camera to input and decode these bar codes. The embodiments are shown in FIGS. 1A and 1B. FIG. 1A describes using a personal digital assistant 100 as the input device. One preferred input device of this type is the Palm V(TM) type hand held computer. A bar code scanner can be used, such as the commercially available Symbol Technology SPT 1700. Alternatively, a camera add-on unit can be added to the Palm V and used as described herein. The PDA includes a screen 145 and user interface 146 all within the same housing 99. FIG. 1B shows the client being a cellular telephone which also may include a screen 145 and user interface 150. The cellular telephone can be

Element	'897 Claim Recitation	Exemplary Citations in Reference
		associated with either a dedicated bar code scanner or an image sensor 160 of the type used to obtain photographs for video telephony." Harris at FIG. 4 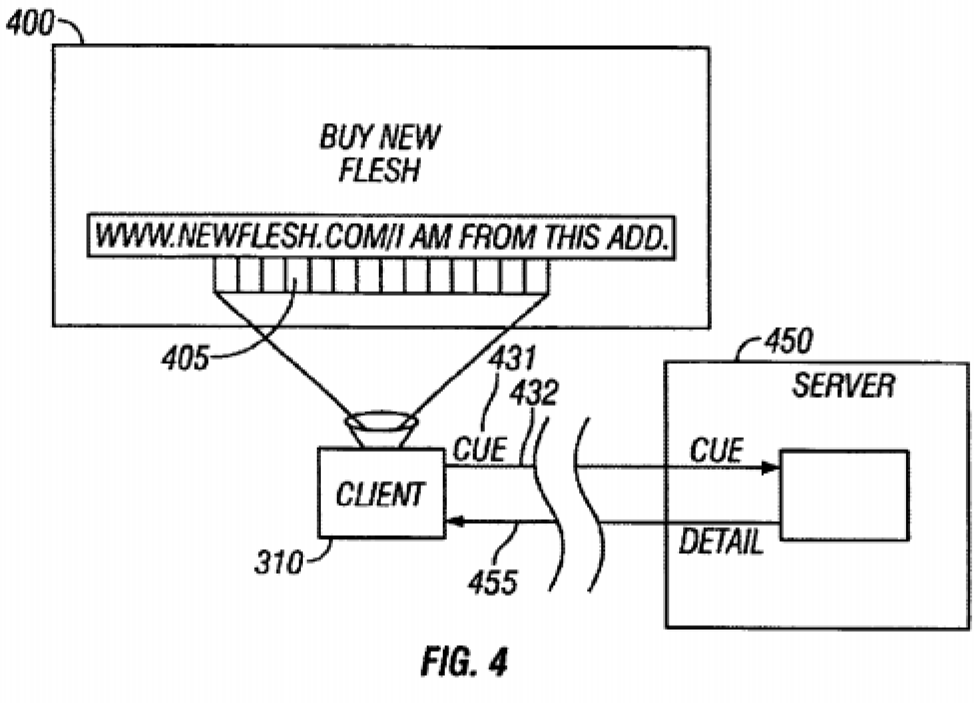 FIG. 4 Harris at 4:15-30 - "The advertisement 400 includes a bar code 405 therein. The bar code 405 is associated with the advertisement, and includes some additional information about the advertisement. For example, the bar code may include the web site address of the company preparing the advertisement, or appointment information about the advertisement, or a "vcf file" or the

Element	'897 Claim Recitation	Exemplary Citations in Reference
		like which is an importable file with address information about the company or author sponsoring the advertisement. In operation, the user brings one of the clients, either the cellular telephone 310 or PDA 315, into range of the bar code 405. The client reads the bar code and decodes it as noted above. The decoded information can represent ASCII information, a compressed file such as a zip file, G code information, or any other compressed or non-compressed information." To the extent it is found that Harris does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Harris does not anticipate this claim, Harris renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Harris with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1B]	displaying a displayed image on the display of the device;	Harris discloses displaying a displayed image on the display of the device. Harris at 4: 9-15 • "Another embodiment shown in FIG. 4 uses bar codes to enter information into the computer. The FIG. 4 embodiment stores scannable non-alphanumeric information, e.g. bar code information, as some part of a communication—here an advertisement. The advertisement can be a

Element	**'897 Claim Recitation**	**Exemplary Citations in Reference**
		print advertisement, a television advertisement, or an Internet advertisement for example." To the extent it is found that Harris does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Harris does not anticipate this claim, Harris renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Harris with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1C]	presenting the displayed image proximate to an optical sensor; and	Harris discloses presenting the displayed image proximate to an optical sensor. Harris at FIG. 1A and FIG. 1B

Element	'897 Claim Recitation	Exemplary Citations in Reference
		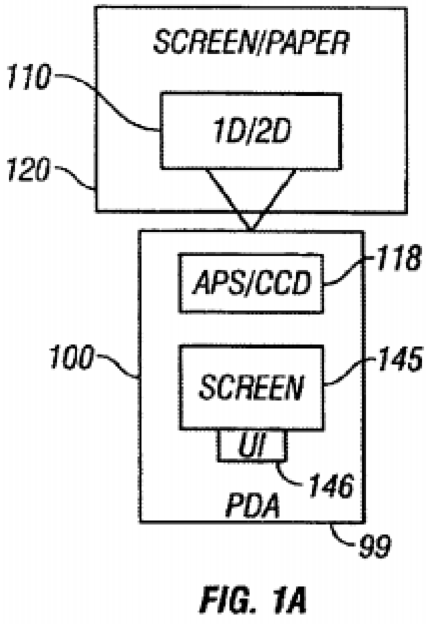 FIG. 1A 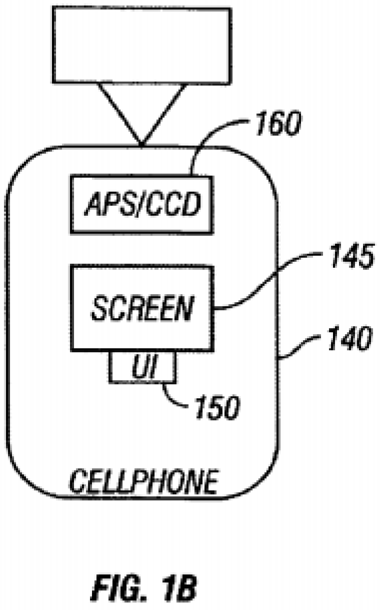 FIG. 1B Harris at 2: 3-19 • "The bar codes can be imaged/scanned in a number of different ways. One embodiment discloses using a camera to input and decode these bar codes. The embodiments are shown in FIGS. 1A and 1B. FIG. 1A describes using a personal digital assistant 100 as the input device. One preferred input device of this type is the Palm V(TM) type hand held computer. A bar code scanner can be used, such as the commercially available Symbol Technology SPT 1700. Alternatively, a camera add-on unit can be added to the Palm V and used as described herein. The PDA includes a screen 145 and user interface 146 all within the same housing 99. FIG. 1B shows the client being a cellular telephone which also may include a screen 145 and user interface 150. The cellular telephone can be

Element	'897 Claim Recitation	Exemplary Citations in Reference
		associated with either a dedicated bar code scanner or an image sensor 160 of the type used to obtain photographs for video telephony." Harris at FIG. 4 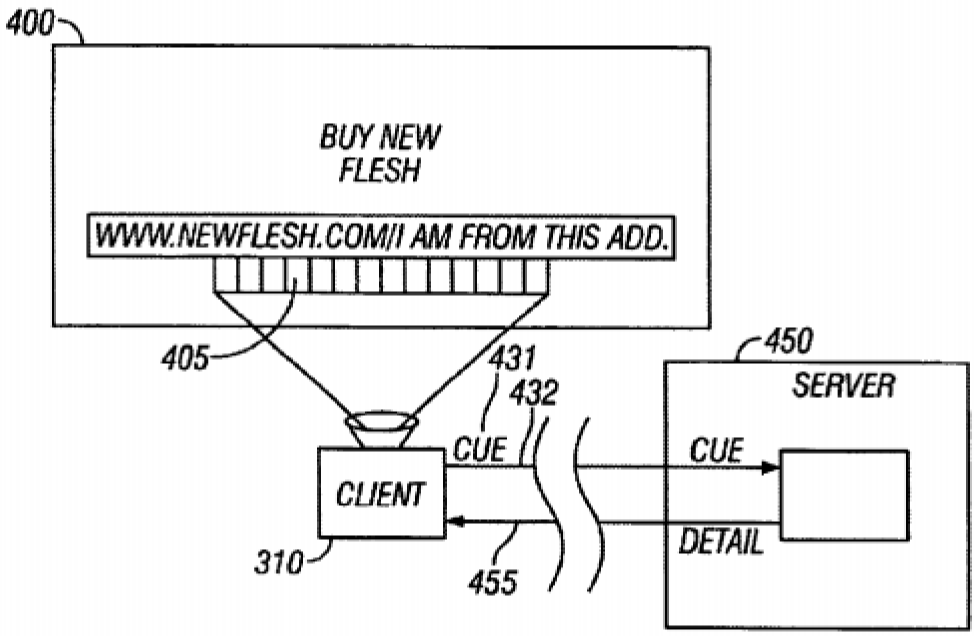 FIG. 4 Harris at 4:15-30 - "The advertisement 400 includes a bar code 405 therein. The bar code 405 is associated with the advertisement, and includes some additional information about the advertisement. For example, the bar code may include the web site address of the company preparing the advertisement, or appointment information about the advertisement, or a "vcf file" or the

Element	'897 Claim Recitation	Exemplary Citations in Reference
		like which is an importable file with address information about the company or author sponsoring the advertisement. In operation, the user brings one of the clients, either the cellular telephone 310 or PDA 315, into range of the bar code 405. The client reads the bar code and decodes it as noted above. The decoded information can represent ASCII information, a compressed file such as a zip file, G code information, or any other compressed or non-compressed information." To the extent it is found that Harris does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Harris does not anticipate this claim, Harris renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Harris with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[1D]	enabling, via the interactive system, identification of features from the displayed image, recognition of a target based on the features, association of the target with target information pertinent to the target, and performance of a transaction based on the target information.	Harris discloses enabling, via the interactive system, identification of features from the displayed image, recognition of a target based on the features, association of the target with target information pertinent to the target, and performance of a transaction based on the target information. Harris at FIG. 4

Element	'897 Claim Recitation	Exemplary Citations in Reference
		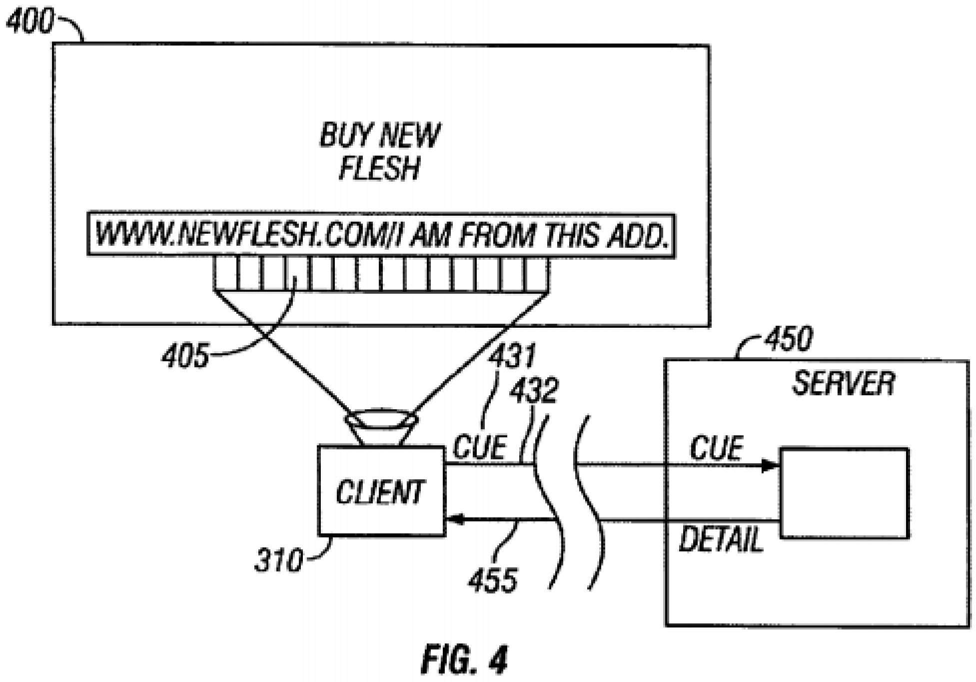 FIG. 4 Harris at 4:41-48 - "In this system, the information can represent a cue, e.g. a number, that is associated with a special function. The cue is used during a hot sync to access a publicly available information network, e.g. by direct connection or by the Internet. The cue 431 is sent over the Internet 432 to **the server 450**, and addresses more detailed information in a memory of the server. The server returns that information as 455, and the client receives the more detailed information." To the extent it is found that Harris does not expressly disclose this limitation,

Element	'897 Claim Recitation	Exemplary Citations in Reference
		the limitation is inherent. And to the extent it is found that Harris does not anticipate this claim, Harris renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Harris with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

2. Claim 30

Claim	'897 Claim Recitation	Exemplary Citations in Reference
30	The method of claim 25, wherein the step of displaying the displayed image includes displaying a time varying image.	To the extent it is found that Harris does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Harris does not anticipate this claim, Harris renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Harris with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

3. Claim 33

Claim	'897 Claim Recitation	Exemplary Citations in Reference
33	The method of claim 25, wherein the step of enabling performance of the transaction includes initiating a verification.	To the extent it is found that Harris does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Harris does not anticipate this claim, Harris renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Harris with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

4. Claim 34

Claim	'897 Claim Recitation	Exemplary Citations in Reference
34	The method of claim 25 wherein the step of enabling performance of the transaction includes initiating an authorization.	To the extent it is found that Harris does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Harris does not anticipate this claim, Harris renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Harris with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

5. Claim 38

Claim	'897 Claim Recitation	Exemplary Citations in Reference
38	The method of claim 25, wherein the step of enabling performance of the interaction includes initiating a transaction with an account.	Harris discloses the step of enabling performance of the interaction includes initiating a transaction with an account. Harris does not explicitly discloses a transaction with an account. Harris at 4:15-30 • "The advertisement 400 includes a bar code 405 therein. The bar code 405 is associated with the advertisement, and includes some additional information about the advertisement. For example, the bar code may include the web site address of the company preparing the advertisement, or appointment information about the advertisement, or a "vcf file" or the like which is an importable file with address information about the company or author sponsoring the advertisement. In operation, **the user** brings one of the clients, either the cellular telephone 310 or PDA 315, into range of the bar code 405. The client reads the bar code and decodes it as noted above. The decoded information can represent ASCII information, a compressed file such as a zip file, G code information, or any other compressed or non-compressed information." To the extent it is found that Harris does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Harris does not anticipate this claim, Harris renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Preliminary Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Harris with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

6. Claim 39

Claim	'897 Claim Recitation	Exemplary Citations in Reference
39	The method of claim 38, wherein the account includes at least one of the following: an on-line account, an account linked with the portable device, an account linked to a user of the portable device, a bank account, and a credit card account.	Harris discloses the account includes at least one of the following: an on-line account, an account linked with the portable device, an account linked to a user of the portable device, a bank account, and a credit card account. Harris does not explicitly disclose the step associated with an account. Harris at 4:15-30 • "The advertisement 400 includes a bar code 405 therein. The bar code 405 is associated with the advertisement, and includes some additional information about the advertisement. For example, the bar code may include the web site address of the company preparing the advertisement, or appointment information about the advertisement, or a "vcf file" or the like which is an importable file with address information about the company or author sponsoring the advertisement. In operation, **the user** brings one of the clients, either the cellular telephone 310 or PDA 315, into range of the bar code 405. The client reads the bar code and decodes it as noted above. The decoded information can represent ASCII information, a compressed file such as a zip file, G code information, or any other compressed or non-compressed information." To the extent it is found that Harris does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Harris does not anticipate this claim, Harris renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Final Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Harris with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these

		references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

Nantworks, LLC v. Bank of America Corporation
Central District of California, No. 2:20-cv-7872
Invalidity Chart for U.S. Patent 8,520,897 based on
U.S. Patent No. 6,097,834 to Krouse ("Krouse")

Exhibit H-18

Bank of America Corporation ("BofA" or "Defendant") provides this chart subject to all reservations, objections, statements, and disclaimers set forth herein and in BofA's Final Invalidity Contentions Cover Pleading, as well as any amendment, supplement, or modification thereof, which are incorporated herein by reference in their entirety.

On information and belief, the application that matured into Krouse was filed on June 13, 1997, claiming the priority date of June 13, 1997. Krouse published at least as of August 1, 2000, and is available as prior art at least under 35 U.S.C. § 102(a) and (e).

As shown in the chart below, Krouse anticipates asserted claims 25, 30, 31, 33, 34, 39 of the '897 Patent. To the extent Krouse is found not to expressly disclose certain limitations in the asserted claims, such limitations are inherent. To the extent Krouse is found not to anticipate any asserted claims or claim elements of the '897 Patent, the reference nevertheless renders those claims or claim elements obvious under 35 U.S.C. § 103, either alone or in combination with other art identified herein or in the cover pleading of BofA's Invalidity Contentions, for the reasons identified in the cover pleading. These Invalidity Contentions are not an admission by BofA that the accused products, including any current or past versions of the accused products, are covered by, or infringe the asserted claims, but are based instead on the recognition that if the claims are interpreted to be broad enough to encompass the accused products, the claims must also be construed to have that same scope when considering whether they are invalid.

The following chart is based on the Court's claim construction rulings (Dkt. Nos. 145, 153, 236), and to the extent not contradictory, Plaintiffs' apparent interpretations of the claims in Plaintiffs' "Preliminary" Final Infringement Contentions. BofA notes that in many instances, Plaintiffs' "Preliminary" Final Infringement Contentions ignore and/or contradict the Court's claim construction rulings. BofA does not accept the assumptions concerning the scope and meaning implicit in Plaintiffs' Preliminary Final Infringement Contentions, to the extent those assumptions are discernible, and reserves the right to challenge Plaintiffs' proposed (or implied) constructions. Furthermore, Plaintiffs' "Preliminary" Final Contentions fail to comply with the level of specificity required by SPR 2.1.2 and 2.1.3, and include improper reservation of rights to supplement with "additional ways in which BoA Accused Products infringe." Any alleged deficient discovery is of Plaintiffs own making. Most if not all of the deficiencies in Plaintiffs' Preliminary Final Contentions have been documented by Defendants through multiple letters since at least May 18, 2021 and up to Bank of America's most recent letter of August 19, 2023. Plaintiffs' shifting infringement contentions, deficiencies in its infringement contentions, and changing claim scope have prejudiced BofA from being able to reasonably prepare its defenses, including its invalidity contentions. BofA also reserves the right to revise and supplement these charts if and when Plaintiffs are permitted to provide revised or supplemental Infringement Contentions. Citations given in the chart below are merely representative of the respective elements to provide notice to Plaintiffs of BofA's invalidity theories and are not meant to be exhaustive.

1. Claim 25

Element	'897 Claim Recitation	Exemplary Citations in Reference
[25Pre]	A method of conducting a transaction with an interactive system, the method comprising	Krouse discloses a method of conducting a transaction with an interactive system. Krouse at Abstract • "In accordance with the present invention, financial transaction processing systems and methods are provided. One preferred embodiment of a method according to one aspect of the present invention includes generating an optically scanned image of at least a portion of document containing visual data, in a particular format, representing information related to the financial transaction. Recognition characteristics are generated from the scanned image and are compared to respective sets of reference recognition characteristics generated from respective other transaction documents having different respective formats to determine therefrom whether the particular format of the visual data matches one of the respective formats of the other documents. When such a match is found to exist, location is determined of a field in the scanned image to which optical character recognition may be applied to generate therefrom the information, based upon the respective format found to match the particular format of the visual data. Optical character recognition is then utilized to generate said visual data from said location."
[25A]	providing access to a device having a display;	Krouse discloses providing access to a device having a display. *Id.* at [25Pre]. Krouse at 6:26–34 • "Preferably, device **30** comprises a conventional alphanumeric keypad and/or keyboard for permitting input to the processor **40**, by the human

Element	'897 Claim Recitation	Exemplary Citations in Reference
		operator, of data and commands (which will be described more fully below) related to the transaction being processed, and an LCD/LED or computer monitor for displaying information transmitted thereto from the processor **40** and for echoing information data and commands input via the device **30**." Krouse at 16:66-17:19: • "The generator 224 then causes display 206 to display this information and to prompt via the display 206 for the user/customer at the display 206 to indicate via the user interface 206 whether the displayed information is correct. If the user/customer indicates that the displayed information is correct, the generator 224 then causes the display 206 to prompt the user/customer to provide means to pay the bill whose billing information is displayed on the display 206, otherwise the user/customer is prompted to input to the generator 224 the corrected billing information, which corrected information, once input, is displayed on the display with a request that the user verify correctness of Same. Once the user/customer indicates that the billing information displayed on the display 206 is correct, the generator 224 causes the display 206 to request that the user/customer indicate to the system 200 via the user interface 206 how the user/customer wishes to make payment based upon the displayed billing information. Interface 206 may comprise a conventional point of Sale debit card reading/transaction mechanism (not shown) to permit the user/customer to accomplish this. To the extent it is found that Krouse does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Krouse does not anticipate this claim, Krouse renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been

Element	'897 Claim Recitation	Exemplary Citations in Reference
		motivated to combine Krouse with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[25B]	displaying a displayed image on the display of the device;	Krouse discloses displaying a displayed image on the display of the device. *Id.* at [25Pre-A]. Krouse at 6:26–34 • "Preferably, device **30** comprises a conventional alphanumeric keypad and/or keyboard for permitting input to the processor **40**, by the human operator, of data and commands (which will be described more fully below) related to the transaction being processed, and an LCD/LED or computer monitor for displaying information transmitted thereto from the processor **40** and for echoing information data and commands input via the device **30**." Krouse at 12:63–13:4 • "When the document **400** is placed into the scanner **208**, the scanner **208** indicates same to the display **206**, whereupon the human operator (not shown) of the terminal **202** commands the scanner **208** to scan the entire side **402** of the document **400** containing the visual data **405** to generate a computer-readable scanned image of side **402** of the document 400 (i.e., computer-readable data from which an appropriately provisioned computer may generate a display image of the at least one scanned side **402** of the document **400**)." Krouse at 16:66-17:19:

Element	'897 Claim Recitation	Exemplary Citations in Reference
		• "The generator 224 then causes display 206 to display this information and to prompt via the display 206 for the user/customer at the display 206 to indicate via the user interface 206 whether the displayed information is correct. If the user/customer indicates that the displayed information is correct, the generator 224 then causes the display 206 to prompt the user/customer to provide means to pay the bill whose billing information is displayed on the display 206, otherwise the user/customer is prompted to input to the generator 224 the corrected billing information, which corrected information, once input, is displayed on the display with a request that the user verify correctness of Same. Once the user/customer indicates that the billing information displayed on the display 206 is correct, the generator 224 causes the display 206 to request that the user/customer indicate to the system 200 via the user interface 206 how the user/customer wishes to make payment based upon the displayed billing information. Interface 206 may comprise a conventional point of Sale debit card reading/transaction mechanism (not shown) to permit the user/customer to accomplish this. To the extent it is found that Krouse does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Krouse does not anticipate this claim, Krouse renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Krouse with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

Element	**'897 Claim Recitation**	**Exemplary Citations in Reference**
[25C]	presenting the displayed image proximate to an optical sensor; and	Krouse discloses presenting the displayed image proximate to an optical sensor. *Id.* at [25Pre-B]. Krouse at 12:63–13:4 - "When the document **400** is placed into the scanner **208**, the scanner **208** indicates same to the display **206**, whereupon the human operator (not shown) of the terminal **202** commands the scanner **208** to scan the entire side **402** of the document **400** containing the visual data **405** to generate a computer-readable scanned image of side **402** of the document 400 (i.e., computer-readable data from which an appropriately provisioned computer may generate a display image of the at least one scanned side **402** of the document **400**)." Krouse at 16:66-17:19: - "The generator 224 then causes display 206 to display this information and to prompt via the display 206 for the user/customer at the display 206 to indicate via the user interface 206 whether the displayed information is correct. If the user/customer indicates that the displayed information is correct, the generator 224 then causes the display 206 to prompt the user/customer to provide means to pay the bill whose billing information is displayed on the display 206, otherwise the user/customer is prompted to input to the generator 224 the corrected billing information, which corrected information, once input, is displayed on the display with a request that the user verify correctness of Same. Once the user/customer indicates that the billing information displayed on the display 206 is correct, the generator 224 causes the display 206 to request that the user/customer indicate to the system 200 via the user interface 206 how the user/customer wishes to make payment based upon the displayed

Element	**'897 Claim Recitation**	**Exemplary Citations in Reference**
		billing information. Interface 206 may comprise a conventional point of Sale debit card reading/transaction mechanism (not shown) to permit the user/customer to accomplish this. To the extent it is found that Krouse does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Krouse does not anticipate this claim, Krouse renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Krouse with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[25D]	enabling, via the interactive system, identification of features from the displayed image, recognition of a target based on the features, association of the target with target information pertinent to the target, and performance of a transaction based on the target information.	Krouse discloses enabling, via the interactive system, identification of features from the displayed image, recognition of a target based on the features, association of the target with target information pertinent to the target, and performance of a transaction based on the target information. *Id.* at [25Pre-C]. Krouse at Abstract • "In accordance with the present invention, financial transaction processing systems and methods are provided. One preferred embodiment of a method according to one aspect of the present invention includes generating an optically scanned image of at least a portion of document containing visual data, in a particular format,

Element	'897 Claim Recitation	Exemplary Citations in Reference
		representing information related to the financial transaction. Recognition characteristics are generated from the scanned image and are compared to respective sets of reference recognition characteristics generated from respective other transaction documents having different respective formats to determine therefrom whether the particular format of the visual data matches one of the respective formats of the other documents. When such a match is found to exist, location is determined of a field in the scanned image to which optical character recognition may be applied to generate therefrom the information, based upon the respective format found to match the particular format of the visual data. Optical character recognition is then utilized to generate said visual data from said location." Krouse at 3:43–4:12 • "One preferred embodiment of the system of this aspect of the present invention includes an optical scanner for generating a computer-readable scanned image of the document, and an optical character recognizer for generating, from the scanned image, transaction data indicative of at least a portion of the visual information. This embodiment of the system of this aspect of the present invention also includes an authorization record generator for generating, based upon the transaction data, a record of the transaction for being assented to by the party tendering the document prior to completion of the transaction, and for generating from the record once it is assented to by the party, a computer-readable record of assent of the party to the transaction. A collection system is also included in this embodiment of the system of this aspect of the present invention for generating in response to the transaction data and the computer-readable record, an electronic financial transaction request for being used by an electronic transaction system to execute the electronic financial transaction.

Element	'897 Claim Recitation	Exemplary Citations in Reference
		One embodiment of the method of the first aspect of the present invention includes generating a computer-readable scanned image of the document, and generating, from the scanned image, transaction data indicative of at least a portion of the visual information contained on the document. A record is generated, from the transaction data, of the electronic financial transaction being processed, which record is for being assented to by the party tendering the document, prior to completion of said transaction. Once the record has been assented to by the party tendering the document, a computer-readable record of the party's assent to the transaction is generated from the assented-to record. An electronic financial transaction request is then generated in response to the transaction data and the computer-readable assented-to record, which request is for being used by an electronic financial transaction system to execute the electronic financial transaction." Krouse at 7:38–67 • "One such standardized portion in check **110** is Magnetic Ink Character Recognizable (MICR) line **94**, which is positioned at the bottom of the check **110** and in check **110**, includes three fields **114, 116, 118** of printed digits separated by special control characters. Fields **114, 116** contain the financial institution transaction routing number and account number on which the check **110** is drawn, and field **118** contains the check number of the check **110**. It should be noted, however, that other than the left most nine digits (and associated separator characters) of the MICR line **94**, the contents of said line may vary, depending upon the financial institution upon which the negotiable instrument is drawn. System **10** may be adapted to recognize and parse information contained in different types of negotiable instrument MICR lines (not shown). Based upon the standardized location, size, style, etc. of the MICR line **94** in document **110** and the fields **114, 116, 118** contained therein,

Element	'897 Claim Recitation	Exemplary Citations in Reference
		recognizer **32** utilizes optical character recognition and image processing techniques to extract from the scanned image data the data generated by the scanner **34** from scanning the MICR line portion **94** of the document **110**, and once thus extracted, to generate numeric transaction data indicative of the visual information (i.e., the numbers printed in fields **114, 116, 118**) in MICR line **94**. This numeric data is then transmitted to the processor **40**, which transmits a copy of the numeric data to the storage system **42** where it is stored in association with the scanned image data and the payee and payment amount data previously stored therein." Krouse at 12:63-13:11: • "When the document 400 is placed into the scanner 208, the scanner 208 indicates same to the display 206, where upon the human operator (not shown) of the terminal 202 commands the scanner 208 to Scan the entire side 402 of the document 400 containing the visual data 405 to generate a computer-readable scanned image of side 402 of the document 400 (i.e., computer-readable data from which an appropriately provisioned computer may generate a display image of the at least one scanned side 402 of the document 400). This scanned image of side 402 is then transmitted to the image characterization generator 204 in order to generate therefrom recognition characteristics for permitting the System 200 to determine whether document 400 has a format that matches that of another billing transaction document and to execute further processing based upon this determination, as will be described in greater detail below." To the extent it is found that Krouse does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Krouse does not anticipate this claim, Krouse renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in

Element	'897 Claim Recitation	Exemplary Citations in Reference
		combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Krouse with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

2. Claim 30

Claim	'897 Claim Recitation	Exemplary Citations in Reference
30	The method of claim 25, wherein the step of displaying the displayed image includes displaying a time varying image.	Krouse discloses the step of displaying the displayed image includes displaying a time varying image. *Id.* at [25]. Krouse at 12:63-13:11: • "When the document 400 is placed into the scanner 208, the scanner 208 indicates same to the display 206, where upon the human operator (not shown) of the terminal 202 commands the scanner 208 to Scan the entire side 402 of the document 400 containing the visual data 405 to generate a computer-readable scanned image of side 402 of the document 400 (i.e., computer-readable data from which an appropriately provisioned computer may generate a display image of the at least one scanned side 402 of the document 400). This scanned image of side 402 is then transmitted to the image characterization generator 204 in order to generate therefrom recognition characteristics for permitting the System 200 to determine whether document 400 has a format that matches that of another billing transaction document and to execute further processing based upon this determination, as will be described in greater detail below."

		To the extent it is found that Krouse does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Krouse does not anticipate this claim, Krouse renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Krouse with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

3. Claim 33

Claim	'897 Claim Recitation	Exemplary Citations in Reference
33	The method of claim 25, wherein the step of enabling performance of the transaction includes initiating a verification.	Krouse discloses the step of enabling performance of the transaction includes initiating a verification. Krouse at Abstract • "In accordance with the present invention, financial transaction processing systems and methods are provided. One preferred embodiment of a method according to one aspect of the present invention includes generating an optically scanned image of at least a portion of document containing visual data, in a particular format, representing information related to the financial transaction. Recognition characteristics are generated from the scanned image and are compared to respective sets of reference recognition characteristics generated from respective other transaction documents having different respective formats to determine therefrom whether the particular format of the

		visual data matches one of the respective formats of the other documents. When such a match is found to exist, location is determined of a field in the scanned image to which optical character recognition may be applied to generate therefrom the information, based upon the respective format found to match the particular format of the visual data. Optical character recognition is then utilized to generate said visual data from said location.” Krouse at 3:49–57 • “This embodiment of the system of this aspect of the present invention also includes an authorization record generator for generating, based upon the transaction data, a record of the transaction for being assented to by the party tendering the document prior to completion of the transaction, and for generating from the record once it is assented to by the party, a computer-readable record of assent of the party to the transaction. A collection system is also included in this embodiment of the system of this aspect of the present invention for generating in response to the transaction data and the computer-readable record, an electronic financial transaction request for being used by an electronic transaction system to execute the electronic financial transaction.” Krouse at 8:1–12 • “A second copy of the numeric data is also transmitted from the processor **40** to the authorization record generator **38**, which also retrieves from the storage system **42** the payee and payment amount data associated with this transaction and transmits same to the generator **38**. Processor **40** also supplies to the generator 38 from the storage system **42** data indicative of the initiation of the current transaction. Preferably, each of the terminals **18, 20, 22** of the system **10** is associated with a unique station/terminal identification code or number which is stored in

		the terminal's respective storage system **42**. This terminal identification number is also supplied to the generator **38** from the system **42**." Krouse at 17:3-16: • "If the user/customer indicates that the displayed information is correct, the generator 224 then causes the display 206 to prompt the user/customer to provide means to pay the bill whose billing information is displayed on the display 206, otherwise the user/customer is prompted to input to the generator 224 the corrected billing information, which corrected information, once input, is displayed on the display with a request that the user verify correctness of Same. Once the user/customer indicates that the billing information displayed on the display 206 is correct, the generator 224 causes the display 206 to request that the user/customer indicate to the system 200 via the user interface 206 how the user/customer wishes to make payment based upon the displayed billing information." To the extent it is found that Krouse does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Krouse does not anticipate this claim, Krouse renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Krouse with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

4. Claim 34

Claim	’897 Claim Recitation	Exemplary Citations in Reference
34	The method of claim 25 wherein the step of enabling performance of the transaction includes initiating an authorization.	Krouse discloses the step of enabling performance of the transaction includes initiating an authorization. *Id.* at [25]. Krouse at Abstract • "In accordance with the present invention, financial transaction processing systems and methods are provided. One preferred embodiment of a method according to one aspect of the present invention includes generating an optically scanned image of at least a portion of document containing visual data, in a particular format, representing information related to the financial transaction. Recognition characteristics are generated from the scanned image and are compared to respective sets of reference recognition characteristics generated from respective other transaction documents having different respective formats to determine therefrom whether the particular format of the visual data matches one of the respective formats of the other documents. When such a match is found to exist, location is determined of a field in the scanned image to which optical character recognition may be applied to generate therefrom the information, based upon the respective format found to match the particular format of the visual data. Optical character recognition is then utilized to generate said visual data from said location." Krouse at 3:49–57 • "This embodiment of the system of this aspect of the present invention also includes an authorization record generator for generating, based upon the transaction data, a record of the transaction for being assented to by the party tendering the document prior to completion of the transaction, and for generating from the record once it is assented to by

		the party, a computer-readable record of assent of the party to the transaction. A collection system is also included in this embodiment of the system of this aspect of the present invention for generating in response to the transaction data and the computer-readable record, an electronic financial transaction request for being used by an electronic transaction system to execute the electronic financial transaction." Krouse at 8:1–12 • "A second copy of the numeric data is also transmitted from the processor **40** to the authorization record generator **38**, which also retrieves from the storage system **42** the payee and payment amount data associated with this transaction and transmits same to the generator **38**. Processor **40** also supplies to the generator 38 from the storage system **42** data indicative of the initiation of the current transaction. Preferably, each of the terminals **18, 20, 22** of the system **10** is associated with a unique station/terminal identification code or number which is stored in the terminal's respective storage system **42**. This terminal identification number is also supplied to the generator **38** from the system **42**." To the extent it is found that Krouse does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Krouse does not anticipate this claim, Krouse renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Krouse with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

5. Claim 38 (Non-asserted)

Claim	'897 Claim Recitation	Exemplary Citations in Reference
38	The method of claim 25, wherein the step of enabling performance of the interaction includes initiating a transaction with an account.	Krouse discloses the step of enabling performance of the interaction includes initiating a transaction with an account. Krouse at 12:38–45 • "Processing of a financial transaction using system **200** begins by placing a bill document (such as bill document **400** shown in FIG. 8A) to be processed into scanner **208** so as to permit at least one side **402** of the document **400** containing an optical character recognition line (OCR) **404** which contains numerical data **405** visually displaying information related to the transaction to be processed (e.g., amount due **406** and customer account number **408**)." Krouse at 17:18–45 • "The information obtained from reading the customer's debit card and provision of the customer's PIN via the interface **206** may then be utilized by the generator **224**, together with the aforesaid billing information, to generate an ACH EFT request to transfer funds from the customer's bank account (indicated via the debit card information) to the bill payee's bank account, to pay the customer's bill as indicated on the billing document **400**. This request is transmitted from the generator **224** to the ACH/EFT system via the interface **226**, and a copy of said request is stored in the archive **228** in association with the other billing information and scanned image obtained from the document **400** by the system **200**. Once the ACH system notifies the generator **224** that the request has been executed, the generator **224** provides to the payee system **232** via the interface **230** a copy of the notice, the submitted ACH/EFT request, and a statement of the customer's billing account

		number in order to permit the payee **232** to update its records to reflect the customer's payment to reduce the outstanding balance of the customer's billing account. Alternatively, the ACHIEFT request may be forwarded by the generator **224** to the payee **232** via the interface **230** for submission by the payee **232** directly to the ACH/EFT system. The archive **228** may also be adapted to permit the payee **232** to access the other information associated with the payee (e.g., the particular EFT requests generated for effectuating payment of bills generated by that payee and scanned images of such bills) stored in the archive **228**." To the extent it is found that Krouse does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Krouse does not anticipate this claim, Krouse renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Preliminary Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Krouse with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

6. Claim 39

Claim	'897 Claim Recitation	Exemplary Citations in Reference
39	The method of claim 38, wherein the account includes at least one of the following: an on-line account, an account linked with the portable device, an account linked to a user	Krouse discloses the account includes at least one of the following: an on-line account, an account linked with the portable device, an account linked to a user of the portable device, a bank account, and a credit card account. *Id.* at [38].

	of the portable device, a bank account, and a credit card account.	Krouse at Fig. 6 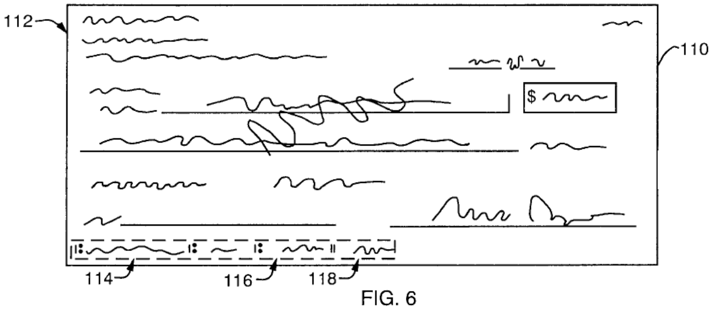 FIG. 6 Krouse at 7:43–53 - "Fields **114, 116** contain the financial institution transaction routing number and account number on which the check **110** is drawn, and field **118** contains the check number of the check **110**. It should be noted, however, that other than the left most nine digits (and associated separator characters) of the MICR line **94**, the contents of said line may vary, depending upon the financial institution upon which the negotiable instrument is drawn. System **10** may be adapted to recognize and parse information contained in different types of negotiable instrument MICR lines (not shown)." Krouse at 12:37–45 - "Processing of a financial transaction using system **200** begins by placing a bill document (such as bill document **400** shown in FIG. 8A) to be processed into scanner **208** so as to permit at least one side **402** of the document **400** containing an optical character recognition line (OCR) **404** which contains numerical data **405** visually displaying information related to the transaction to be processed (e.g., amount due **406** and customer account number **408**)."

		To the extent it is found that Krouse does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that Krouse does not anticipate this claim, Krouse renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine Krouse with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

Nantworks, LLC v. Bank of America Corporation
Central District of California, No. 2:20-cv-7872
Invalidity Chart for U.S. Patent 8,520,897 based on the Query by Image and Video Content ("QBIC") System

Exhibit H-23

Bank of America Corporation ("BofA" or "Defendant") provides this second supplemental chart subject to all reservations, objections, statements, and disclaimers set forth herein and in BofA's Final Invalidity Contentions Cover Pleading, as well as any amendment, supplement, or modification thereof, which are incorporated herein by reference in their entirety. A statement of reasons and good cause for each supplement in this chart is identified in the cover pleading.

QBIC was known or used by others in the United States and publicly used or on sale in the United States by November 5, 2000. *See generally* W. Niblack et al., The QBIC Project: Querying Images By Content Using Color, Texture, and Shape, SPIE Vol. 1908 (1993), pp. 173–187 (IBM0002390–404); M. Flickner et al., Query by Image and Video Content: The QBIC System ("QBIC"), IEEE Sept. 1995, pp. 23–32 (IBM0000893–902); W. Niblack et al, Updates to the QBIC System, SPIE Vol. 3312 (1997), pp. 150–161 (IBM 0002419–430); and Query By Image Content QBIC Demonstration Program, IBM Research, Sept. 1998 (IBM000747). QBIC is available as prior art at least under 35 U.S.C. § 102(a), (b), and (g), having been known or used by others in the United States and publicly used or on sale in the United States by at least 1995 and certainly by November 5, 2000, including, for example, at http://libra.uc.davis.edu through the Art History Department at U.C. Davis; at www..thinker.org/imagebase/index-2 .htm3 through the Fine Arts Museum of San Francisco; through demos at http://wwwqbic.almaden.ibm.com; and through IBM's commercial products, including in IBM's Ultimedia Manager, the DB2 series of products such as the Image Extender in DB2 Universal Database, and the Enterprise Information Portal. *See generally* Deposition of IBM Corporate Representative Myron Flickner, *Nantworks, LLC and Nant Holdings IP, LLC v. Bank of America Corporation and Bank of America, N.A.*, 2:20-cv-07872-GW-PVC (C.D.C.A. 2020) ("Flickner 8/29/23 Depo") at e.g., 32–33, 39–40, 44, 62–63, 66–67,72–74, 103–107, 110–112, 127, 163–174, 183–189, 191–195, 215, 219–220, 241, 307–309, 311–317, Exs. 3–13; *see also* Deposition of IBM Corporate Representative Myron Flickner, *Network-1 Technologies, Inc., v. Google LLC and Youtube LLC*, 1:14-cv-09558-PGG (S.D.N.Y. 2019) ("Flickner Depo"), including, e.g., at 13, 14, 15–18, 29–30, 34, 36, 37, 43–46, 48, 58, 60, 61–62, 67–69, 72, 74–76, 84–86, 88–89, 98–100, 102, 106, 113, 125, 135, 152–154, 160–161, 170–172, 180–181, 194, 198–201, 203–205, 215–216, 220–221, and Exs. 1–16. As shown in the chart below, QBIC anticipates asserted claims 25, 30, 33, 34, and 39 of the '897 Patent. The disclosures/exemplary citations for each element incorporate the disclosures/exemplary citations of the proceeding limitations. To the extent QBIC is found not to expressly disclose certain limitations in the asserted claims, such limitations are inherent. To the extent QBIC is found not to anticipate any asserted claims or claim elements of the '897 Patent, the reference nevertheless renders those claims or claim elements obvious under 35 U.S.C. § 103, either alone or in combination with other art identified herein or in the cover pleading of BofA's Final Invalidity Contentions, for the reasons identified in the cited references and/or in the cover pleading. These Invalidity Contentions are not an admission by BofA that the accused products, including any current or past versions of the accused products, are covered by, or infringe the asserted claims, but are based instead on the recognition that if the claims are interpreted to be broad enough to encompass the accused products, the claims must also be construed to have that same scope when considering whether they are invalid. The following chart incorporates by reference the analysis from Exhibit H-21 for each limitation.

The following chart is based on the Court's claim construction rulings (Dkt. Nos. 145, 153, 236), and to the extent not contradictory, Plaintiffs' apparent interpretations of the claims in Plaintiffs' "Preliminary" Final Infringement Contentions. BofA notes that in many instances, Plaintiffs' "Preliminary" Final Infringement Contentions ignore and/or contradict the Court's claim construction rulings. BofA does not accept the assumptions concerning the scope and meaning implicit in Plaintiffs' Preliminary Final Infringement Contentions, to the extent those assumptions are discernible, and reserves the right to challenge Plaintiffs' proposed (or implied) constructions. Furthermore, Plaintiffs' "Preliminary" Final Contentions fail to comply with the level of specificity required by SPR 2.1.2 and 2.1.3, and include improper reservation of rights to supplement with "additional ways in which BoA Accused Products infringe." Any alleged deficient discovery is of Plaintiffs own making. Most if not all of the deficiencies in Plaintiffs' Preliminary Final Contentions have been documented by Defendants through multiple letters since at least May 18, 2021 and up to Bank of America's most recent letter of August 19, 2023. Plaintiffs' shifting infringement contentions, deficiencies in its infringement contentions, and changing claim scope have prejudiced BofA from being able to reasonably prepare its defenses, including its invalidity contentions. BofA also reserves the right to revise and supplement these charts if and when Plaintiffs are permitted to provide revised or supplemental Infringement Contentions. Citations given in the chart below are merely representative of the respective elements to provide notice to Plaintiffs of BofA's invalidity theories and are not meant to be exhaustive.

1. Claim 25

Element	'897 Claim Recitation	Exemplary Citations in Reference
[25Pre]	A method of conducting a transaction with an interactive system, the method comprising	QBIC discloses a method of conducting a transaction with an interactive system. The QBIC Project: Querying Images by Content Using Color, Texture, and Shape (IBM 0002390–404) at 2390 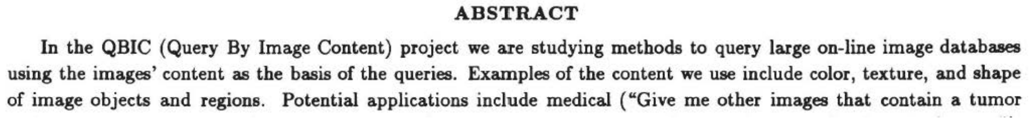 ABSTRACT In the QBIC (Query By Image Content) project we are studying methods to query large on-line image databases using the images' content as the basis of the queries. Examples of the content we use include color, texture, and shape of image objects and regions. Potential applications include medical ("Give me other images that contain a tumor QBIC Demo (IBM000747)  QBIC$^{(TM)}$ -- IBM's Query By Image Content QBIC allows you to make queries of image databases large and small based on visual characteristics -- properties such as color percentages, color layout and textures occurring in the images. The queries match image colors, textures and their positions without having to describe them with in words. QBIC's content-based queries can also be combined with keywords to provide a powerful multimedia search and retrieval tool. As image collections grow larger and become more prolific, QBIC is the best tool to more efficiently search and navigate through them. In this demo, we use QBIC to index about 470 color images based on their color and texture information. You can evaluate the QBIC Search capability by following the link located at left called IBM Image Search Demo. For license and business related questins, please contact Ted Loewenberg by phone 408-927-1202, or by email Flickner Depo at, e.g., 24, 41-42, 62-63 (user could take a picture or video and submit it to the QBIC system).

Element	'897 Claim Recitation	Exemplary Citations in Reference
		Q. And you mentioned that you would extract features from the images? A. We did. Q. How would you do that? A. We ran algorithms on the pixels of the image and from that we created a feature. Q. What sort of features would you extract from the images? A. A QBIC system did color, shape and texture. Q. Okay. So let's take color. How would you, say, extract a color feature from an image? A. We typically would compute -- compute a color histogram. Q. And what's a color histogram? A. It's a measure of -- it's an estimate of the probability density of a color being in an image. Q. So you would take an image and extract that color histogram from it? A. Yup.

		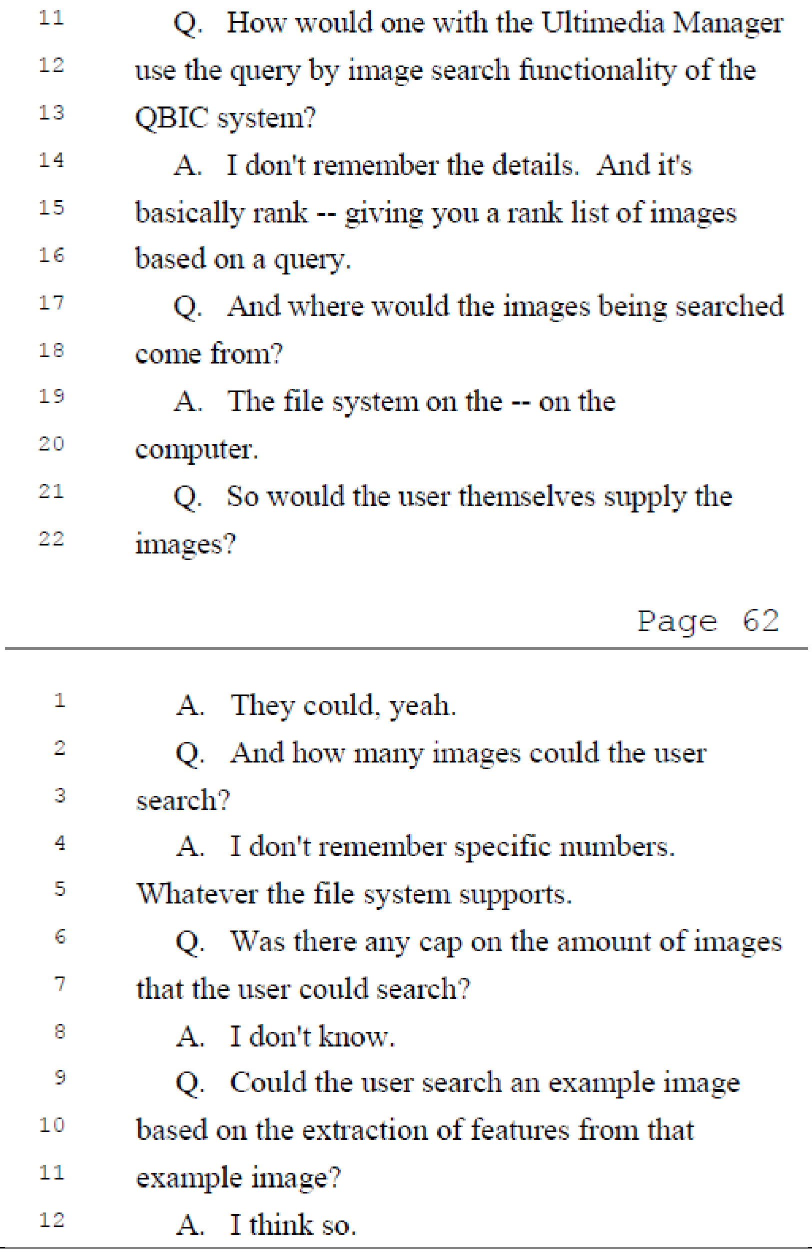 Q. How would one with the Ultimedia Manager use the query by image search functionality of the QBIC system? A. I don't remember the details. And it's basically rank -- giving you a rank list of images based on a query. Q. And where would the images being searched come from? A. The file system on the -- on the computer. Q. So would the user themselves supply the images? Page 62 A. They could, yeah. Q. And how many images could the user search? A. I don't remember specific numbers. Whatever the file system supports. Q. Was there any cap on the amount of images that the user could search? A. I don't know. Q. Could the user search an example image based on the extraction of features from that example image? A. I think so.

Element	'897 Claim Recitation	Exemplary Citations in Reference
		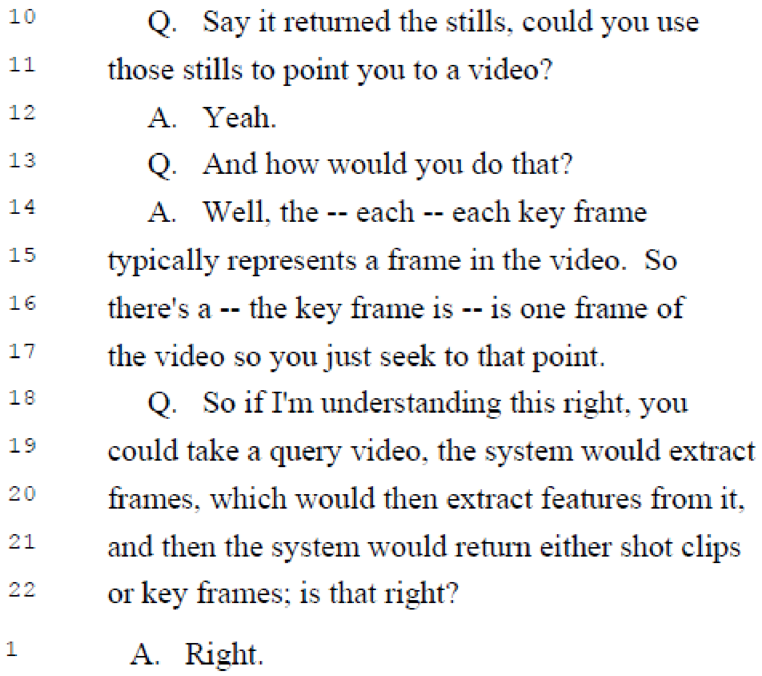 Q. Say it returned the stills, could you use those stills to point you to a video? A. Yeah. Q. And how would you do that? A. Well, the -- each -- each key frame typically represents a frame in the video. So there's a -- the key frame is -- is one frame of the video so you just seek to that point. Q. So if I'm understanding this right, you could take a query video, the system would extract frames, which would then extract features from it, and then the system would return either shot clips or key frames; is that right? A. Right. Flickner 3/29/23 Depo at 42

Element	'897 Claim Recitation	Exemplary Citations in Reference
		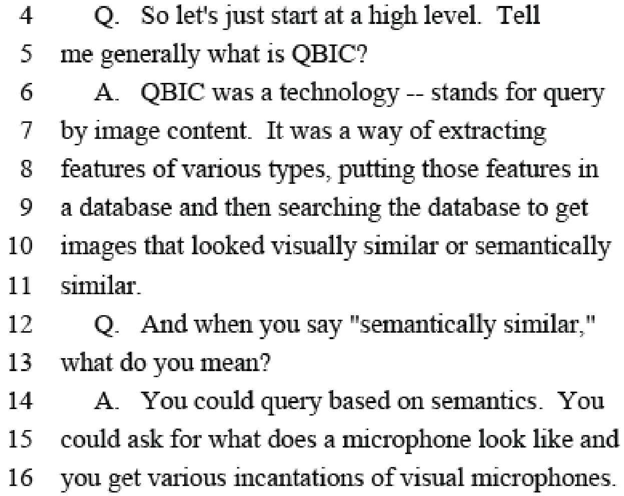 Q. So let's just start at a high level. Tell me generally what is QBIC? A. QBIC was a technology -- stands for query by image content. It was a way of extracting features of various types, putting those features in a database and then searching the database to get images that looked visually similar or semantically similar. Q. And when you say "semantically similar," what do you mean? A. You could query based on semantics. You could ask for what does a microphone look like and you get various incantations of visual microphones. *Id.* at 205–206

Element	'897 Claim Recitation	Exemplary Citations in Reference
		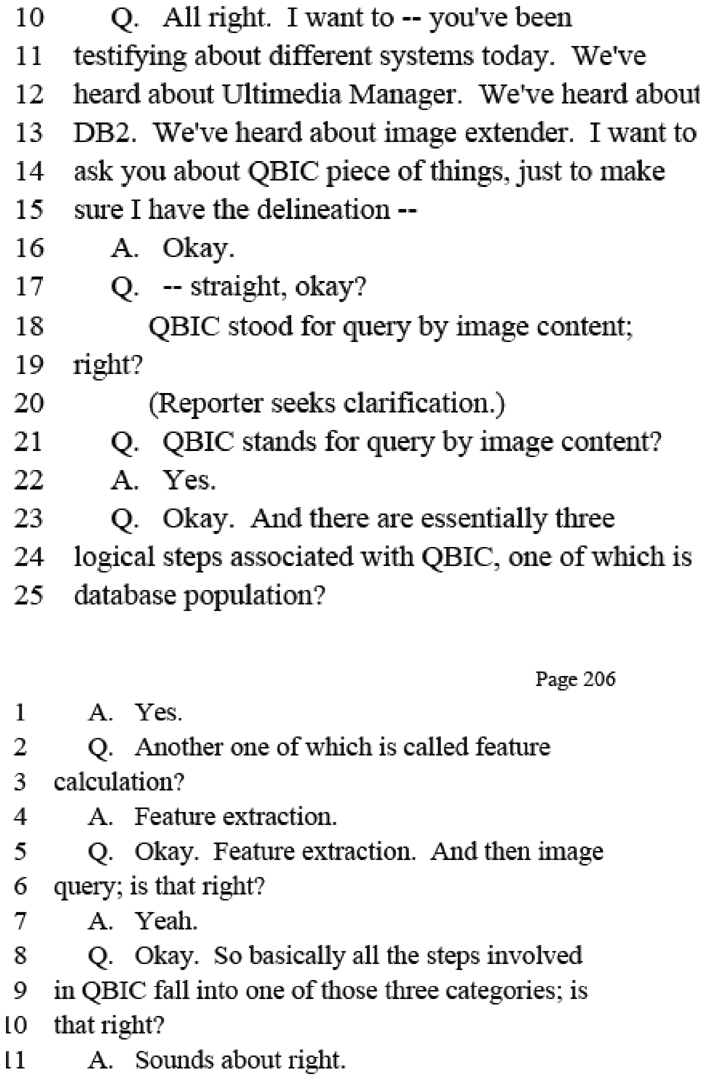 Q. All right. I want to -- you've been testifying about different systems today. We've heard about Ultimedia Manager. We've heard about DB2. We've heard about image extender. I want to ask you about QBIC piece of things, just to make sure I have the delineation -- A. Okay. Q. -- straight, okay? QBIC stood for query by image content; right? (Reporter seeks clarification.) Q. QBIC stands for query by image content? A. Yes. Q. Okay. And there are essentially three logical steps associated with QBIC, one of which is database population? Page 206 A. Yes. Q. Another one of which is called feature calculation? A. Feature extraction. Q. Okay. Feature extraction. And then image query; is that right? A. Yeah. Q. Okay. So basically all the steps involved in QBIC fall into one of those three categories; is that right? A. Sounds about right. *Id.* at 197–198

Element	'897 Claim Recitation	Exemplary Citations in Reference
		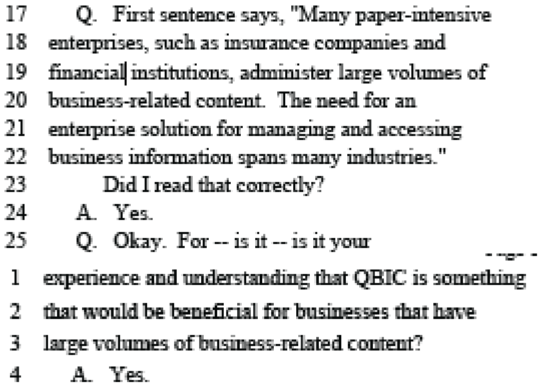 Q. First sentence says, "Many paper-intensive enterprises, such as insurance companies and financial institutions, administer large volumes of business-related content. The need for an enterprise solution for managing and accessing business information spans many industries." Did I read that correctly? A. Yes. Q. Okay. For -- is it -- is it your experience and understanding that QBIC is something that would be beneficial for businesses that have large volumes of business-related content? A. Yes. *Id.* at 307 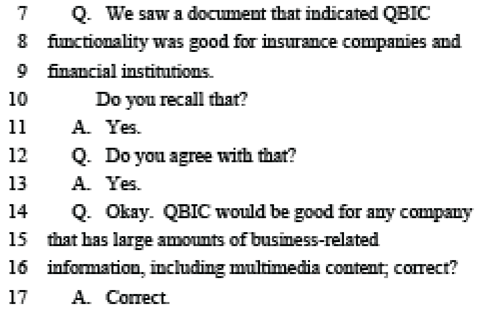 Q. We saw a document that indicated QBIC functionality was good for insurance companies and financial institutions. Do you recall that? A. Yes. Q. Do you agree with that? A. Yes. Q. Okay. QBIC would be good for any company that has large amounts of business-related information, including multimedia content; correct? A. Correct. A person of ordinary skill in the art would have been motivated to improve, modify, or combine QBIC with prior art reference teachings to arrive at the limitations in this claim at least because QBIC was designed for applications that included large amounts of business-related information, including multi-media content, including for financial institutions. *See, e.g.*, IBM 000789; Flickner Depo Rough Tr. at 191:16-192:8, 293:14-294:8. The QBIC system was implemented in IBM commercial products, including in IBM's Ultimedia Manager, the DB2 series of products such as the Image Extender in DB2 Universal Database, and the Enterprise Information

Element	'897 Claim Recitation	Exemplary Citations in Reference
		Portal. All of the features referenced in this chart were incorporated into these commercial implementations before November 6, 2000, as reflected in a combination of scientific papers, IBM's internal documentation, and the QBIC demo software applications. *See, e.g.*, Flickner 8/29/23 Depo Tr. at 33:9-24, 71:21-72:10, 164:21-166:2, 166:11-175:24, 178:5-186:17, 186:21-192:4, 192:14-194:1, 195:16-196:4, 194:1-308:1; IBM 000001-617; IBM 000747; IBM 000777-880; IBM 000893-902; IBM 0002315-320; IBM 0002390-404; IBM 0002418-430.
[25A]	providing access to a device having a display;	Flickner Depo at, e.g., 24, 41-42, 62-63 (user could take a picture or video and submit it to the QBIC system). Flickner 8/29/23 Depo at 75–76 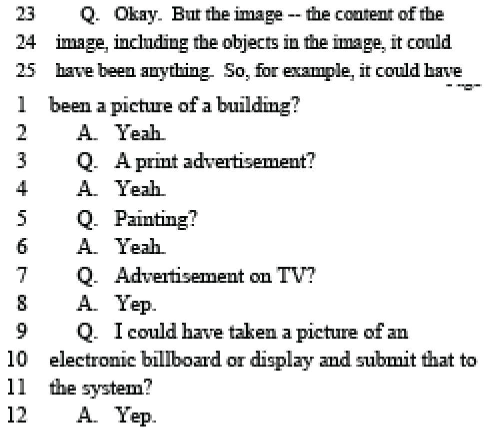 Q. Okay. But the image -- the content of the image, including the objects in the image, it could have been anything. So, for example, it could have been a picture of a building? A. Yeah. Q. A print advertisement? A. Yeah. Q. Painting? A. Yeah. Q. Advertisement on TV? A. Yep. Q. I could have taken a picture of an electronic billboard or display and submit that to the system? A. Yep. To the extent it is found that QBIC does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that QBIC does not anticipate this claim, QBIC renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine QBIC with any of the other prior art references identified in

Element	'897 Claim Recitation	Exemplary Citations in Reference
		the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[25B]	displaying a displayed image on the display of the device;	Flickner Depo at, e.g., 24, 41-42, 62-63 (user could take a picture or video and submit it to the QBIC system). Flickner 8/29/23 Depo at 75–76 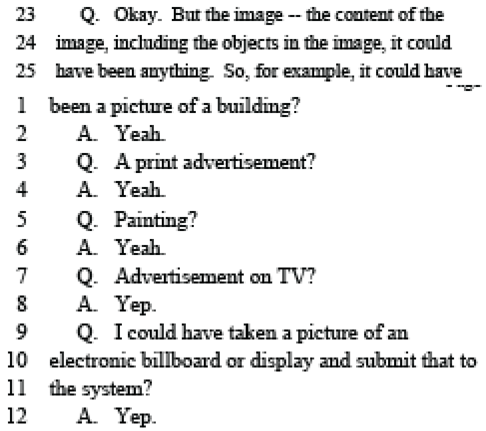 Q. Okay. But the image -- the content of the image, including the objects in the image, it could have been anything. So, for example, it could have been a picture of a building? A. Yeah. Q. A print advertisement? A. Yeah. Q. Painting? A. Yeah. Q. Advertisement on TV? A. Yep. Q. I could have taken a picture of an electronic billboard or display and submit that to the system? A. Yep. To the extent it is found that QBIC does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that QBIC does not anticipate this claim, QBIC renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine QBIC with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the

Element	'897 Claim Recitation	Exemplary Citations in Reference
		reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[25C]	presenting the displayed image proximate to an optical sensor; and	Flickner Depo at, e.g., 24, 41-42, 62-63 (user could take a picture or video and submit it to the QBIC system). Flickner 8/29/23 Depo at 75–76 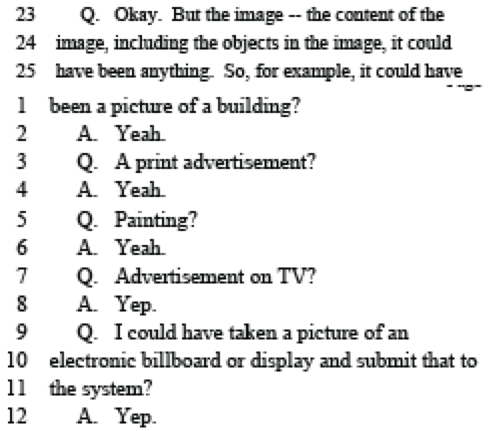 Q. Okay. But the image -- the content of the image, including the objects in the image, it could have been anything. So, for example, it could have been a picture of a building? A. Yeah. Q. A print advertisement? A. Yeah. Q. Painting? A. Yeah. Q. Advertisement on TV? A. Yep. Q. I could have taken a picture of an electronic billboard or display and submit that to the system? A. Yep. *Id.* at 53

Element	'897 Claim Recitation	Exemplary Citations in Reference
		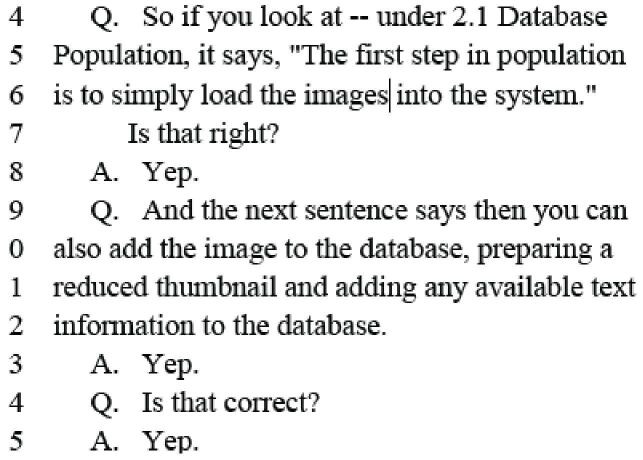 Q. So if you look at -- under 2.1 Database Population, it says, "The first step in population is to simply load the images into the system." Is that right? A. Yep. Q. And the next sentence says then you can also add the image to the database, preparing a reduced thumbnail and adding any available text information to the database. A. Yep. Q. Is that correct? A. Yep. *Id.* at 60–61

Element	'897 Claim Recitation	Exemplary Citations in Reference
		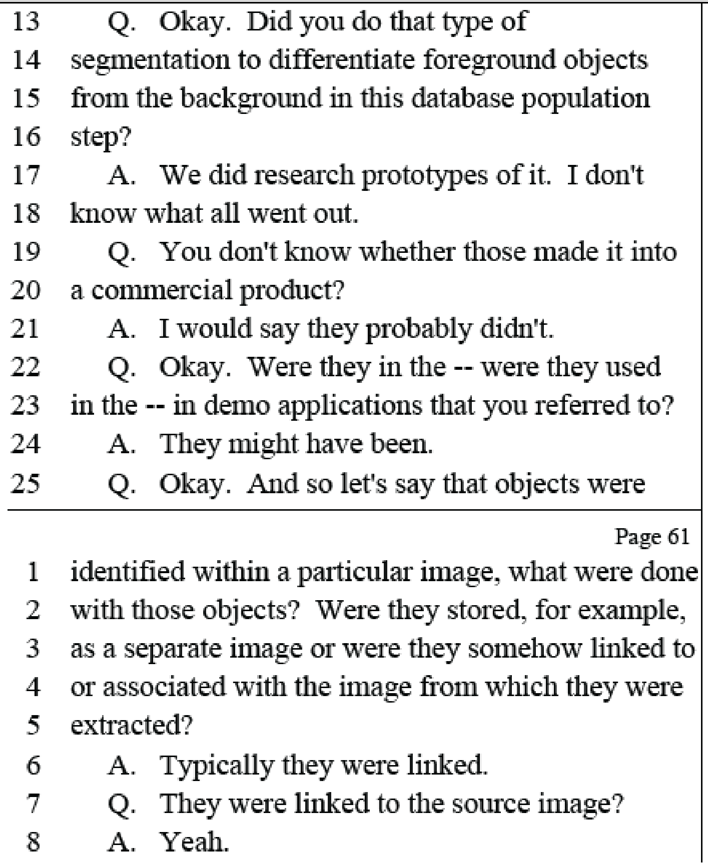 Q. Okay. Did you do that type of segmentation to differentiate foreground objects from the background in this database population step? A. We did research prototypes of it. I don't know what all went out. Q. You don't know whether those made it into a commercial product? A. I would say they probably didn't. Q. Okay. Were they in the -- were they used in the -- in demo applications that you referred to? A. They might have been. Q. Okay. And so let's say that objects were Page 61 identified within a particular image, what were done with those objects? Were they stored, for example, as a separate image or were they somehow linked to or associated with the image from which they were extracted? A. Typically they were linked. Q. They were linked to the source image? A. Yeah. *Id.* at 32–33

Element	'897 Claim Recitation	Exemplary Citations in Reference
		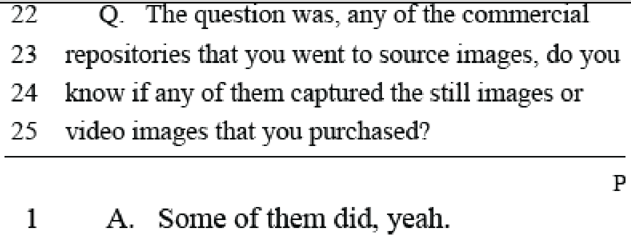 Q. The question was, any of the commercial repositories that you went to source images, do you know if any of them captured the still images or video images that you purchased? P A. Some of them did, yeah. *Id.* at 74–75

			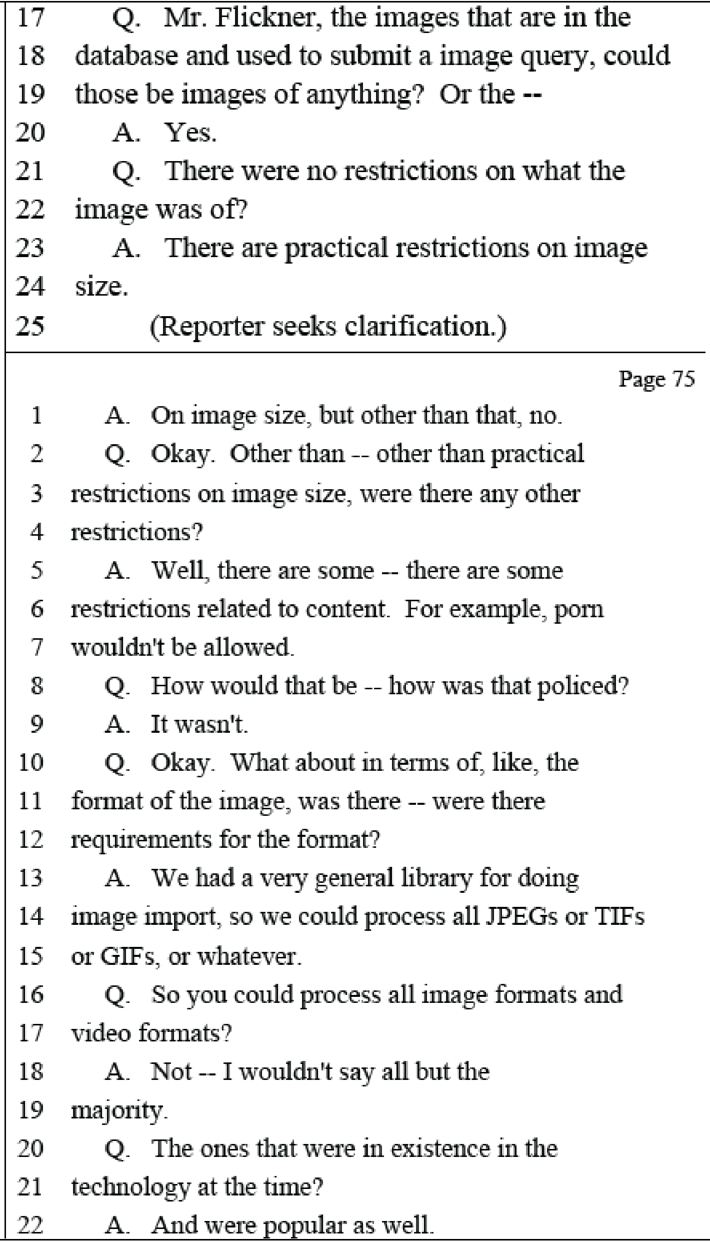 Q. Mr. Flickner, the images that are in the database and used to submit a image query, could those be images of anything? Or the -- A. Yes. Q. There were no restrictions on what the image was of? A. There are practical restrictions on image size. (Reporter seeks clarification.) Page 75 A. On image size, but other than that, no. Q. Okay. Other than -- other than practical restrictions on image size, were there any other restrictions? A. Well, there are some -- there are some restrictions related to content. For example, porn wouldn't be allowed. Q. How would that be -- how was that policed? A. It wasn't. Q. Okay. What about in terms of, like, the format of the image, was there -- were there requirements for the format? A. We had a very general library for doing image import, so we could process all JPEGs or TIFs or GIFs, or whatever. Q. So you could process all image formats and video formats? A. Not -- I wouldn't say all but the majority. Q. The ones that were in existence in the technology at the time? A. And were popular as well.

Element	'897 Claim Recitation	Exemplary Citations in Reference
		Id. at 114 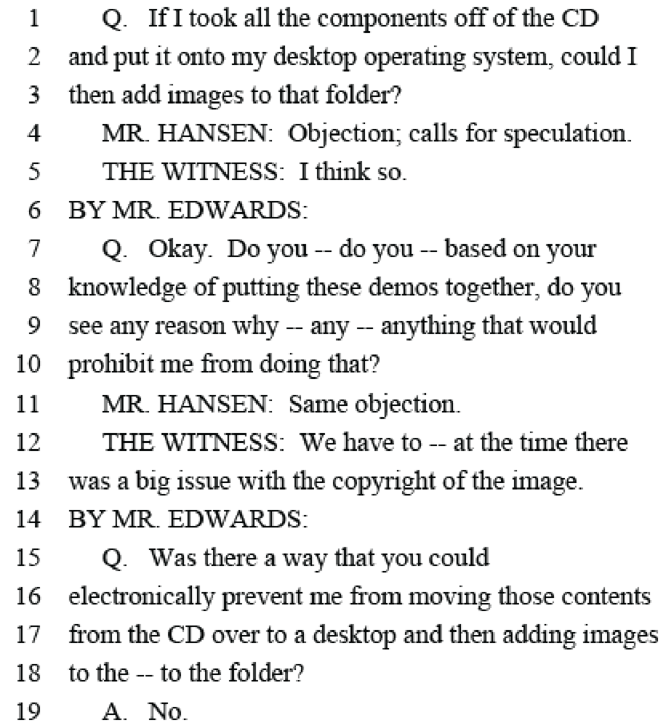 Q. If I took all the components off of the CD and put it onto my desktop operating system, could I then add images to that folder? MR. HANSEN: Objection; calls for speculation. THE WITNESS: I think so. BY MR. EDWARDS: Q. Okay. Do you -- do you -- based on your knowledge of putting these demos together, do you see any reason why -- any -- anything that would prohibit me from doing that? MR. HANSEN: Same objection. THE WITNESS: We have to -- at the time there was a big issue with the copyright of the image. BY MR. EDWARDS: Q. Was there a way that you could electronically prevent me from moving those contents from the CD over to a desktop and then adding images to the -- to the folder? A. No. *Id.* at 202

Element	’897 Claim Recitation	Exemplary Citations in Reference
		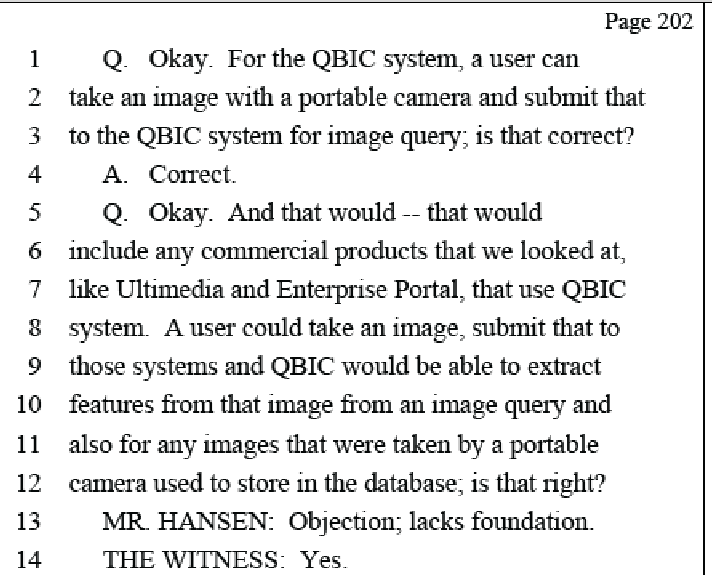 Page 202 Q. Okay. For the QBIC system, a user can take an image with a portable camera and submit that to the QBIC system for image query; is that correct? A. Correct. Q. Okay. And that would -- that would include any commercial products that we looked at, like Ultimedia and Enterprise Portal, that use QBIC system. A user could take an image, submit that to those systems and QBIC would be able to extract features from that image from an image query and also for any images that were taken by a portable camera used to store in the database; is that right? MR. HANSEN: Objection; lacks foundation. THE WITNESS: Yes. *Id.* at 300 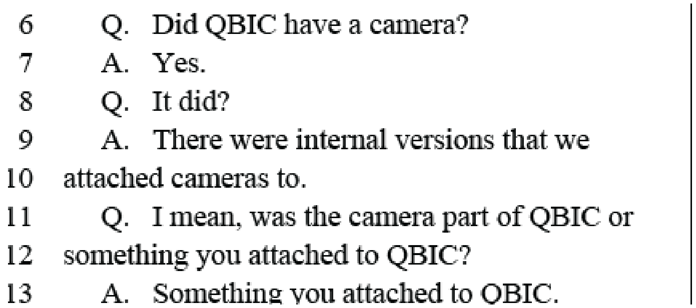 Q. Did QBIC have a camera? A. Yes. Q. It did? A. There were internal versions that we attached cameras to. Q. I mean, was the camera part of QBIC or something you attached to QBIC? A. Something you attached to QBIC. *Id.* at 312

Element	'897 Claim Recitation	Exemplary Citations in Reference
		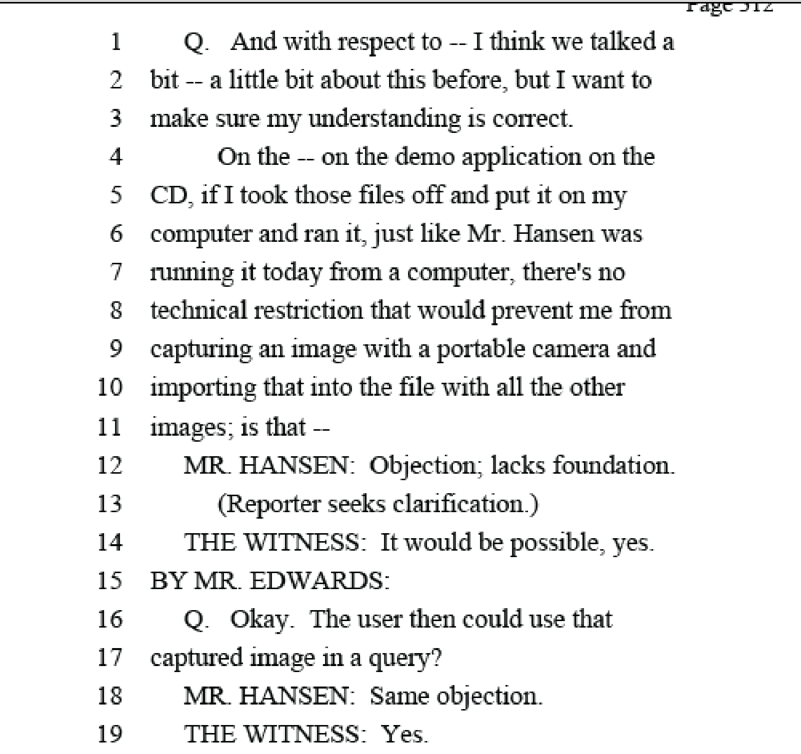 Page 312 Q. And with respect to -- I think we talked a bit -- a little bit about this before, but I want to make sure my understanding is correct. On the -- on the demo application on the CD, if I took those files off and put it on my computer and ran it, just like Mr. Hansen was running it today from a computer, there's no technical restriction that would prevent me from capturing an image with a portable camera and importing that into the file with all the other images; is that -- MR. HANSEN: Objection; lacks foundation. (Reporter seeks clarification.) THE WITNESS: It would be possible, yes. BY MR. EDWARDS: Q. Okay. The user then could use that captured image in a query? MR. HANSEN: Same objection. THE WITNESS: Yes. To the extent it is found that QBIC does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that QBIC does not anticipate this claim, QBIC renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine QBIC with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these

Element	'897 Claim Recitation	Exemplary Citations in Reference
		references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.
[25D]	enabling, via the interactive system, identification of features from the displayed image, recognition of a target based on the features, association of the target with target information pertinent to the target, and performance of a transaction based on the target information.	QBIC discloses enabling, via the interactive system, identification of features from the displayed image, recognition of a target based on the features, association of the target with target information pertinent to the target, and performance of a transaction based on the target information. The QBIC Project: Querying Images by Content Using Color, Texture, and Shape (IBM 0002390–404) at 2395-96 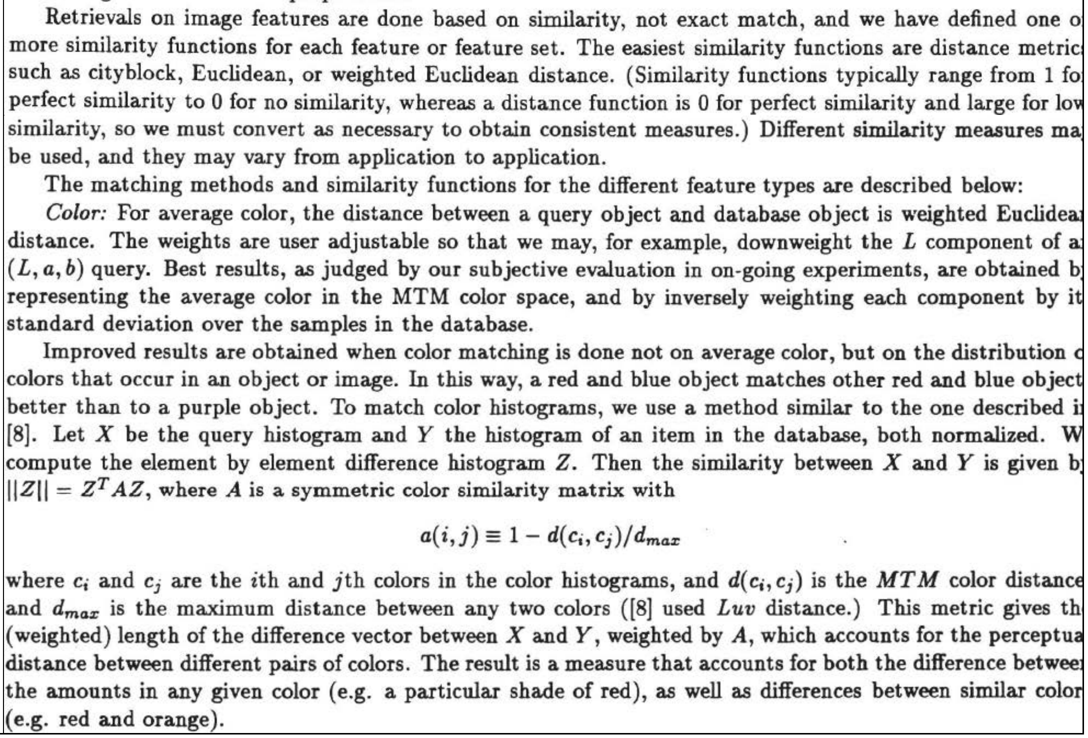 Retrievals on image features are done based on similarity, not exact match, and we have defined one o more similarity functions for each feature or feature set. The easiest similarity functions are distance metric such as cityblock, Euclidean, or weighted Euclidean distance. (Similarity functions typically range from 1 fo perfect similarity to 0 for no similarity, whereas a distance function is 0 for perfect similarity and large for low similarity, so we must convert as necessary to obtain consistent measures.) Different similarity measures ma be used, and they may vary from application to application. The matching methods and similarity functions for the different feature types are described below: *Color:* For average color, the distance between a query object and database object is weighted Euclidea distance. The weights are user adjustable so that we may, for example, downweight the L component of a (L, a, b) query. Best results, as judged by our subjective evaluation in on-going experiments, are obtained b representing the average color in the MTM color space, and by inversely weighting each component by it standard deviation over the samples in the database. Improved results are obtained when color matching is done not on average color, but on the distribution c colors that occur in an object or image. In this way, a red and blue object matches other red and blue object better than to a purple object. To match color histograms, we use a method similar to the one described i [8]. Let X be the query histogram and Y the histogram of an item in the database, both normalized. W compute the element by element difference histogram Z. Then the similarity between X and Y is given b $\|Z\| = Z^T A Z$, where A is a symmetric color similarity matrix with $$a(i,j) \equiv 1 - d(c_i, c_j)/d_{max}$$ where c_i and c_j are the ith and jth colors in the color histograms, and $d(c_i, c_j)$ is the MTM color distance and d_{max} is the maximum distance between any two colors ([8] used Luv distance.) This metric gives th (weighted) length of the difference vector between X and Y, weighted by A, which accounts for the perceptua distance between different pairs of colors. The result is a measure that accounts for both the difference betwee the amounts in any given color (e.g. a particular shade of red), as well as differences between similar color (e.g. red and orange).

Element	'897 Claim Recitation	Exemplary Citations in Reference
		The QBIC Project: Querying Images by Content Using Color, Texture, and Shape (IBM 0002390–404) at 2396 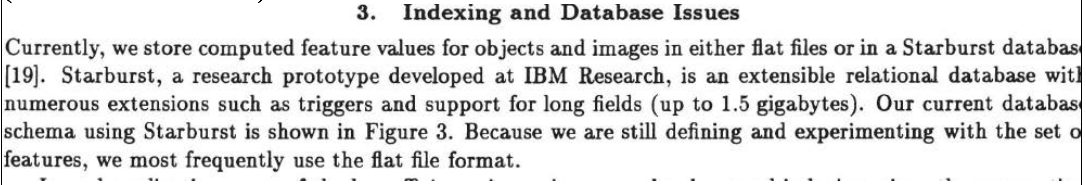 **3. Indexing and Database Issues** Currently, we store computed feature values for objects and images in either flat files or in a Starburst databas [19]. Starburst, a research prototype developed at IBM Research, is an extensible relational database wit numerous extensions such as triggers and support for long fields (up to 1.5 gigabytes). Our current databas schema using Starburst is shown in Figure 3. Because we are still defining and experimenting with the set o features, we most frequently use the flat file format. … 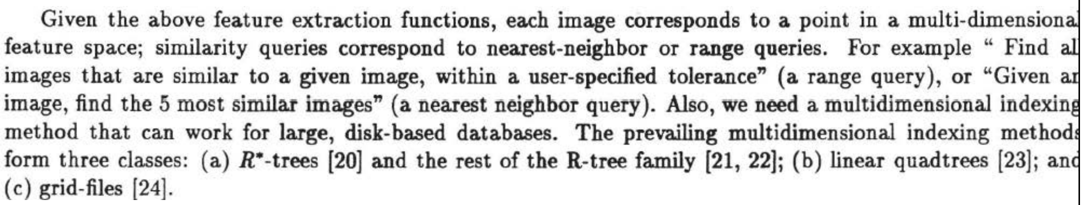 Given the above feature extraction functions, each image corresponds to a point in a multi-dimensiona feature space; similarity queries correspond to nearest-neighbor or range queries. For example " Find al images that are similar to a given image, within a user-specified tolerance" (a range query), or "Given ar image, find the 5 most similar images" (a nearest neighbor query). Also, we need a multidimensional indexing method that can work for large, disk-based databases. The prevailing multidimensional indexing method form three classes: (a) R^*-trees [20] and the rest of the R-tree family [21, 22]; (b) linear quadtrees [23]; and (c) grid-files [24]. QBIC (Query by Image Content) at 1

Element	'897 Claim Recitation	Exemplary Citations in Reference
		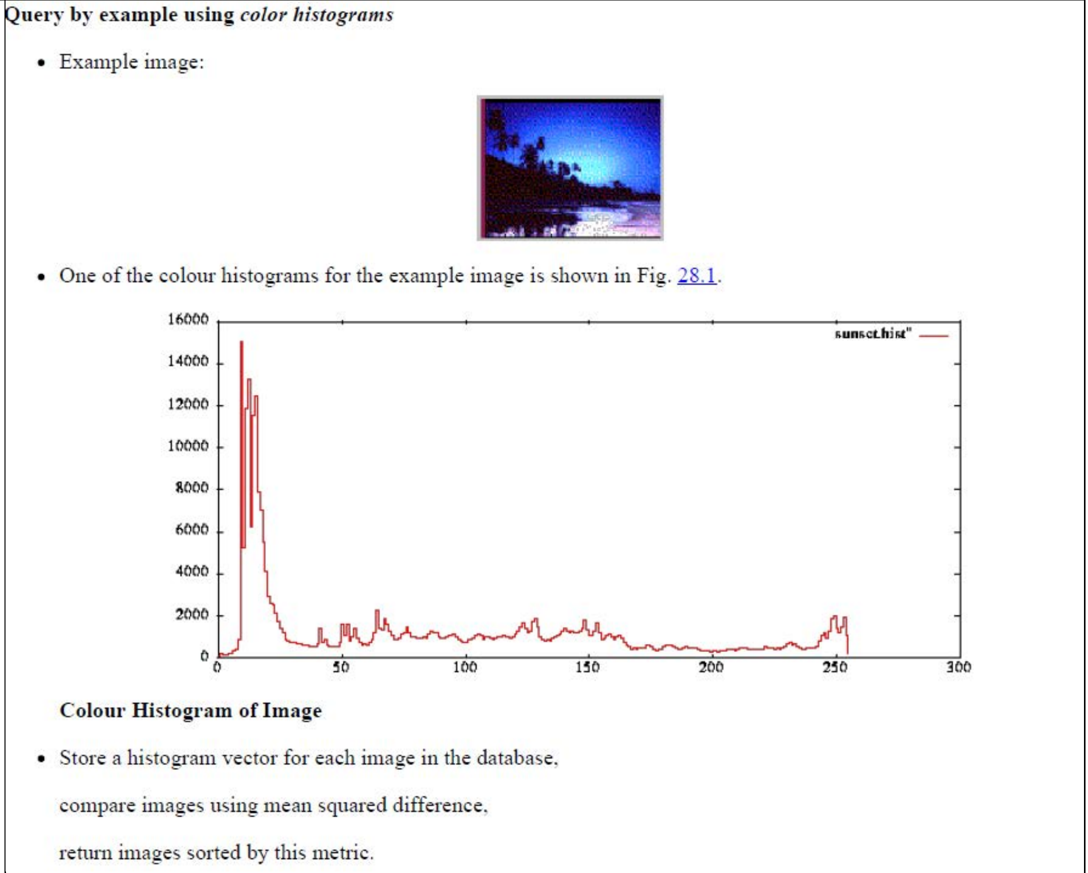 **Query by example using *color histograms*** • Example image: • One of the colour histograms for the example image is shown in Fig. 28.1. **Colour Histogram of Image** • Store a histogram vector for each image in the database, compare images using mean squared difference, return images sorted by this metric. Query by Image and Video Content: The QBIC System (IBM 0000893–902): "We have developed the QBIC (Query by Image Content) system to explore content-based retrieval methods. QBIC allows queries on large image and video databases based on • example images, • user-constructed sketches and drawings, • selected color and texture patterns, • camera and object motion, and • other graphical information." *Id*. at pp. 1–3.

Element	’897 Claim Recitation	Exemplary Citations in Reference
		“To achieve this functionality, QBIC has two main components: database population (the process of creating an image database) and database query. During the population, images and videos are processed to extract features describing their content-colors, textures, shapes, and camera and object motion-and the features are stored in a database. During the query, the user composes a query graphically. Features are generated from the graphical query and then input to a matching engine that finds images or videos from the database with similar features. Figure 2 shows the system architecture.” *Id.* at p. 3. 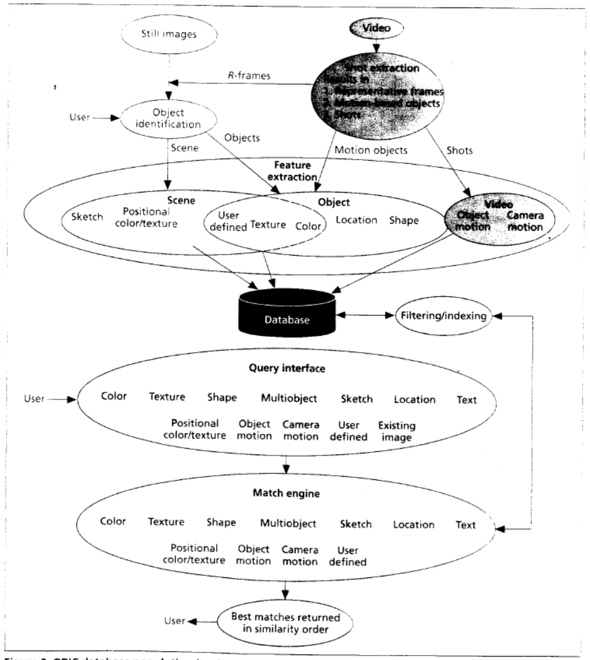 Figure 2. QBIC database population (top) and query (bottom) architecture. *Id.* at Fig. 2. “For both population and query, the QBIC data model has • still images or scenes (full images) that contain objects (subsets of an image), and • video

Element	'897 Claim Recitation	Exemplary Citations in Reference
		shots that consist of sets of contiguous frames and contain motion objects. For still images, the QBIC data model distinguishes between 'scenes' (or images) and 'objects.' A scene is an image or single representative frame of video. An object is a part of a scene-for example, the fox in Figure 3-or a moving entity in a video. For still image database population, features are extracted from images and objects and stored in a database as shown in the top left part of Figure 2. Videos are broken into clips called shots. Representative frames, or r-frames, are generated for each extracted shot. R-frames are treated as still images, and features are extracted and stored in the database. Further processing of shots generates motion objects-for example, a car moving across the screen. Queries are allowed on objects ('Find images with a red, round object'), scenes ('Find images that have approximately 30-percent red and IS-percent blue colors'), shots ('Find all shots panning from left to right'), or any combination ('Find images that have 30 percent red and contain a blue textured object')." *Id*. at p. 3. "**DATABASE POPULATION** In still image database population, the images are reduced to a standard-sized icon called a thumbnail and annotated with any available text information. Object identification is an optional but key part of this step. It lets users manually, semiautomatically, or fully automatically identify interesting regions—which we call objects—in the images. Internally, each object is represented as a binary mask. There may be an arbitrary number of objects per image. Objects can overlap and can consist of multiple disconnected components like the set of dots on a polka-dot dress. Text, like 'baby on beach,' can be associated with an outlined object or with the scene as a whole. **Object-outlining tools** Ideally, object identification would be automatic, but this is generally difficult. The alternative—manual identification—is tedious and can inhibit query-by-content applications. As a result, we have devoted considerable effort to

Element	'897 Claim Recitation	Exemplary Citations in Reference
		developing tools to aid in this step. In recent work, we have successfully used fully automatic unsupervised segmentation methods along with a foreground/background model to identify objects in a restricted class of images. The images, typical of museums and retail catalogs, have a small number of foreground objects on a generally separable background. Figure 4 shows example results. Even in this domain, robust algorithms are required because of the textured and variegated backgrounds. We also provide semiautomatic tools for identifying objects. One is an enhanced flood-fill technique. Flood-fill methods, found in most photo-editing programs, start from a single object pixel and repeatedly add adjacent pixels whose values are within some given threshold of the original pixel. Selecting the threshold, which must change from image to image and object to object, is tedious. We automatically calculate a dynamic threshold by having the user click on background as well as object points. For reasonably uniform objects that are distinct from the background, this operation allows fast object identification without manually adjusting a threshold. The example in Figure 3 shows an object, a fox, identified by using only a few clicks. We designed another outlining tool to help users track object edges. This tool takes a user-drawn curve and automatically aligns it with nearby image edges. Based on the 'snakes' concept developed in recent computer vision research, the tool finds the curve that maximizes the image gradient magnitude along the curve. The spline snake formulation we use allows for smooth solutions to the resulting nonlinear minimization problem. The computation is done at interactive speeds so that, as the user draws a curve, it is 'rubber-banded' to lie along object boundaries. . . . **Video data** For video data, database population has three major components: • shot detection, • representative frame creation for each shot, and • derivation of a layered representation of coherently moving structures/ objects. . . . Once the shot boundaries have been detected, each shot is represented using

Element	'897 Claim Recitation	Exemplary Citations in Reference
		an r-frame. R-frames are used for several purposes. First, during database population, r-frames are treated as still images in which objects can be identified by using the previously described methods." *Id.* at pp. 3–6. "**SAMPLE QUERIES** For each full-scene image, identified image object, r-frame, and identified video object resulting from the above processing, a set of features is computed to allow content-based queries. The features are computed and stored during database population. We present a brief description of the features and the associated queries. Mathematical details on the features and matching methods can be found in Ashley et al. and Niblack et al. Average color queries let users find images or objects that are similar to a selected color, say from a color wheel, or to the color of an object. The feature used in the query is a 3D vector of Munsell color coordinates. Histogram color queries return items with matching color distributions—say, a fabric pattern with approximately 40 percent red and 20 percent blue. For this case, the underlying feature is a 256-element histogram computed over a quantized version of the color space. Figure 7 shows a histogram query on still images and a color query on video r-frames. Note that in the query specification for the histogram query of Figure 7, the user has selected percentages of two colors (blue and white) by adjusting sliders. Using such a query, an advertising agent could, for example, search for a picture of a beach scene, one predominantly blue (for sky and water) and white (for sand and clouds); or find images with similar color spreads for a uniform ad campaign. The average color query demonstrates a query against a video shot database where the user is searching for red r-frames. Again, the query specification is on the left and the best hits are on the right. Figure 8 shows an example texture query. In this case, the query is specified by selecting from a sampler—a set of prestored example images. The underlying texture features are mathematical representations of coarseness, contrast, and directionality features. Coarseness measures the scale of a texture

Element	'897 Claim Recitation	Exemplary Citations in Reference
		(pebbles versus boulders), contrast describes its vividness, and directionality describes whether it has a favored direction (like grass) or not (like a smooth object). An object shape query is shown in Figure 8. In this case, the query specification is the drawn shape on the left. Area, circularity, eccentricity, major-axis direction, features derived from the object moments, and a set of tangent angles around the object perimeter are the features used to characterize and match shapes. Figure 9 illustrates query by sketch. In this case, the query specification is a freehand drawing of the dominant lines and edges in the image. The sketch feature is an automatically extracted reduced-resolution 'edge map.' Matching is done by using a template-matching technique. A multiobject query asking for images that contain both a red round object and a green textured object is shown in the bottom of Figure 9. The features are standard color and texture. The matching is done by combining the color and texture distances. Combining distances is applied to arbitrary sets of objects and features to implement logical And semantics. . . . . **ANNOTATION AND DATABASE POPULATION TOOLS**. Automatic methods (such as our Positional Color query) that don't rely on object identification, methods that identify objects automatically as in the museum image example, fast and easy-to-use semiautomatic outlining tools, and motion-based segmentation algorithms will enable additional application areas. **FEATURE EXTRACTION AND MATCHING METHODS**. New mathematical representations of video, image, and object attributes that capture 'interesting' features for retrieval are needed. Features that describe new image properties such as alternate texture measures or that are based on fractals or wavelet representations, for example, may offer advantages of representation, indexability, and ease of similarity matching. **INTEGRATION WITH TEXT AND PARAMETRIC ANNOTATION**.

Element	'897 Claim Recitation	Exemplary Citations in Reference
		Query by visual content complements and extends existing query methods. Systems must be able to integrate queries combining date, subject matter, price, and availability with content properties such as color, texture, and shape." *Id.* at pp. 7–8. Image, Audio, and Video Extenders Administration and Programming (IBM 000001–617): "DB2 (DB2) Universal Database (UDB) is a powerful, object-relational database manager. It stores and protects traditional numeric and character data, as well as large, complex objects (LOBs). DB2 extenders help you exploit DB2 's object-relational features. The extenders define distinct data types and special functions for image, audio, video, and text objects. … Because of that, the extenders give your applications a single point of access to any or all of these types of data, along with traditional numeric and character data." *Id.* at p. 3. "Your applications can even search for images by content. Imagine an application that uses visual examples to search for images. With such an application, users could select an example image and have the application find other images that have colors or textures similar to those in the example. With DB2 extenders' Query by Image Content (QBIC) capability, you can create applications that search for images in this visual way." *Id.* at p. 4. "A QBIC catalog is a set of files that hold data about the visual features of images. The Image Extender uses this data to search for images by content. You create a QBIC catalog for each column of images in a user table that you want to make available for searching by content. When you create a QBIC catalog you identify the features for which you want the Image Extender to analyze, store, and later query data. You can also add or drop features from a QBIC catalog after the catalog is created. A QBIC catalog can hold data for the following image features:

Element	'897 Claim Recitation	Exemplary Citations in Reference
		Average color The sum of the color values for all pixels in an image divided by the number of pixels in the image. (A pixel is the smallest element of an image that can be assigned color and intensity.) For example, if 50% of an image consists of blue pixels and the other 50% red pixels, the image has an average color value of purple. Average color is used to search for images that have a predominant color. If an image has a predominant color, the average color will be similar to the predominant color. **Histogram color** Measures the distribution of colors in an image against a spectrum of 64 colors. For each of the 64 colors, histogram color identifies the percentage of pixels in an image that have that color. For example, the histogram color of an image might be 40% white pixels, 50% blue, and 10% red; none of the pixels in the image have any of the remaining colors in the histogram spectrum. Histogram color is used to search for images that have a variety of colors. **Positional color** The average color value for the pixels in a specified area in an image. For example, the upper right-hand corner of an image might show a bright yellow sun; the positional color of this area of the image is bright yellow. Positional color is used to search for images that have a predominant color in a particular area. **Texture** Measures the coarseness, contrast, and directionality of an image. Coarseness indicates the size of repeating items in an image (for example, pebbles versus boulders). Contrast identifies the brightness variations in an image (light versus dark). Directionality indicates whether a direction predominates in an image (as in the vertical direction of a picket fence) or does not predominate (as in an image of sand). Texture is used to search for images that have a particular pattern. To make an image available for searching by content, you catalog the image. When you catalog an image, the Image Extender analyzes the image, by

Element	'897 Claim Recitation	Exemplary Citations in Reference
		computing the feature values for the image, and stores the values in a QBIC catalog. When you search for an image by content, your query identifies one or more features for the search (such as average color), a source for each feature (such as an example image), and a target set of cataloged images. The Image Extender computes the feature value of the source and compares it to the cataloged feature values for the target images. It then computes a score that indicates how similar the feature values of the target images are to the source. You can have the Image Extender return the images whose features are most similar to the source. The Image Extender will return the handle of each image and the image score. You can also have the Image Extender return only the score of a single image." *Id.* at pp. 20–21. 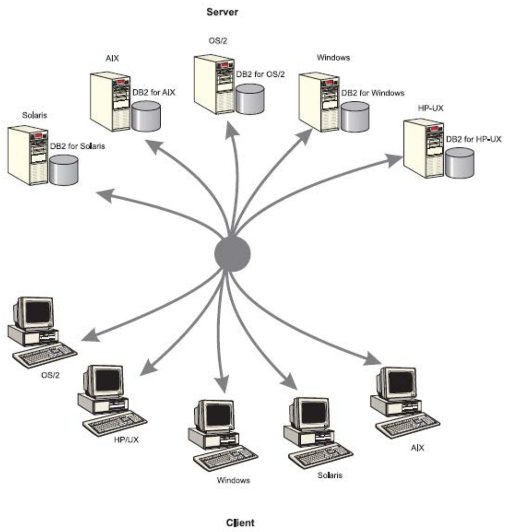 *Figure 6. DB2 extender platforms* *Id.* at p. 12 (Fig. 6). "The DB2 Extenders run in the DB2 client/server environment. This environment consists of a database server and one or more remote database

Element	'897 Claim Recitation	Exemplary Citations in Reference
		clients. The DB2 extender services run on the server. Before you can access them, you have to start them." *Id.* at p. 47. Managing Enterprise Information Portal (IBM 000777–880): "Many paper-intensive enterprises, such as insurance companies and financial institutions, administer large volumes of business-related content. The need for an enterprise solution for managing and accessing business information spans many industries. A content server is a software system that stores multimedia, business forms, documents, and related data, along with metadata that allows employees to process and work with the content." *Id.* at p. 3. "The image search server uses IBM's QBIC (query by image content) technology to help you search for objects by certain visual properties, such as color and texture. The image search server analyzes images and stores the image information in a database. Then users can run image queries, which use the visual properties of images, to match colors, textures, and their positions without describing them in words. You can combine content-based queries with text and keyword searches for more powerful retrieval of image and multimedia data." *Id.* at p. 63.

Element	'897 Claim Recitation	Exemplary Citations in Reference
		QBIC Demo (IBM000747) 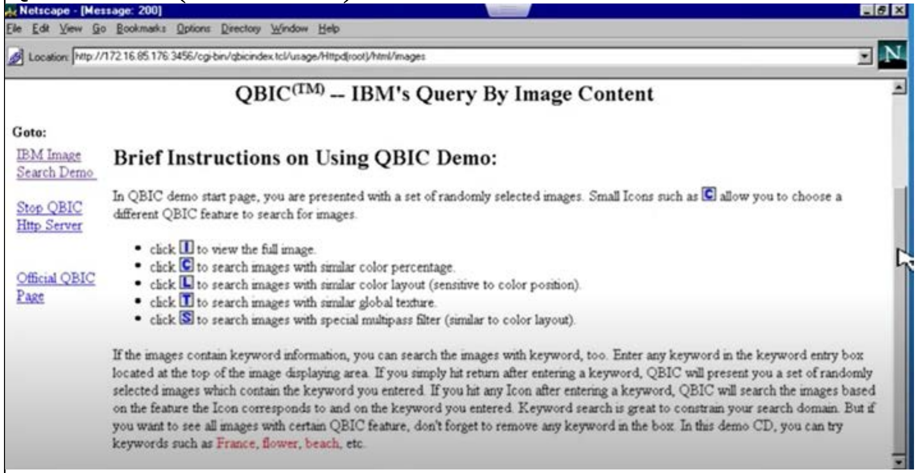 Netscape - [Message: 200] File Edit View Go Bookmarks Options Directory Window Help Location: http://172.16.85.176:3456/cgi-bin/qbicindex.tcl/usage/Httpd(root)/html/images QBIC(TM) -- IBM's Query By Image Content Goto: IBM Image Search Demo Stop QBIC Http Server Official QBIC Page **Brief Instructions on Using QBIC Demo:** In QBIC demo start page, you are presented with a set of randomly selected images. Small Icons such as C allow you to choose a different QBIC feature to search for images. - click I to view the full image. - click C to search images with similar color percentage. - click L to search images with similar color layout (sensitive to color position). - click T to search images with similar global texture. - click S to search images with special multipass filter (similar to color layout). If the images contain keyword information, you can search the images with keyword, too. Enter any keyword in the keyword entry box located at the top of the image displaying area. If you simply hit return after entering a keyword, QBIC will present you a set of randomly selected images which contain the keyword you entered. If you hit any Icon after entering a keyword, QBIC will search the images based on the feature the Icon corresponds to and on the keyword you entered. Keyword search is great to constrain your search domain. But if you want to see all images with certain QBIC feature, don't forget to remove any keyword in the box. In this demo CD, you can try keywords such as France, flower, beach, etc.  Netscape - [Message: 200] File Edit View Go Bookmarks Options Directory Window Help Back Forward Home Edit Reload Images Open Print Find Stop Location: http://172.16.85.176:3456/cgi-bin/ibm.tcl/search?fname=Color+Layout&france77%2Ejpg.x=0 IBM QBIC™ © IBM Corporation Usage: I Get Info C Color Histogram L Layout T Texture S Special Hybrid Keywords: Previous Next How to Use QBIC Back To Start Page Applet ResultApplet running Start qbic QBICTCL Netscape - [Message... Netscape - [http://wwwqbi... 1:47 PM QBIC Demo (IBM000747)

Element	'897 Claim Recitation	Exemplary Citations in Reference
		 Location: http://172.16.85.176:3456/cgi-bin/ibm.tcl/histpicker Keywords: Red: 59 Green: 131 Blue: 59 9% (255,255,255) 8% (59,131,59) 4% (106,182,106) To specify your color histogram query, "paint" a graphical representation of your pattern of colors and approximate percentages that you want. First select a *percentage* by clicking in the painting area to block off a rectangle of the desired size (percentage of the full area). All new regions are painted white, initially. Then select a color from the QBIC Demo (IBM000747) 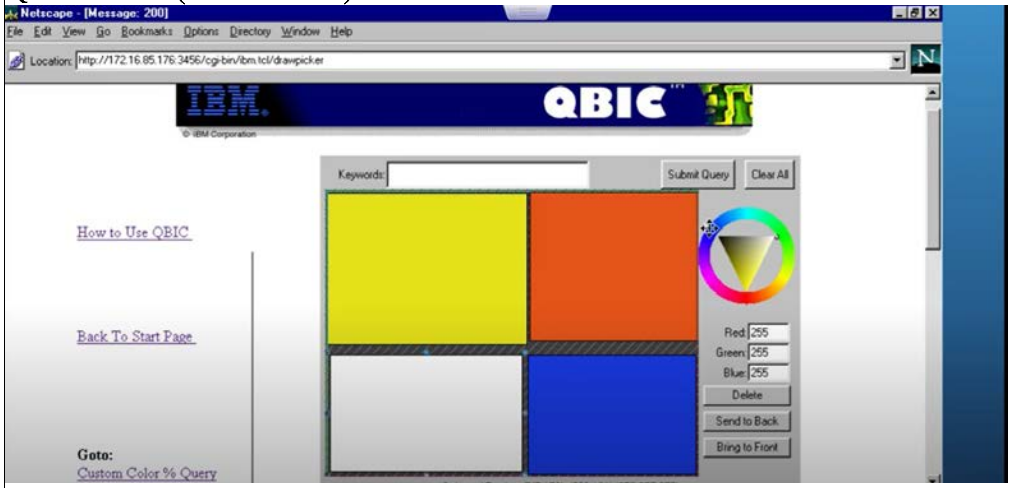 Location: http://172.16.85.176:3456/cgi-bin/ibm.tcl/drawpicker Keywords: Red: 255 Green: 255 Blue: 255

Element	'897 Claim Recitation	Exemplary Citations in Reference
		 Netscape - [Message: 200] Location: http://172.16.85.176:3456/cgi-bin/ibm.tcl/search?fname=HybridColor&france10%2Ejpg.x=0 Goto: Custom Color % Query Custom Paint Query Random Images Official QBIC Page Query was: Example: =france10.jpg Query Type: QbHybridColorFeatureClass 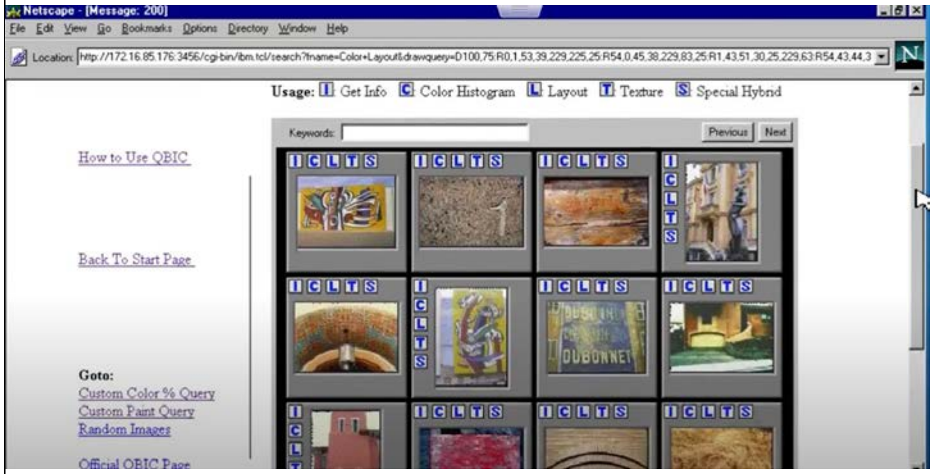 Netscape - [Message: 200] Location: http://172.16.85.176:3456/cgi-bin/ibm.tcl/search?fname=Color+Layout&drawquery=D100,75:R0,1,53,39,229,225,25:R54,0,45,38,229,83,25:R1,43,51,30,25,229,63:R54,43,44,3 Usage: I Get Info C Color Histogram L Layout T Texture S Special Hybrid Keywords: Previous Next How to Use QBIC Back To Start Page Goto: Custom Color % Query Custom Paint Query Random Images Official QBIC Page

Element	'897 Claim Recitation	Exemplary Citations in Reference
		 Location: http://172.16.85.176:3456/cgi-bin/ibm.tcl/info?france50%2Ejpg Image france50.jpg Flickner Depo at 54-55

Element	'897 Claim Recitation	Exemplary Citations in Reference
		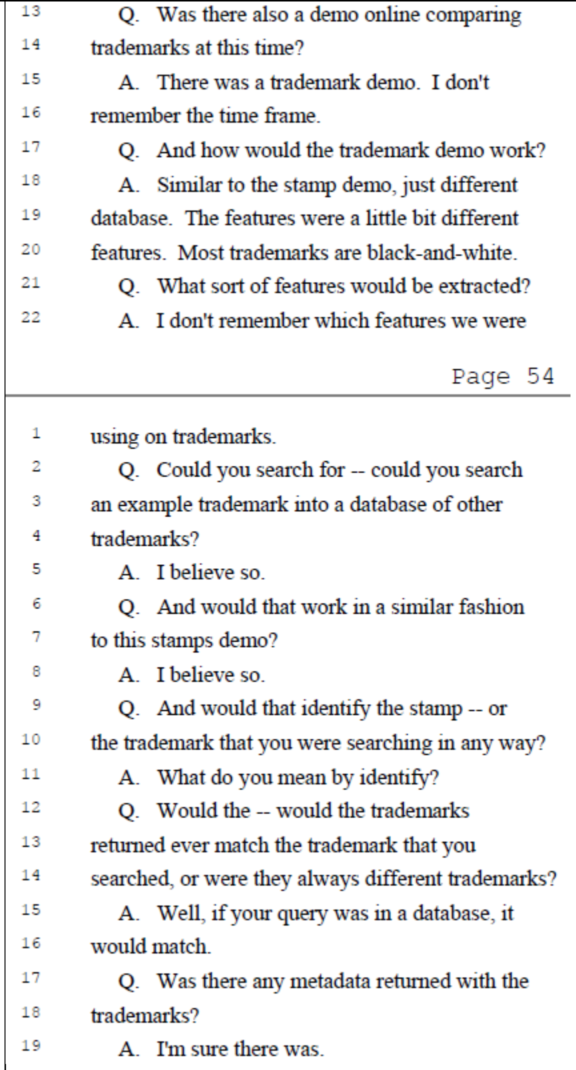 Q. Was there also a demo online comparing trademarks at this time? A. There was a trademark demo. I don't remember the time frame. Q. And how would the trademark demo work? A. Similar to the stamp demo, just different database. The features were a little bit different features. Most trademarks are black-and-white. Q. What sort of features would be extracted? A. I don't remember which features we were Page 54 using on trademarks. Q. Could you search for -- could you search an example trademark into a database of other trademarks? A. I believe so. Q. And would that work in a similar fashion to this stamps demo? A. I believe so. Q. And would that identify the stamp -- or the trademark that you were searching in any way? A. What do you mean by identify? Q. Would the -- would the trademarks returned ever match the trademark that you searched, or were they always different trademarks? A. Well, if your query was in a database, it would match. Q. Was there any metadata returned with the trademarks? A. I'm sure there was. Flickner Depo at 89-90

Element	'897 Claim Recitation	Exemplary Citations in Reference
		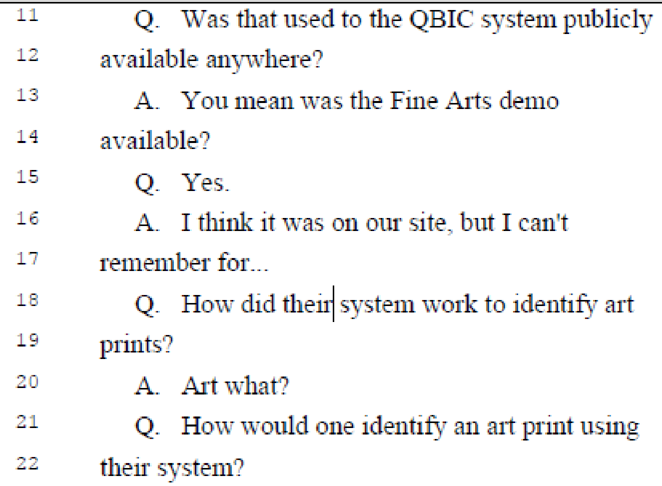 Q. Was that used to the QBIC system publicly available anywhere? A. You mean was the Fine Arts demo available? Q. Yes. A. I think it was on our site, but I can't remember for... Q. How did their system work to identify art prints? A. Art what? Q. How would one identify an art print using their system? 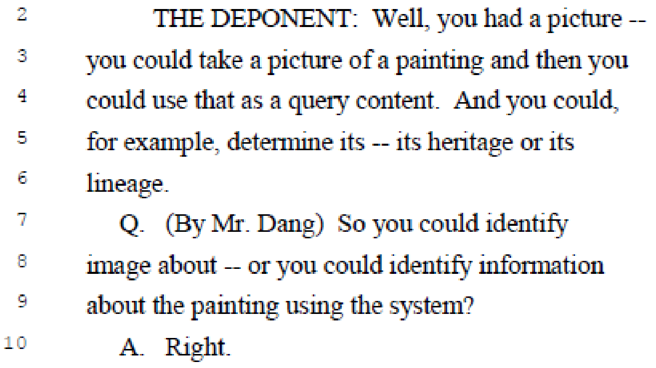 THE DEPONENT: Well, you had a picture -- you could take a picture of a painting and then you could use that as a query content. And you could, for example, determine its -- its heritage or its lineage. Q. (By Mr. Dang) So you could identify image about -- or you could identify information about the painting using the system? A. Right. Flickner Depo at 88

Element	'897 Claim Recitation	Exemplary Citations in Reference
		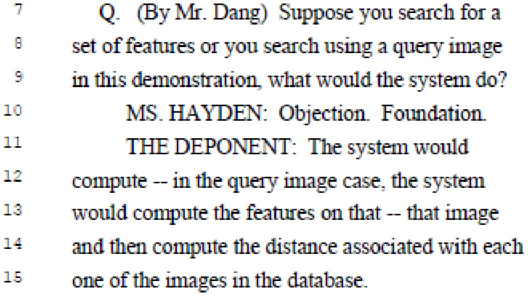 Q. (By Mr. Dang) Suppose you search for a set of features or you search using a query image in this demonstration, what would the system do? MS. HAYDEN: Objection. Foundation. THE DEPONENT: The system would compute -- in the query image case, the system would compute the features on that -- that image and then compute the distance associated with each one of the images in the database. Flickner Depo at 40-42

		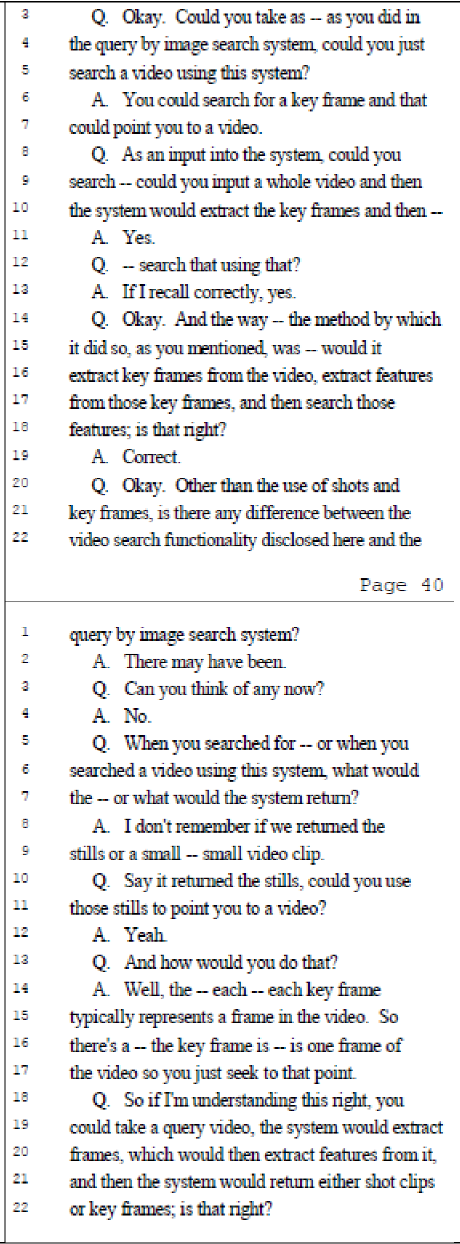 Q. Okay. Could you take as -- as you did in the query by image search system, could you just search a video using this system? A. You could search for a key frame and that could point you to a video. Q. As an input into the system, could you search -- could you input a whole video and then the system would extract the key frames and then -- A. Yes. Q. -- search that using that? A. If I recall correctly, yes. Q. Okay. And the way -- the method by which it did so, as you mentioned, was -- would it extract key frames from the video, extract features from those key frames, and then search those features; is that right? A. Correct. Q. Okay. Other than the use of shots and key frames, is there any difference between the video search functionality disclosed here and the Page 40 query by image search system? A. There may have been. Q. Can you think of any now? A. No. Q. When you searched for -- or when you searched a video using this system, what would the -- or what would the system return? A. I don't remember if we returned the stills or a small -- small video clip. Q. Say it returned the stills, could you use those stills to point you to a video? A. Yeah. Q. And how would you do that? A. Well, the -- each -- each key frame typically represents a frame in the video. So there's a -- the key frame is -- is one frame of the video so you just seek to that point. Q. So if I'm understanding this right, you could take a query video, the system would extract frames, which would then extract features from it, and then the system would return either shot clips or key frames; is that right?

Element	'897 Claim Recitation	Exemplary Citations in Reference
		 A. Right. Flickner Depo at 87. 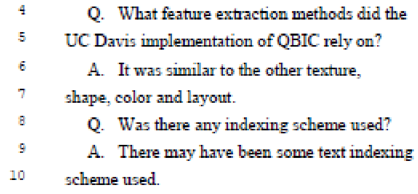 Q. What feature extraction methods did the UC Davis implementation of QBIC rely on? A. It was similar to the other texture, shape, color and layout. Q. Was there any indexing scheme used? A. There may have been some text indexing scheme used. Flickner Depo at 78-79

Element	'897 Claim Recitation	Exemplary Citations in Reference
		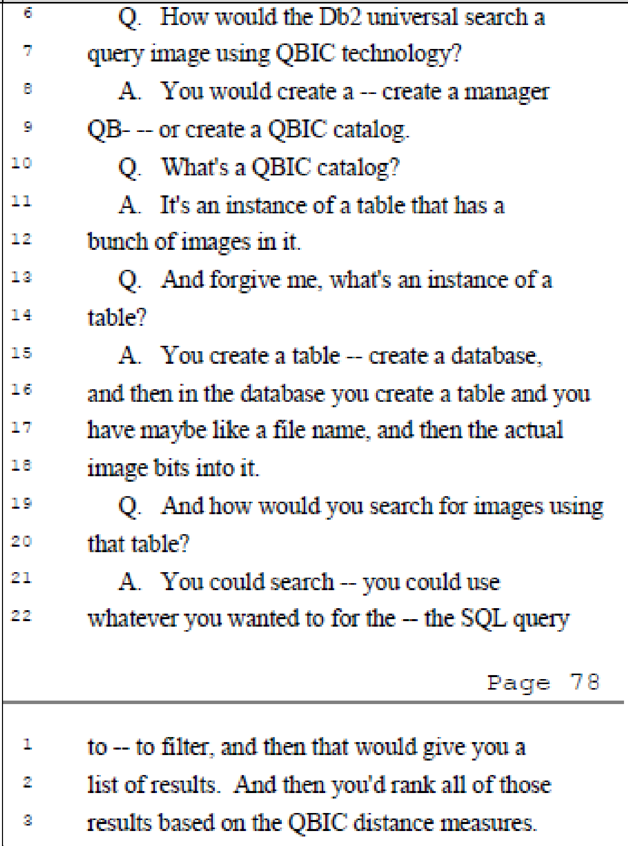 Q. How would the Db2 universal search a query image using QBIC technology? A. You would create a -- create a manager QB- -- or create a QBIC catalog. Q. What's a QBIC catalog? A. It's an instance of a table that has a bunch of images in it. Q. And forgive me, what's an instance of a table? A. You create a table -- create a database, and then in the database you create a table and you have maybe like a file name, and then the actual image bits into it. Q. And how would you search for images using that table? A. You could search -- you could use whatever you wanted to for the -- the SQL query Page 78 to -- to filter, and then that would give you a list of results. And then you'd rank all of those results based on the QBIC distance measures. Flickner Depo at 201-202

Element	'897 Claim Recitation	Exemplary Citations in Reference
		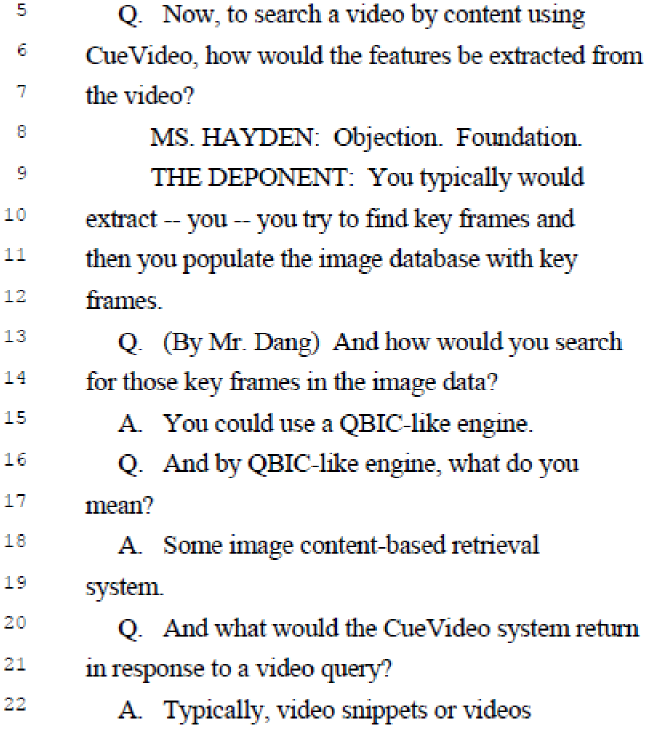 Q. Now, to search a video by content using CueVideo, how would the features be extracted from the video? MS. HAYDEN: Objection. Foundation. THE DEPONENT: You typically would extract -- you -- you try to find key frames and then you populate the image database with key frames. Q. (By Mr. Dang) And how would you search for those key frames in the image data? A. You could use a QBIC-like engine. Q. And by QBIC-like engine, what do you mean? A. Some image content-based retrieval system. Q. And what would the CueVideo system return in response to a video query? A. Typically, video snippets or videos

Element	'897 Claim Recitation	Exemplary Citations in Reference
		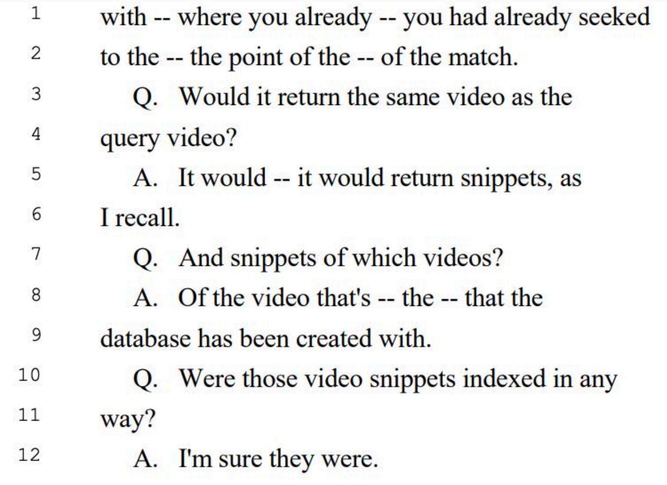with -- where you already -- you had already seeked to the -- the point of the -- of the match. Q. Would it return the same video as the query video? A. It would -- it would return snippets, as I recall. Q. And snippets of which videos? A. Of the video that's -- the -- that the database has been created with. Q. Were those video snippets indexed in any way? A. I'm sure they were. Flickner Depo at 22-25

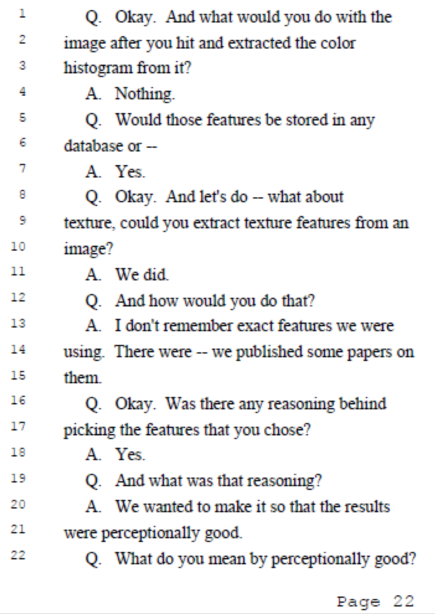

Q. Okay. And what would you do with the image after you hit and extracted the color histogram from it?
A. Nothing.
Q. Would those features be stored in any database or --
A. Yes.
Q. Okay. And let's do -- what about texture, could you extract texture features from an image?
A. We did.
Q. And how would you do that?
A. I don't remember exact features we were using. There were -- we published some papers on them.
Q. Okay. Was there any reasoning behind picking the features that you chose?
A. Yes.
Q. And what was that reasoning?
A. We wanted to make it so that the results were perceptionally good.
Q. What do you mean by perceptionally good?

Page 22

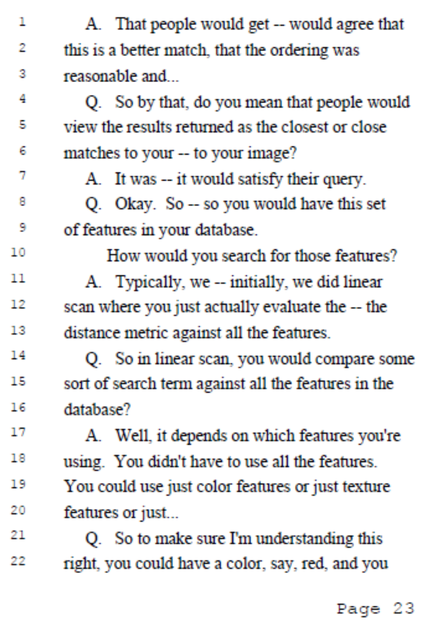

A. That people would get -- would agree that this is a better match, that the ordering was reasonable and...
Q. So by that, do you mean that people would view the results returned as the closest or close matches to your -- to your image?
A. It was -- it would satisfy their query.
Q. Okay. So -- so you would have this set of features in your database.
How would you search for those features?
A. Typically, we -- initially, we did linear scan where you just actually evaluate the -- the distance metric against all the features.
Q. So in linear scan, you would compare some sort of search term against all the features in the database?
A. Well, it depends on which features you're using. You didn't have to use all the features. You could use just color features or just texture features or just...
Q. So to make sure I'm understanding this right, you could have a color, say, red, and you

Page 23

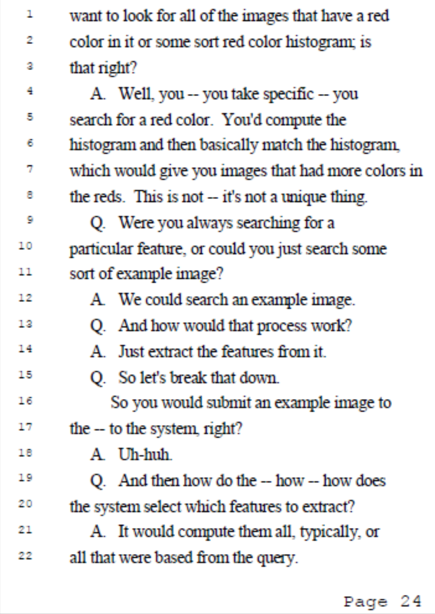

want to look for all of the images that have a red color in it or some sort red color histogram; is that right?
A. Well, you -- you take specific -- you search for a red color. You'd compute the histogram and then basically match the histogram, which would give you images that had more colors in the reds. This is not -- it's not a unique thing.
Q. Were you always searching for a particular feature, or could you just search some sort of example image?
A. We could search an example image.
Q. And how would that process work?
A. Just extract the features from it.
Q. So let's break that down.
So you would submit an example image to the -- to the system, right?
A. Uh-huh.
Q. And then how do the -- how -- how does the system select which features to extract?
A. It would compute them all, typically, or all that were based from the query.

Page 24

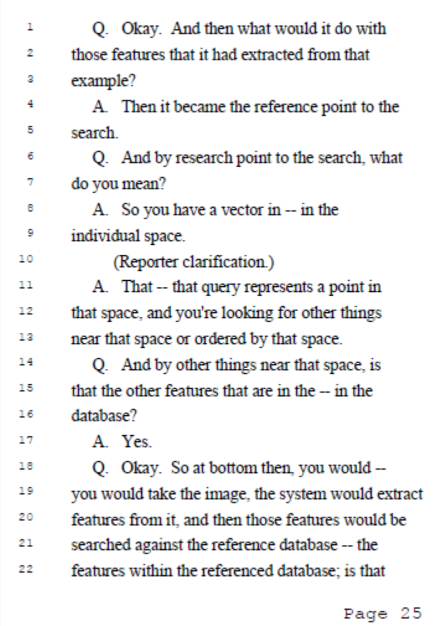

Q. Okay. And then what would it do with those features that it had extracted from that example?
A. Then it became the reference point to the search.
Q. And by research point to the search, what do you mean?
A. So you have a vector in -- in the individual space.
(Reporter clarification.)
A. That -- that query represents a point in that space, and you're looking for other things near that space or ordered by that space.
Q. And by other things near that space, is that the other features that are in the -- in the database?
A. Yes.
Q. Okay. So at bottom then, you would -- you would take the image, the system would extract features from it, and then those features would be searched against the reference database -- the features within the referenced database; is that

Page 25

Element	'897 Claim Recitation	Exemplary Citations in Reference
		Flickner Depo at 159 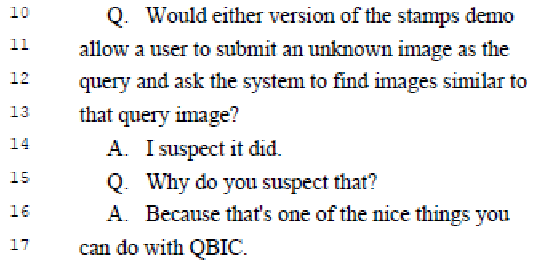 Q. Would either version of the stamps demo allow a user to submit an unknown image as the query and ask the system to find images similar to that query image? A. I suspect it did. Q. Why do you suspect that? A. Because that's one of the nice things you can do with QBIC. Flickner Depo at 209-210 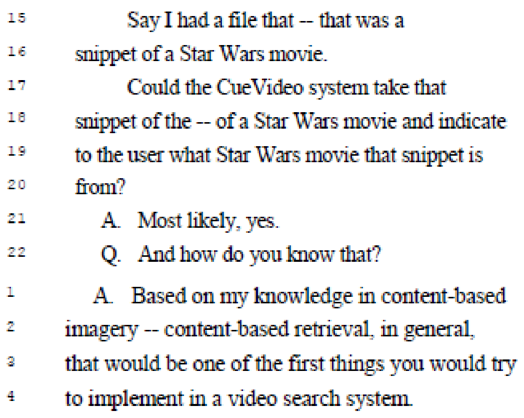 Q. Okay. Let me ask -- let me ask that question using -- using an example. Say I had a file that -- that was a snippet of a Star Wars movie. Could the CueVideo system take that snippet of the -- of a Star Wars movie and indicate to the user what Star Wars movie that snippet is from? A. Most likely, yes. Q. And how do you know that? A. Based on my knowledge in content-based imagery -- content-based retrieval, in general, that would be one of the first things you would try to implement in a video search system. Flickner Depo at 21

Element	'897 Claim Recitation	Exemplary Citations in Reference
		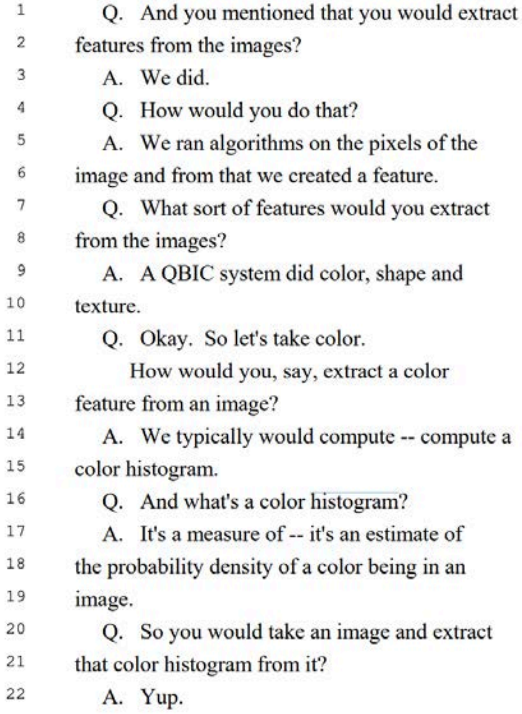 Q. And you mentioned that you would extract features from the images? A. We did. Q. How would you do that? A. We ran algorithms on the pixels of the image and from that we created a feature. Q. What sort of features would you extract from the images? A. A QBIC system did color, shape and texture. Q. Okay. So let's take color. How would you, say, extract a color feature from an image? A. We typically would compute -- compute a color histogram. Q. And what's a color histogram? A. It's a measure of -- it's an estimate of the probability density of a color being in an image. Q. So you would take an image and extract that color histogram from it? A. Yup. Flickner Depo at 36 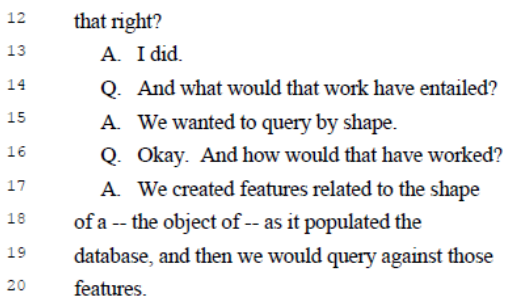 Q. So underlying this paper, you mentioned you did more work on the shape side of things; is that right? A. I did. Q. And what would that work have entailed? A. We wanted to query by shape. Q. Okay. And how would that have worked? A. We created features related to the shape of a -- the object of -- as it populated the database, and then we would query against those features.

Element	'897 Claim Recitation	Exemplary Citations in Reference
		Flickner Depo at 38-41 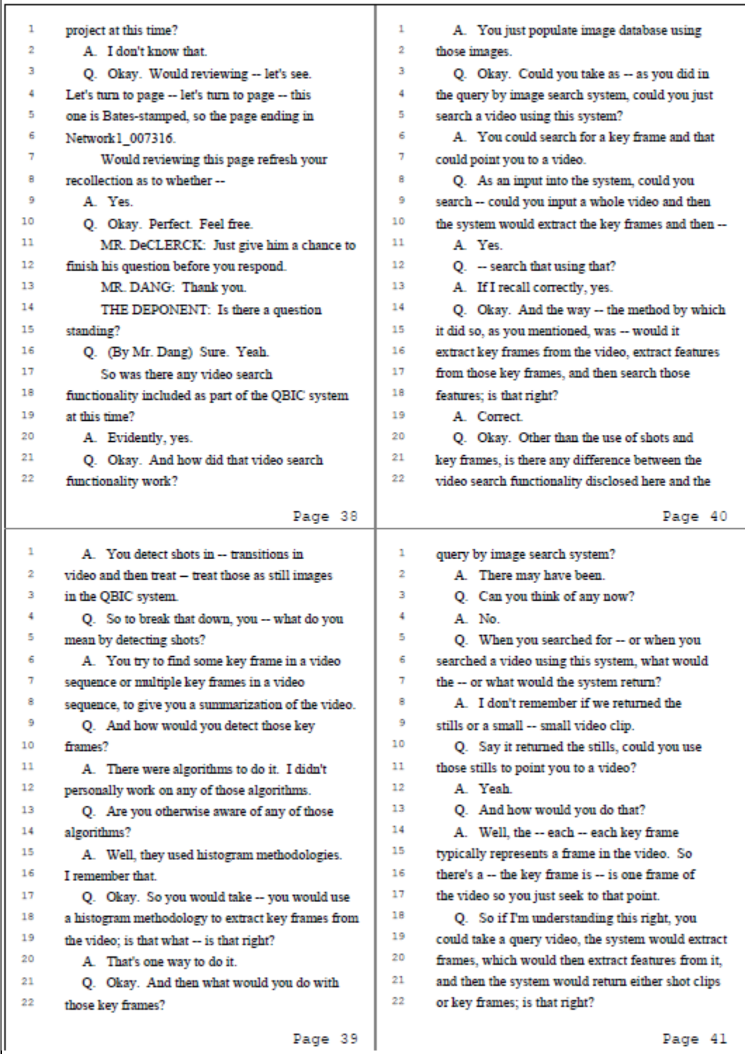

project at this time?
A. I don't know that.
Q. Okay. Would reviewing -- let's see. Let's turn to page -- let's turn to page -- this one is Bates-stamped, so the page ending in Network1_007316.
Would reviewing this page refresh your recollection as to whether --
A. Yes.
Q. Okay. Perfect. Feel free.
MR. DeCLERCK: Just give him a chance to finish his question before you respond.
MR. DANG: Thank you.
THE DEPONENT: Is there a question standing?
Q. (By Mr. Dang) Sure. Yeah.
So was there any video search functionality included as part of the QBIC system at this time?
A. Evidently, yes.
Q. Okay. And how did that video search functionality work?

Page 38

A. You detect shots in -- transitions in video and then treat -- treat those as still images in the QBIC system.
Q. So to break that down, you -- what do you mean by detecting shots?
A. You try to find some key frame in a video sequence or multiple key frames in a video sequence, to give you a summarization of the video.
Q. And how would you detect those key frames?
A. There were algorithms to do it. I didn't personally work on any of those algorithms.
Q. Are you otherwise aware of any of those algorithms?
A. Well, they used histogram methodologies. I remember that.
Q. Okay. So you would take -- you would use a histogram methodology to extract key frames from the video; is that what -- is that right?
A. That's one way to do it.
Q. Okay. And then what would you do with those key frames?

Page 39

A. You just populate image database using those images.
Q. Okay. Could you take as -- as you did in the query by image search system, could you just search a video using this system?
A. You could search for a key frame and that could point you to a video.
Q. As an input into the system, could you search -- could you input a whole video and then the system would extract the key frames and then --
A. Yes.
Q. -- search that using that?
A. If I recall correctly, yes.
Q. Okay. And the way -- the method by which it did so, as you mentioned, was -- would it extract key frames from the video, extract features from those key frames, and then search those features; is that right?
A. Correct.
Q. Okay. Other than the use of shots and key frames, is there any difference between the video search functionality disclosed here and the

Page 40

query by image search system?
A. There may have been.
Q. Can you think of any now?
A. No.
Q. When you searched for -- or when you searched a video using this system, what would the -- or what would the system return?
A. I don't remember if we returned the stills or a small -- small video clip.
Q. Say it returned the stills, could you use those stills to point you to a video?
A. Yeah.
Q. And how would you do that?
A. Well, the -- each -- each key frame typically represents a frame in the video. So there's a -- the key frame is -- is one frame of the video so you just seek to that point.
Q. So if I'm understanding this right, you could take a query video, the system would extract frames, which would then extract features from it, and then the system would return either shot clips or key frames; is that right?

Page 41

Element	'897 Claim Recitation	Exemplary Citations in Reference
		Flickner Depo at 50 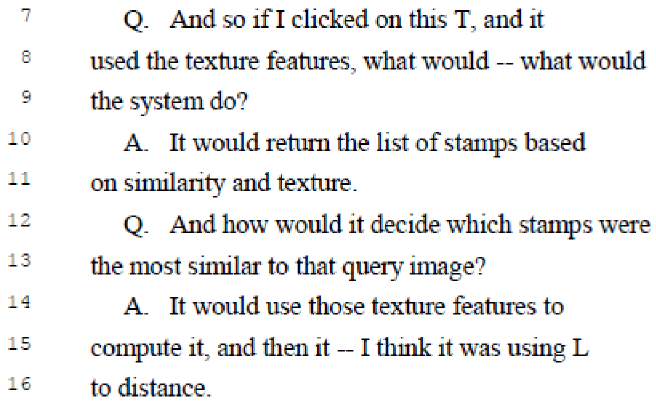 Q. And so if I clicked on this T, and it used the texture features, what would -- what would the system do? A. It would return the list of stamps based on similarity and texture. Q. And how would it decide which stamps were the most similar to that query image? A. It would use those texture features to compute it, and then it -- I think it was using L to distance. Flickner Depo at 75-76 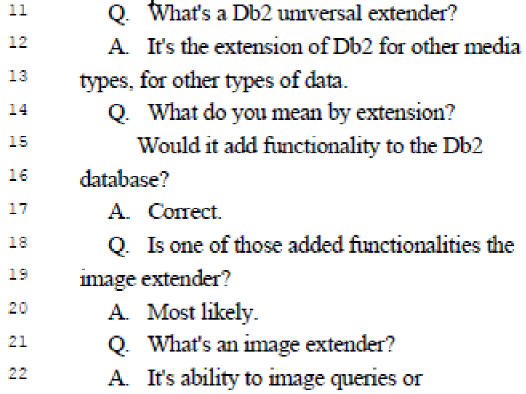 Q. What's a Db2 universal extender? A. It's the extension of Db2 for other media types, for other types of data. Q. What do you mean by extension? Would it add functionality to the Db2 database? A. Correct. Q. Is one of those added functionalities the image extender? A. Most likely. Q. What's an image extender? A. It's ability to image queries or

Element	'897 Claim Recitation	Exemplary Citations in Reference
		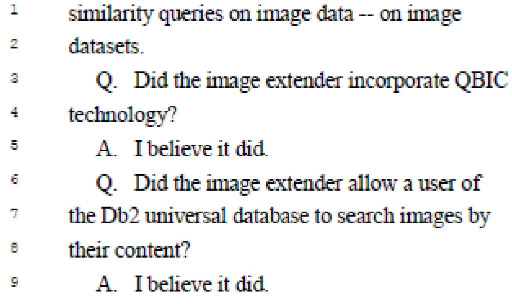 similarity queries on image data -- on image datasets. Q. Did the image extender incorporate QBIC technology? A. I believe it did. Q. Did the image extender allow a user of the Db2 universal database to search images by their content? A. I believe it did. Flickner Depo at 131-132 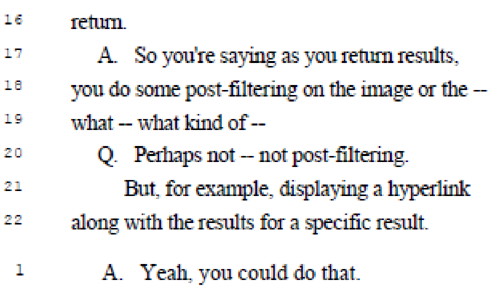 Q. An action that the QBIC system would automatically take based on the results it would return. A. So you're saying as you return results, you do some post-filtering on the image or the -- what -- what kind of -- Q. Perhaps not -- not post-filtering. But, for example, displaying a hyperlink along with the results for a specific result. A. Yeah, you could do that. Flickner 8/29/23 Depo at 33–37

Element	'897 Claim Recitation	Exemplary Citations in Reference
		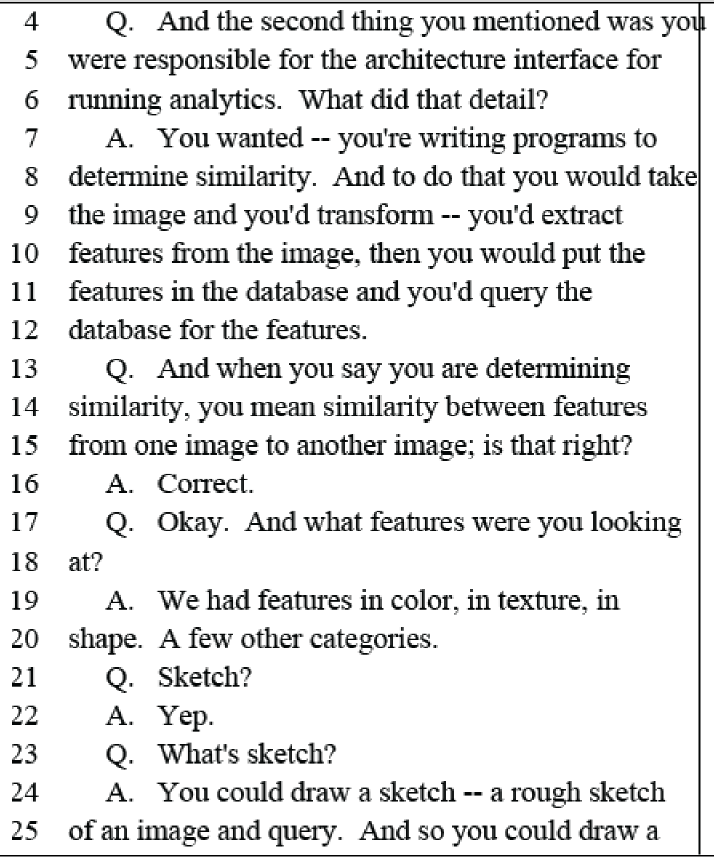 Q. And the second thing you mentioned was you were responsible for the architecture interface for running analytics. What did that detail? A. You wanted -- you're writing programs to determine similarity. And to do that you would take the image and you'd transform -- you'd extract features from the image, then you would put the features in the database and you'd query the database for the features. Q. And when you say you are determining similarity, you mean similarity between features from one image to another image; is that right? A. Correct. Q. Okay. And what features were you looking at? A. We had features in color, in texture, in shape. A few other categories. Q. Sketch? A. Yep. Q. What's sketch? A. You could draw a sketch -- a rough sketch of an image and query. And so you could draw a

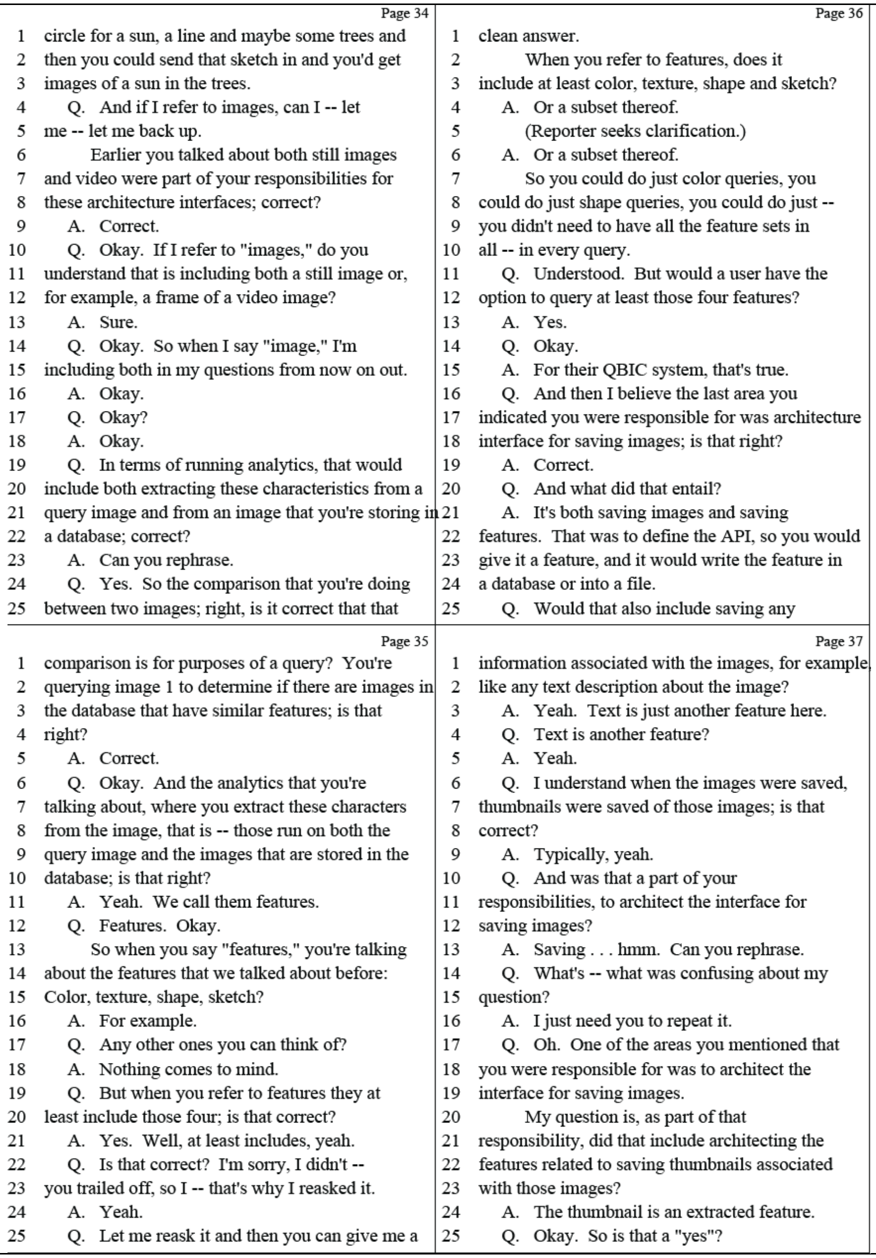

Page 34

circle for a sun, a line and maybe some trees and then you could send that sketch in and you'd get images of a sun in the trees.
Q. And if I refer to images, can I -- let me -- let me back up.
Earlier you talked about both still images and video were part of your responsibilities for these architecture interfaces; correct?
A. Correct.
Q. Okay. If I refer to "images," do you understand that is including both a still image or, for example, a frame of a video image?
A. Sure.
Q. Okay. So when I say "image," I'm including both in my questions from now on out.
A. Okay.
Q. Okay?
A. Okay.
Q. In terms of running analytics, that would include both extracting these characteristics from a query image and from an image that you're storing in a database; correct?
A. Can you rephrase.
Q. Yes. So the comparison that you're doing between two images; right, is it correct that that

Page 35

comparison is for purposes of a query? You're querying image 1 to determine if there are images in the database that have similar features; is that right?
A. Correct.
Q. Okay. And the analytics that you're talking about, where you extract these characters from the image, that is -- those run on both the query image and the images that are stored in the database; is that right?
A. Yeah. We call them features.
Q. Features. Okay.
So when you say "features," you're talking about the features that we talked about before: Color, texture, shape, sketch?
A. For example.
Q. Any other ones you can think of?
A. Nothing comes to mind.
Q. But when you refer to features they at least include those four; is that correct?
A. Yes. Well, at least includes, yeah.
Q. Is that correct? I'm sorry, I didn't -- you trailed off, so I -- that's why I reasked it.
A. Yeah.
Q. Let me reask it and then you can give me a

Page 36

clean answer.
When you refer to features, does it include at least color, texture, shape and sketch?
A. Or a subset thereof.
(Reporter seeks clarification.)
A. Or a subset thereof.
So you could do just color queries, you could do just shape queries, you could do just -- you didn't need to have all the feature sets in all -- in every query.
Q. Understood. But would a user have the option to query at least those four features?
A. Yes.
Q. Okay.
A. For their QBIC system, that's true.
Q. And then I believe the last area you indicated you were responsible for was architecture interface for saving images; is that right?
A. Correct.
Q. And what did that entail?
A. It's both saving images and saving features. That was to define the API, so you would give it a feature, and it would write the feature in a database or into a file.
Q. Would that also include saving any

Page 37

information associated with the images, for example, like any text description about the image?
A. Yeah. Text is just another feature here.
Q. Text is another feature?
A. Yeah.
Q. I understand when the images were saved, thumbnails were saved of those images; is that correct?
A. Typically, yeah.
Q. And was that a part of your responsibilities, to architect the interface for saving images?
A. Saving . . . hmm. Can you rephrase.
Q. What's -- what was confusing about my question?
A. I just need you to repeat it.
Q. Oh. One of the areas you mentioned that you were responsible for was to architect the interface for saving images.
My question is, as part of that responsibility, did that include architecting the features related to saving thumbnails associated with those images?
A. The thumbnail is an extracted feature.
Q. Okay. So is that a "yes"?

Element	'897 Claim Recitation	Exemplary Citations in Reference
		Id. at 52–53 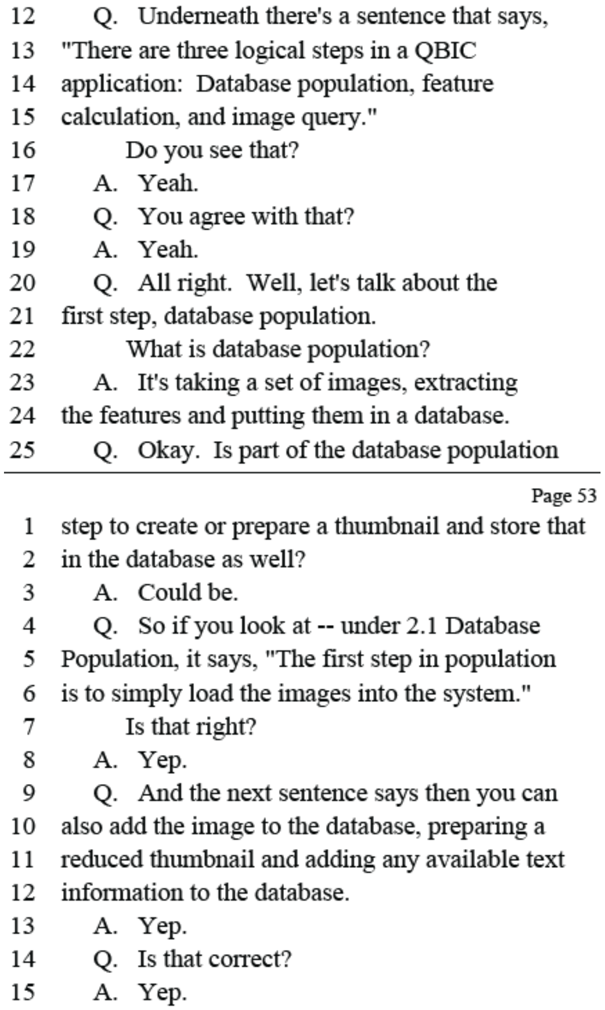 Q. Underneath there's a sentence that says, "There are three logical steps in a QBIC application: Database population, feature calculation, and image query." Do you see that? A. Yeah. Q. You agree with that? A. Yeah. Q. All right. Well, let's talk about the first step, database population. What is database population? A. It's taking a set of images, extracting the features and putting them in a database. Q. Okay. Is part of the database population step to create or prepare a thumbnail and store that in the database as well? Page 53 A. Could be. Q. So if you look at -- under 2.1 Database Population, it says, "The first step in population is to simply load the images into the system." Is that right? A. Yep. Q. And the next sentence says then you can also add the image to the database, preparing a reduced thumbnail and adding any available text information to the database. A. Yep. Q. Is that correct? A. Yep.

Element	'897 Claim Recitation	Exemplary Citations in Reference
		Id. at 77–78 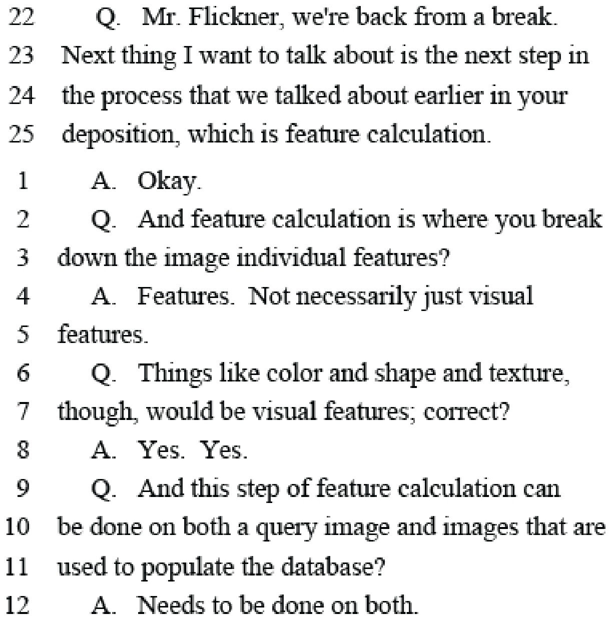 Q. Mr. Flickner, we're back from a break. Next thing I want to talk about is the next step in the process that we talked about earlier in your deposition, which is feature calculation. A. Okay. Q. And feature calculation is where you break down the image individual features? A. Features. Not necessarily just visual features. Q. Things like color and shape and texture, though, would be visual features; correct? A. Yes. Yes. Q. And this step of feature calculation can be done on both a query image and images that are used to populate the database? A. Needs to be done on both. *Id.* at 85–99 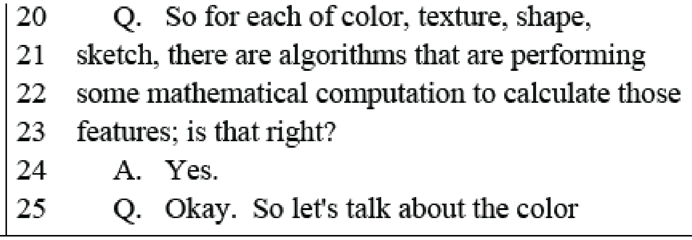 Q. So for each of color, texture, shape, sketch, there are algorithms that are performing some mathematical computation to calculate those features; is that right? A. Yes. Q. Okay. So let's talk about the color

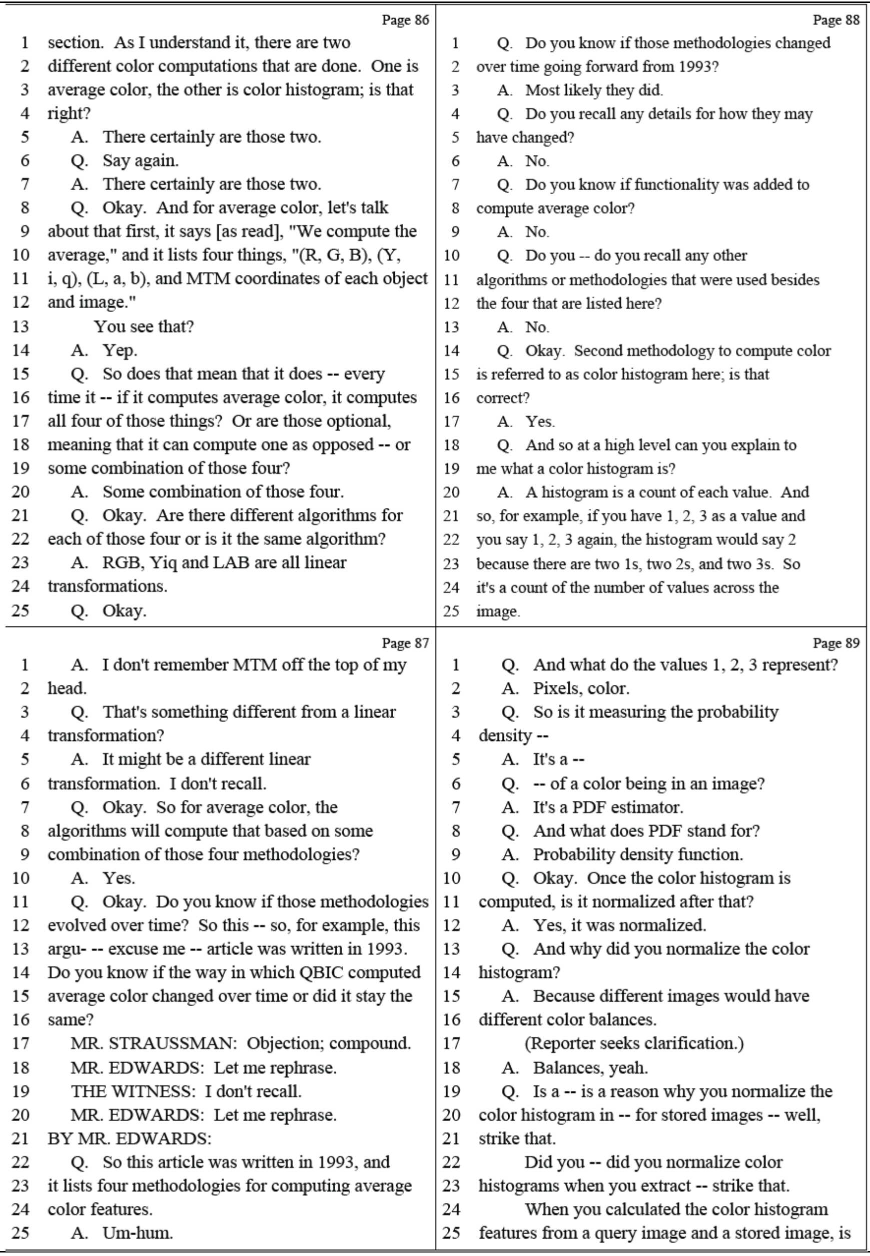

Page 86

section. As I understand it, there are two different color computations that are done. One is average color, the other is color histogram; is that right?
A. There certainly are those two.
Q. Say again.
A. There certainly are those two.
Q. Okay. And for average color, let's talk about that first, it says [as read], "We compute the average," and it lists four things, "(R, G, B), (Y, i, q), (L, a, b), and MTM coordinates of each object and image."
You see that?
A. Yep.
Q. So does that mean that it does -- every time it -- if it computes average color, it computes all four of those things? Or are those optional, meaning that it can compute one as opposed -- or some combination of those four?
A. Some combination of those four.
Q. Okay. Are there different algorithms for each of those four or is it the same algorithm?
A. RGB, Yiq and LAB are all linear transformations.
Q. Okay.

Page 87

A. I don't remember MTM off the top of my head.
Q. That's something different from a linear transformation?
A. It might be a different linear transformation. I don't recall.
Q. Okay. So for average color, the algorithms will compute that based on some combination of those four methodologies?
A. Yes.
Q. Okay. Do you know if those methodologies evolved over time? So this -- so, for example, this argu- -- excuse me -- article was written in 1993. Do you know if the way in which QBIC computed average color changed over time or did it stay the same?
MR. STRAUSSMAN: Objection; compound.
MR. EDWARDS: Let me rephrase.
THE WITNESS: I don't recall.
MR. EDWARDS: Let me rephrase.
BY MR. EDWARDS:
Q. So this article was written in 1993, and it lists four methodologies for computing average color features.
A. Um-hum.

Page 88

Q. Do you know if those methodologies changed over time going forward from 1993?
A. Most likely they did.
Q. Do you recall any details for how they may have changed?
A. No.
Q. Do you know if functionality was added to compute average color?
A. No.
Q. Do you -- do you recall any other algorithms or methodologies that were used besides the four that are listed here?
A. No.
Q. Okay. Second methodology to compute color is referred to as color histogram here; is that correct?
A. Yes.
Q. And so at a high level can you explain to me what a color histogram is?
A. A histogram is a count of each value. And so, for example, if you have 1, 2, 3 as a value and you say 1, 2, 3 again, the histogram would say 2 because there are two 1s, two 2s, and two 3s. So it's a count of the number of values across the image.

Page 89

Q. And what do the values 1, 2, 3 represent?
A. Pixels, color.
Q. So is it measuring the probability density --
A. It's a --
Q. -- of a color being in an image?
A. It's a PDF estimator.
Q. And what does PDF stand for?
A. Probability density function.
Q. Okay. Once the color histogram is computed, is it normalized after that?
A. Yes, it was normalized.
Q. And why did you normalize the color histogram?
A. Because different images would have different color balances.
(Reporter seeks clarification.)
A. Balances, yeah.
Q. Is a -- is a reason why you normalize the color histogram in -- for stored images -- well, strike that.
Did you -- did you normalize color histograms when you extract -- strike that.
When you calculated the color histogram features from a query image and a stored image, is

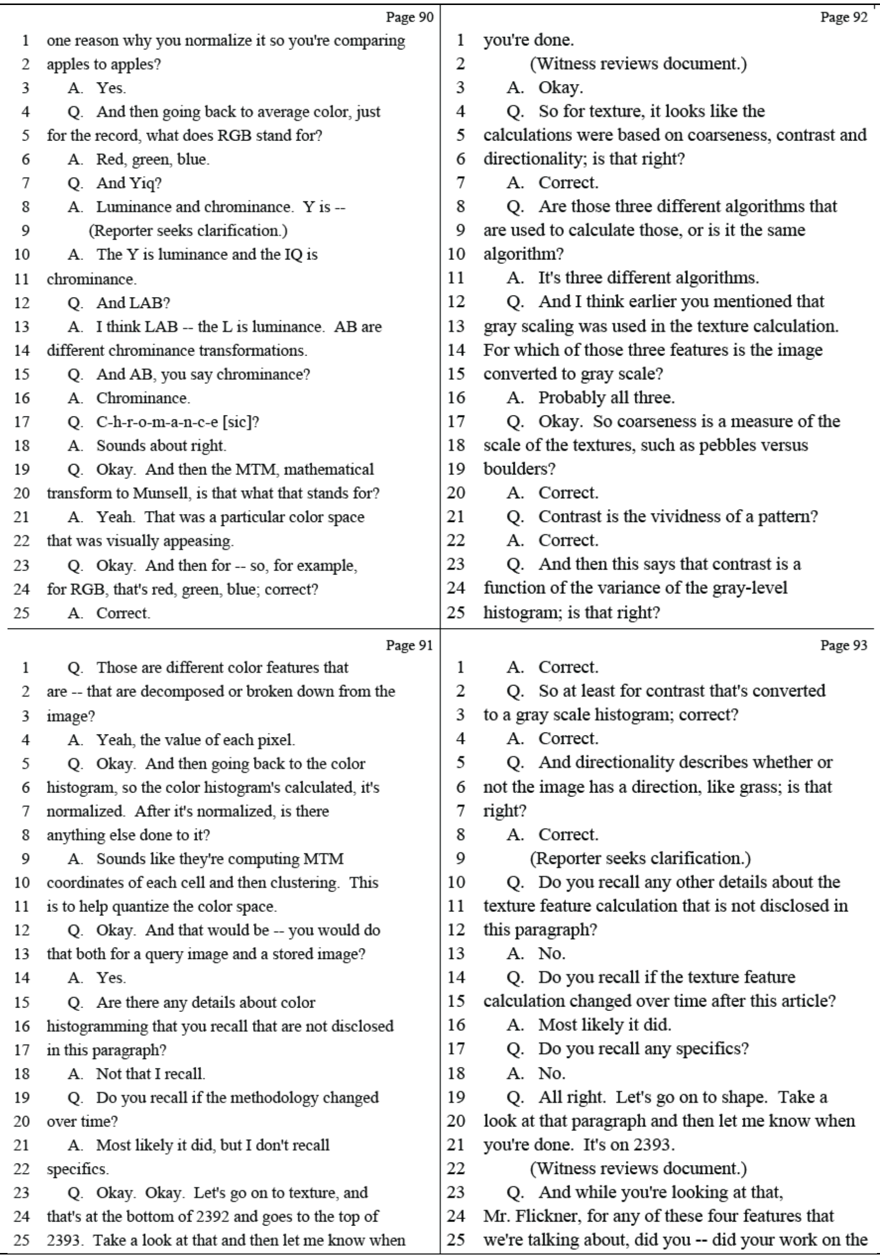

Page 90

one reason why you normalize it so you're comparing apples to apples?
A. Yes.
Q. And then going back to average color, just for the record, what does RGB stand for?
A. Red, green, blue.
Q. And Yiq?
A. Luminance and chrominance. Y is --
(Reporter seeks clarification.)
A. The Y is luminance and the IQ is chrominance.
Q. And LAB?
A. I think LAB -- the L is luminance. AB are different chrominance transformations.
Q. And AB, you say chrominance?
A. Chrominance.
Q. C-h-r-o-m-a-n-c-e [sic]?
A. Sounds about right.
Q. Okay. And then the MTM, mathematical transform to Munsell, is that what that stands for?
A. Yeah. That was a particular color space that was visually appeasing.
Q. Okay. And then for -- so, for example, for RGB, that's red, green, blue; correct?
A. Correct.

Page 91

Q. Those are different color features that are -- that are decomposed or broken down from the image?
A. Yeah, the value of each pixel.
Q. Okay. And then going back to the color histogram, so the color histogram's calculated, it's normalized. After it's normalized, is there anything else done to it?
A. Sounds like they're computing MTM coordinates of each cell and then clustering. This is to help quantize the color space.
Q. Okay. And that would be -- you would do that both for a query image and a stored image?
A. Yes.
Q. Are there any details about color histogramming that you recall that are not disclosed in this paragraph?
A. Not that I recall.
Q. Do you recall if the methodology changed over time?
A. Most likely it did, but I don't recall specifics.
Q. Okay. Okay. Let's go on to texture, and that's at the bottom of 2392 and goes to the top of 2393. Take a look at that and then let me know when

Page 92

you're done.
(Witness reviews document.)
A. Okay.
Q. So for texture, it looks like the calculations were based on coarseness, contrast and directionality; is that right?
A. Correct.
Q. Are those three different algorithms that are used to calculate those, or is it the same algorithm?
A. It's three different algorithms.
Q. And I think earlier you mentioned that gray scaling was used in the texture calculation. For which of those three features is the image converted to gray scale?
A. Probably all three.
Q. Okay. So coarseness is a measure of the scale of the textures, such as pebbles versus boulders?
A. Correct.
Q. Contrast is the vividness of a pattern?
A. Correct.
Q. And then this says that contrast is a function of the variance of the gray-level histogram; is that right?

Page 93

A. Correct.
Q. So at least for contrast that's converted to a gray scale histogram; correct?
A. Correct.
Q. And directionality describes whether or not the image has a direction, like grass; is that right?
A. Correct.
(Reporter seeks clarification.)
Q. Do you recall any other details about the texture feature calculation that is not disclosed in this paragraph?
A. No.
Q. Do you recall if the texture feature calculation changed over time after this article?
A. Most likely it did.
Q. Do you recall any specifics?
A. No.
Q. All right. Let's go on to shape. Take a look at that paragraph and then let me know when you're done. It's on 2393.
(Witness reviews document.)
Q. And while you're looking at that, Mr. Flickner, for any of these four features that we're talking about, did you -- did your work on the

Page 94

QBIC project -- were you more involved in one or more of these features than the other?
A. Yes.
Q. Which one?
A. Shape.
Q. Okay.
(Witness reviews document.)
A. So one thing I can state is originally I couldn't remember the name of the edge operator -- and the name is Canny, the Canny edge operator -- when we were talking about finding edges.
Q. Okay. So this -- the paragraph on IBM 2393 refreshed your memory that shape calculation used Canny edge detection?
A. No, that's sketch.
Q. Oh, I'm sorry.
A. I'm just stating that originally I was asked a question about standard edge operators. I couldn't remember the name. The name was Canny.
Q. Canny. And Canny is used in the sketch calculation?
A. Yeah, it's used in the sketch calculation.
Q. Understood. Okay. Let's go back to shape.
And for shape, is it a first step -- well,

Page 95

back up, strike that.
To calculate the shape feature in a query image or a stored database image, is the first step to convert that image into a binary image?
A. I don't think so.
Q. So if you look at the second paragraph -- and the reason why I ask is -- it's the second sentence in the second paragraph under "Shape features." It says [as read], "All shapes are assumed to be non-occluded planar shapes allowing for each shape to be represented as a binary image."
See that?
A. Yeah.
Q. And then the next sentence says [as read], "The area is computed as the number of pixels set in the binary image."
A. Right.
Q. So that's why I ask is it the first step in calculating shape to convert the image to binary?
A. Yes.
Q. And then once you convert it to binary image, then what do you do next in order to calculate the shape?
A. You're computing something called perround, which is perimeter squared over the area,

Page 96

of the object, of the blob.
Q. Can you say the name of the word again?
A. I think it's called perround.
Q. Per round?
A. Yeah.
Q. Those -- are those two words, p-e-r, r-o-u-n-d?
A. P-e-r -- one word.
Q. Okay. Perround, one word.
A. It's basically the perimeter squared over area.
(Reporter seeks clarification.)
A. Perimeter squared over area.
(Reporter seeks clarification.)
A. Length of the -- over area.
Q. Okay. So let's step back a little bit. So it says -- the second paragraph says [as read], "Our shape features are based on a combination of heuristic" -- h-e-u-r-i-s-t-i-c -- "shape features area, circularity, eccentricity, major access orientation and a set of algebraic moment invariants," i-n-v-a-r-i-a-n-t-s.
Did I read that correctly?
A. Um-hum.
Q. Okay. For circularity, that is computed

Page 97

as perimeter squared over area; correct?
A. Right.
Q. All right. Area is the number of pixels in the binary image?
A. In the binary blob.
Q. In the binary what?
A. Blob.
Q. Blob. And when you say "blob," what does that refer to?
A. The outline of the object that you're doing a shape query on.
Q. And so are each of these, area, circularity, eccentricity, major axis orientation and algebraic moment invariants, are those each different algorithms or is it all the same algorithm?
A. It's different algorithms.
Q. Okay. And since you worked on the shape feature the most, do you recall any other details that are not disclosed in this paragraph?
A. Well, this evolved and we ended up doing an implementation of something that's not in here.
Q. Okay. And what was that implementation?
A. Something called turning angle per round, turning angle, which is -- it looks at the tangent

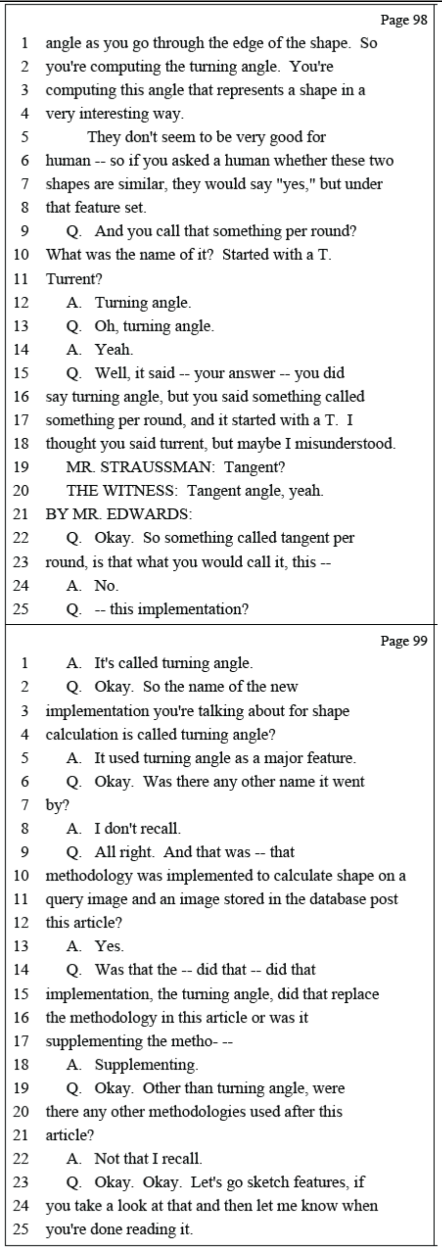

Page 98

angle as you go through the edge of the shape. So you're computing the turning angle. You're computing this angle that represents a shape in a very interesting way.

They don't seem to be very good for human -- so if you asked a human whether these two shapes are similar, they would say "yes," but under that feature set.

Q. And you call that something per round? What was the name of it? Started with a T. Turrent?

A. Turning angle.

Q. Oh, turning angle.

A. Yeah.

Q. Well, it said -- your answer -- you did say turning angle, but you said something called something per round, and it started with a T. I thought you said turrent, but maybe I misunderstood.

MR. STRAUSSMAN: Tangent?

THE WITNESS: Tangent angle, yeah.

BY MR. EDWARDS:

Q. Okay. So something called tangent per round, is that what you would call it, this --

A. No.

Q. -- this implementation?

Page 99

A. It's called turning angle.

Q. Okay. So the name of the new implementation you're talking about for shape calculation is called turning angle?

A. It used turning angle as a major feature.

Q. Okay. Was there any other name it went by?

A. I don't recall.

Q. All right. And that was -- that methodology was implemented to calculate shape on a query image and an image stored in the database post this article?

A. Yes.

Q. Was that the -- did that -- did that implementation, the turning angle, did that replace the methodology in this article or was it supplementing the metho- --

A. Supplementing.

Q. Okay. Other than turning angle, were there any other methodologies used after this article?

A. Not that I recall.

Q. Okay. Okay. Let's go sketch features, if you take a look at that and then let me know when you're done reading it.

Element	'897 Claim Recitation	Exemplary Citations in Reference
		Id. at 119–128 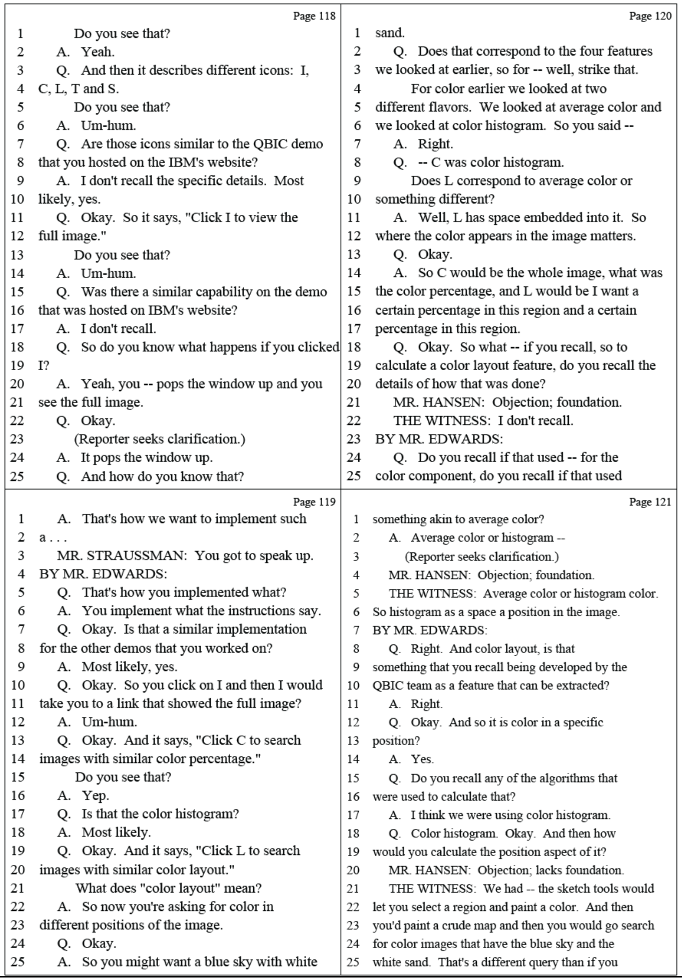

Page 118

Do you see that?
A. Yeah.
Q. And then it describes different icons: I, C, L, T and S.
Do you see that?
A. Um-hum.
Q. Are those icons similar to the QBIC demo that you hosted on the IBM's website?
A. I don't recall the specific details. Most likely, yes.
Q. Okay. So it says, "Click I to view the full image."
Do you see that?
A. Um-hum.
Q. Was there a similar capability on the demo that was hosted on IBM's website?
A. I don't recall.
Q. So do you know what happens if you clicked I?
A. Yeah, you -- pops the window up and you see the full image.
Q. Okay.
(Reporter seeks clarification.)
A. It pops the window up.
Q. And how do you know that?

Page 119

A. That's how we want to implement such a . . .
MR. STRAUSSMAN: You got to speak up.
BY MR. EDWARDS:
Q. That's how you implemented what?
A. You implement what the instructions say.
Q. Okay. Is that a similar implementation for the other demos that you worked on?
A. Most likely, yes.
Q. Okay. So you click on I and then I would take you to a link that showed the full image?
A. Um-hum.
Q. Okay. And it says, "Click C to search images with similar color percentage."
Do you see that?
A. Yep.
Q. Is that the color histogram?
A. Most likely.
Q. Okay. And it says, "Click L to search images with similar color layout."
What does "color layout" mean?
A. So now you're asking for color in different positions of the image.
Q. Okay.
A. So you might want a blue sky with white

Page 120

sand.
Q. Does that correspond to the four features we looked at earlier, so for -- well, strike that.
For color earlier we looked at two different flavors. We looked at average color and we looked at color histogram. So you said --
A. Right.
Q. -- C was color histogram.
Does L correspond to average color or something different?
A. Well, L has space embedded into it. So where the color appears in the image matters.
Q. Okay.
A. So C would be the whole image, what was the color percentage, and L would be I want a certain percentage in this region and a certain percentage in this region.
Q. Okay. So what -- if you recall, so to calculate a color layout feature, do you recall the details of how that was done?
MR. HANSEN: Objection; foundation.
THE WITNESS: I don't recall.
BY MR. EDWARDS:
Q. Do you recall if that used -- for the color component, do you recall if that used

Page 121

something akin to average color?
A. Average color or histogram --
(Reporter seeks clarification.)
MR. HANSEN: Objection; foundation.
THE WITNESS: Average color or histogram color. So histogram as a space a position in the image.
BY MR. EDWARDS:
Q. Right. And color layout, is that something that you recall being developed by the QBIC team as a feature that can be extracted?
A. Right.
Q. Okay. And so it is color in a specific position?
A. Yes.
Q. Do you recall any of the algorithms that were used to calculate that?
A. I think we were using color histogram.
Q. Color histogram. Okay. And then how would you calculate the position aspect of it?
MR. HANSEN: Objection; lacks foundation.
THE WITNESS: We had -- the sketch tools would let you select a region and paint a color. And then you'd paint a crude map and then you would go search for color images that have the blue sky and the white sand. That's a different query than if you

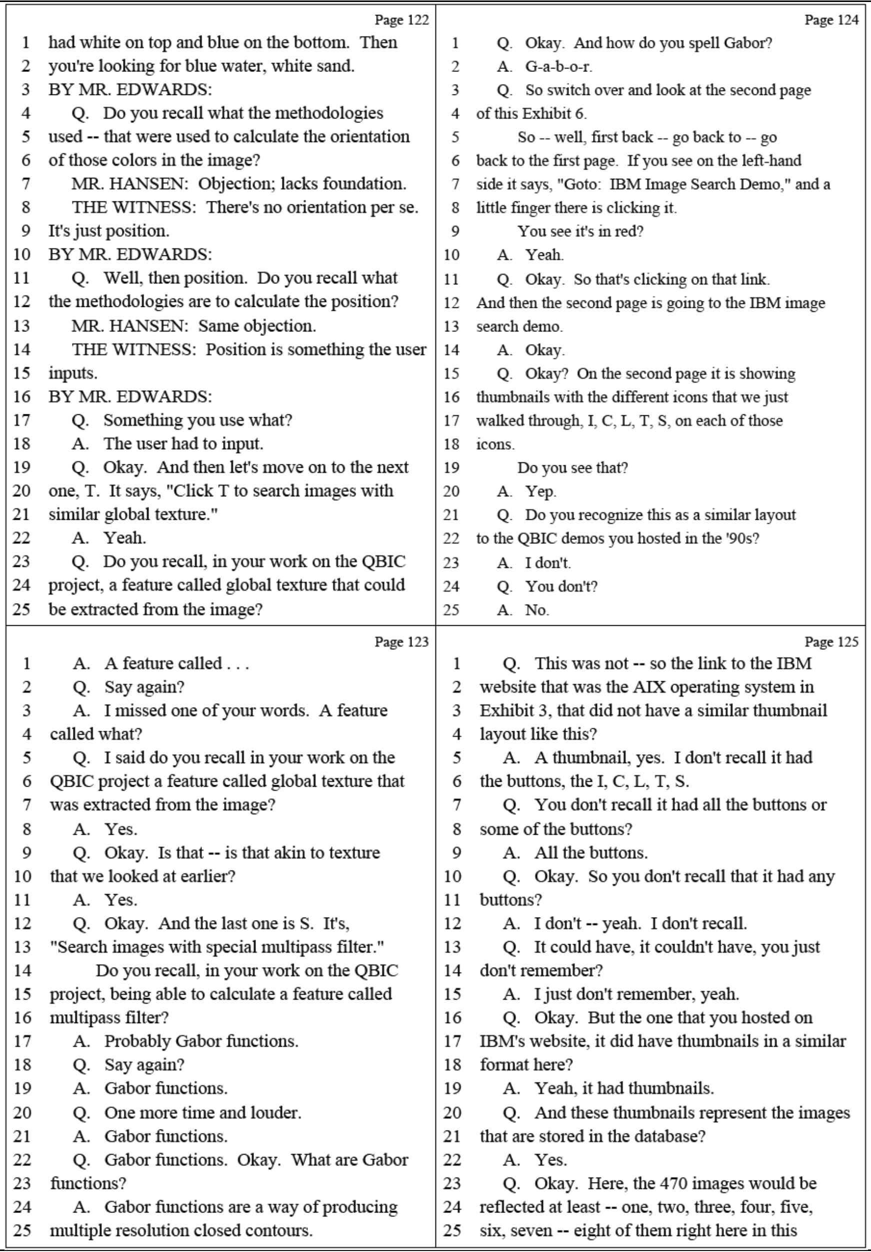

Page 122

had white on top and blue on the bottom. Then you're looking for blue water, white sand.
BY MR. EDWARDS:
Q. Do you recall what the methodologies used -- that were used to calculate the orientation of those colors in the image?
MR. HANSEN: Objection; lacks foundation.
THE WITNESS: There's no orientation per se. It's just position.
BY MR. EDWARDS:
Q. Well, then position. Do you recall what the methodologies are to calculate the position?
MR. HANSEN: Same objection.
THE WITNESS: Position is something the user inputs.
BY MR. EDWARDS:
Q. Something you use what?
A. The user had to input.
Q. Okay. And then let's move on to the next one, T. It says, "Click T to search images with similar global texture."
A. Yeah.
Q. Do you recall, in your work on the QBIC project, a feature called global texture that could be extracted from the image?

Page 123

A. A feature called . . .
Q. Say again?
A. I missed one of your words. A feature called what?
Q. I said do you recall in your work on the QBIC project a feature called global texture that was extracted from the image?
A. Yes.
Q. Okay. Is that -- is that akin to texture that we looked at earlier?
A. Yes.
Q. Okay. And the last one is S. It's, "Search images with special multipass filter."
Do you recall, in your work on the QBIC project, being able to calculate a feature called multipass filter?
A. Probably Gabor functions.
Q. Say again?
A. Gabor functions.
Q. One more time and louder.
A. Gabor functions.
Q. Gabor functions. Okay. What are Gabor functions?
A. Gabor functions are a way of producing multiple resolution closed contours.

Page 124

Q. Okay. And how do you spell Gabor?
A. G-a-b-o-r.
Q. So switch over and look at the second page of this Exhibit 6.
So -- well, first back -- go back to -- go back to the first page. If you see on the left-hand side it says, "Goto: IBM Image Search Demo," and a little finger there is clicking it.
You see it's in red?
A. Yeah.
Q. Okay. So that's clicking on that link. And then the second page is going to the IBM image search demo.
A. Okay.
Q. Okay? On the second page it is showing thumbnails with the different icons that we just walked through, I, C, L, T, S, on each of those icons.
Do you see that?
A. Yep.
Q. Do you recognize this as a similar layout to the QBIC demos you hosted in the '90s?
A. I don't.
Q. You don't?
A. No.

Page 125

Q. This was not -- so the link to the IBM website that was the AIX operating system in Exhibit 3, that did not have a similar thumbnail layout like this?
A. A thumbnail, yes. I don't recall it had the buttons, the I, C, L, T, S.
Q. You don't recall it had all the buttons or some of the buttons?
A. All the buttons.
Q. Okay. So you don't recall that it had any buttons?
A. I don't -- yeah. I don't recall.
Q. It could have, it couldn't have, you just don't remember?
A. I just don't remember, yeah.
Q. Okay. But the one that you hosted on IBM's website, it did have thumbnails in a similar format here?
A. Yeah, it had thumbnails.
Q. And these thumbnails represent the images that are stored in the database?
A. Yes.
Q. Okay. Here, the 470 images would be reflected at least -- one, two, three, four, five, six, seven -- eight of them right here in this

Element	'897 Claim Recitation	Exemplary Citations in Reference
		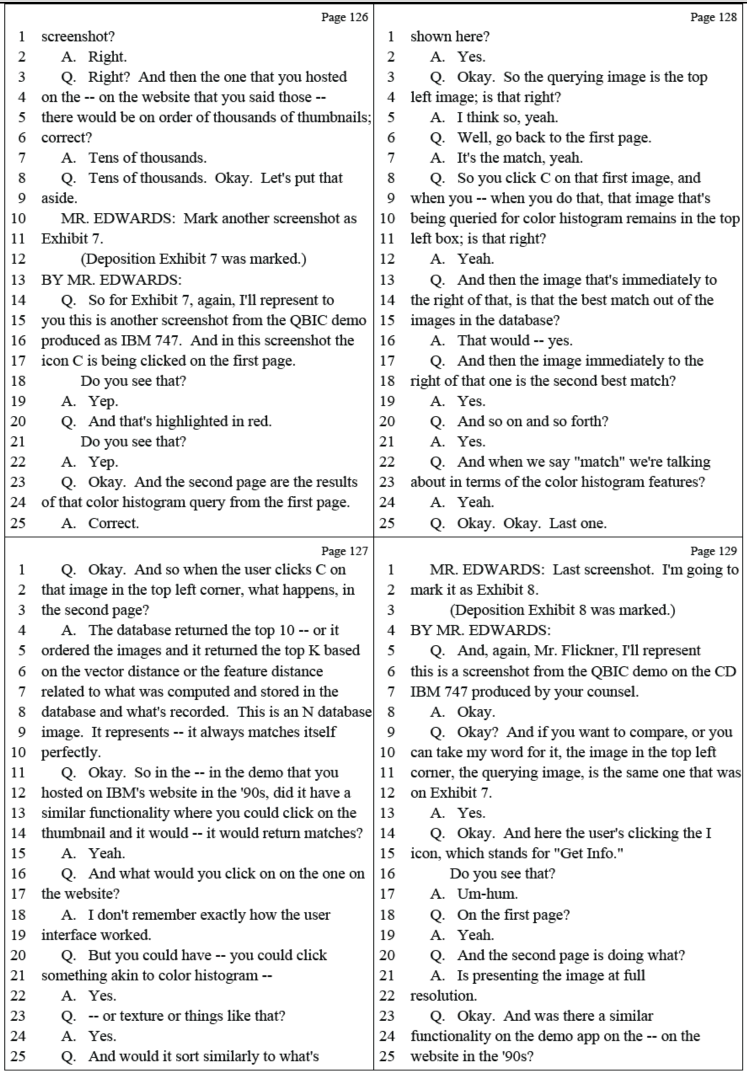Page 126 screenshot? A. Right. Q. Right? And then the one that you hosted on the -- on the website that you said those -- there would be on order of thousands of thumbnails; correct? A. Tens of thousands. Q. Tens of thousands. Okay. Let's put that aside. MR. EDWARDS: Mark another screenshot as Exhibit 7. (Deposition Exhibit 7 was marked.) BY MR. EDWARDS: Q. So for Exhibit 7, again, I'll represent to you this is another screenshot from the QBIC demo produced as IBM 747. And in this screenshot the icon C is being clicked on the first page. Do you see that? A. Yep. Q. And that's highlighted in red. Do you see that? A. Yep. Q. Okay. And the second page are the results of that color histogram query from the first page. A. Correct. Page 127 Q. Okay. And so when the user clicks C on that image in the top left corner, what happens, in the second page? A. The database returned the top 10 -- or it ordered the images and it returned the top K based on the vector distance or the feature distance related to what was computed and stored in the database and what's recorded. This is an N database image. It represents -- it always matches itself perfectly. Q. Okay. So in the -- in the demo that you hosted on IBM's website in the '90s, did it have a similar functionality where you could click on the thumbnail and it would -- it would return matches? A. Yeah. Q. And what would you click on on the one on the website? A. I don't remember exactly how the user interface worked. Q. But you could have -- you could click something akin to color histogram -- A. Yes. Q. -- or texture or things like that? A. Yes. Q. And would it sort similarly to what's Page 128 shown here? A. Yes. Q. Okay. So the querying image is the top left image; is that right? A. I think so, yeah. Q. Well, go back to the first page. A. It's the match, yeah. Q. So you click C on that first image, and when you -- when you do that, that image that's being queried for color histogram remains in the top left box; is that right? A. Yeah. Q. And then the image that's immediately to the right of that, is that the best match out of the images in the database? A. That would -- yes. Q. And then the image immediately to the right of that one is the second best match? A. Yes. Q. And so on and so forth? A. Yes. Q. And when we say "match" we're talking about in terms of the color histogram features? A. Yeah. Q. Okay. Okay. Last one. Page 129 MR. EDWARDS: Last screenshot. I'm going to mark it as Exhibit 8. (Deposition Exhibit 8 was marked.) BY MR. EDWARDS: Q. And, again, Mr. Flickner, I'll represent this is a screenshot from the QBIC demo on the CD IBM 747 produced by your counsel. A. Okay. Q. Okay? And if you want to compare, or you can take my word for it, the image in the top left corner, the querying image, is the same one that was on Exhibit 7. A. Yes. Q. Okay. And here the user's clicking the I icon, which stands for "Get Info." Do you see that? A. Um-hum. Q. On the first page? A. Yeah. Q. And the second page is doing what? A. Is presenting the image at full resolution. Q. Okay. And was there a similar functionality on the demo app on the -- on the website in the '90s? *Id*. at 135 – 136

Element	’897 Claim Recitation	Exemplary Citations in Reference
		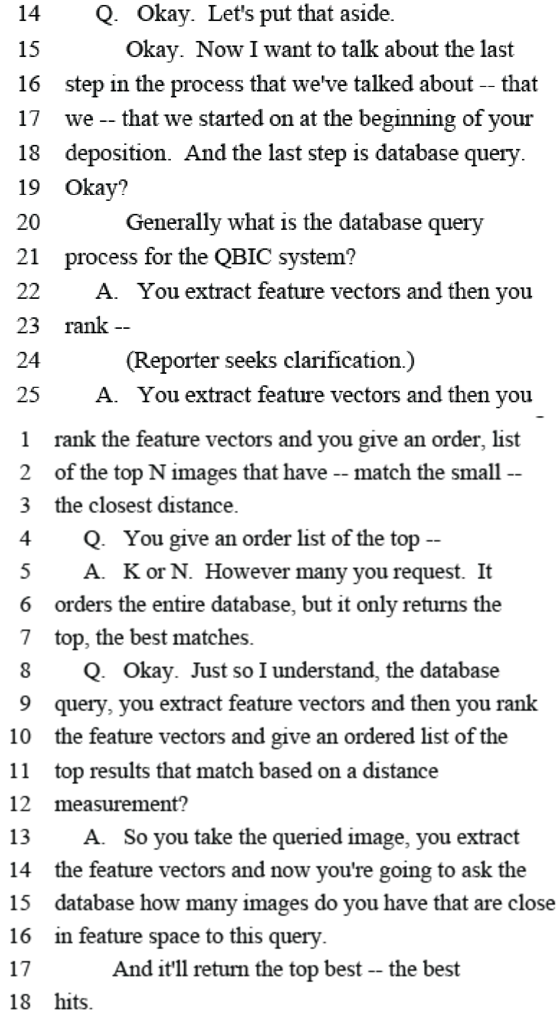 Q. Okay. Let's put that aside. Okay. Now I want to talk about the last step in the process that we've talked about -- that we -- that we started on at the beginning of your deposition. And the last step is database query. Okay? Generally what is the database query process for the QBIC system? A. You extract feature vectors and then you rank -- (Reporter seeks clarification.) A. You extract feature vectors and then you rank the feature vectors and you give an order, list of the top N images that have -- match the small -- the closest distance. Q. You give an order list of the top -- A. K or N. However many you request. It orders the entire database, but it only returns the top, the best matches. Q. Okay. Just so I understand, the database query, you extract feature vectors and then you rank the feature vectors and give an ordered list of the top results that match based on a distance measurement? A. So you take the queried image, you extract the feature vectors and now you're going to ask the database how many images do you have that are close in feature space to this query. And it'll return the top best -- the best hits. *Id.* at 139–165

Element	'897 Claim Recitation	Exemplary Citations in Reference
		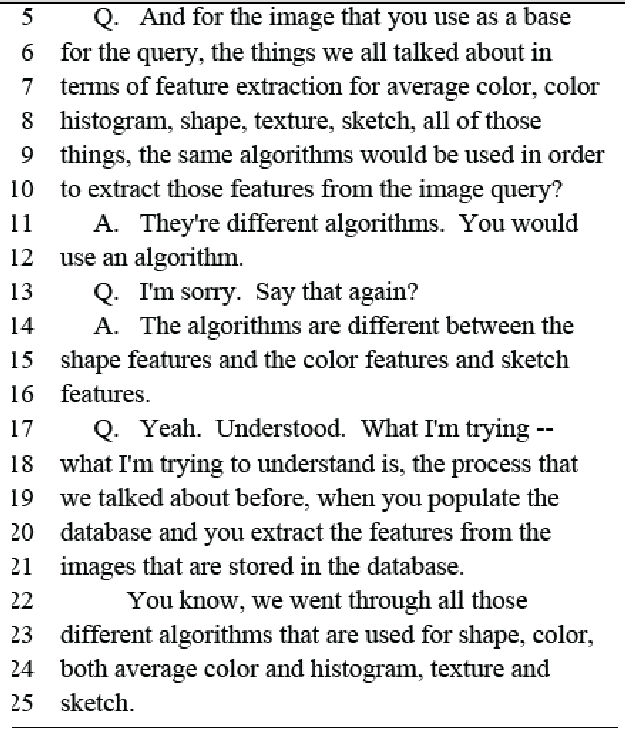 Q. And for the image that you use as a base for the query, the things we all talked about in terms of feature extraction for average color, color histogram, shape, texture, sketch, all of those things, the same algorithms would be used in order to extract those features from the image query? A. They're different algorithms. You would use an algorithm. Q. I'm sorry. Say that again? A. The algorithms are different between the shape features and the color features and sketch features. Q. Yeah. Understood. What I'm trying -- what I'm trying to understand is, the process that we talked about before, when you populate the database and you extract the features from the images that are stored in the database. You know, we went through all those different algorithms that are used for shape, color, both average color and histogram, texture and sketch.

Element	'897 Claim Recitation	Exemplary Citations in Reference
		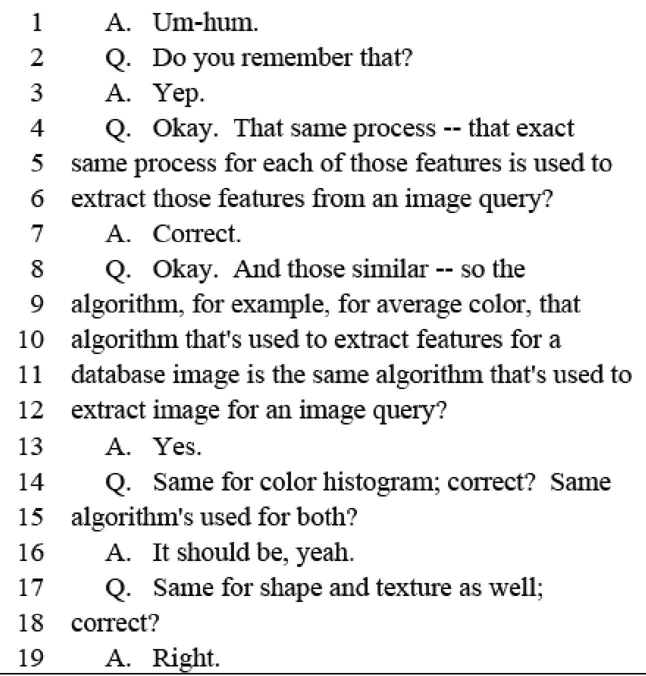 A. Um-hum. Q. Do you remember that? A. Yep. Q. Okay. That same process -- that exact same process for each of those features is used to extract those features from an image query? A. Correct. Q. Okay. And those similar -- so the algorithm, for example, for average color, that algorithm that's used to extract features for a database image is the same algorithm that's used to extract image for an image query? A. Yes. Q. Same for color histogram; correct? Same algorithm's used for both? A. It should be, yeah. Q. Same for shape and texture as well; correct? A. Right.

			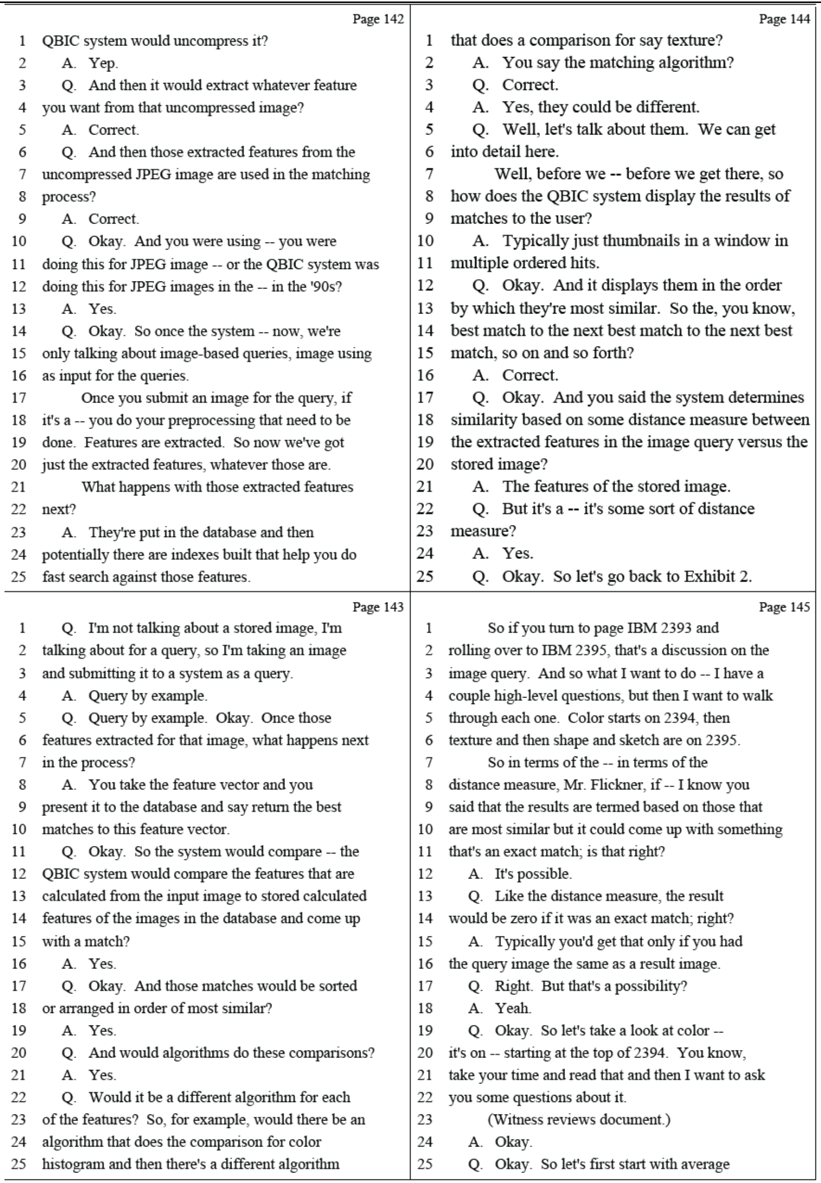

Page 142

QBIC system would uncompress it?
A. Yep.
Q. And then it would extract whatever feature you want from that uncompressed image?
A. Correct.
Q. And then those extracted features from the uncompressed JPEG image are used in the matching process?
A. Correct.
Q. Okay. And you were using -- you were doing this for JPEG image -- or the QBIC system was doing this for JPEG images in the -- in the '90s?
A. Yes.
Q. Okay. So once the system -- now, we're only talking about image-based queries, image using as input for the queries.
Once you submit an image for the query, if it's a -- you do your preprocessing that need to be done. Features are extracted. So now we've got just the extracted features, whatever those are.
What happens with those extracted features next?
A. They're put in the database and then potentially there are indexes built that help you do fast search against those features.

Page 143

Q. I'm not talking about a stored image, I'm talking about for a query, so I'm taking an image and submitting it to a system as a query.
A. Query by example.
Q. Query by example. Okay. Once those features extracted for that image, what happens next in the process?
A. You take the feature vector and you present it to the database and say return the best matches to this feature vector.
Q. Okay. So the system would compare -- the QBIC system would compare the features that are calculated from the input image to stored calculated features of the images in the database and come up with a match?
A. Yes.
Q. Okay. And those matches would be sorted or arranged in order of most similar?
A. Yes.
Q. And would algorithms do these comparisons?
A. Yes.
Q. Would it be a different algorithm for each of the features? So, for example, would there be an algorithm that does the comparison for color histogram and then there's a different algorithm

Page 144

that does a comparison for say texture?
A. You say the matching algorithm?
Q. Correct.
A. Yes, they could be different.
Q. Well, let's talk about them. We can get into detail here.
Well, before we -- before we get there, so how does the QBIC system display the results of matches to the user?
A. Typically just thumbnails in a window in multiple ordered hits.
Q. Okay. And it displays them in the order by which they're most similar. So the, you know, best match to the next best match to the next best match, so on and so forth?
A. Correct.
Q. Okay. And you said the system determines similarity based on some distance measure between the extracted features in the image query versus the stored image?
A. The features of the stored image.
Q. But it's a -- it's some sort of distance measure?
A. Yes.
Q. Okay. So let's go back to Exhibit 2.

Page 145

So if you turn to page IBM 2393 and rolling over to IBM 2395, that's a discussion on the image query. And so what I want to do -- I have a couple high-level questions, but then I want to walk through each one. Color starts on 2394, then texture and then shape and sketch are on 2395.
So in terms of the -- in terms of the distance measure, Mr. Flickner, if -- I know you said that the results are termed based on those that are most similar but it could come up with something that's an exact match; is that right?
A. It's possible.
Q. Like the distance measure, the result would be zero if it was an exact match; right?
A. Typically you'd get that only if you had the query image the same as a result image.
Q. Right. But that's a possibility?
A. Yeah.
Q. Okay. So let's take a look at color -- it's on -- starting at the top of 2394. You know, take your time and read that and then I want to ask you some questions about it.
(Witness reviews document.)
A. Okay.
Q. Okay. So let's first start with average

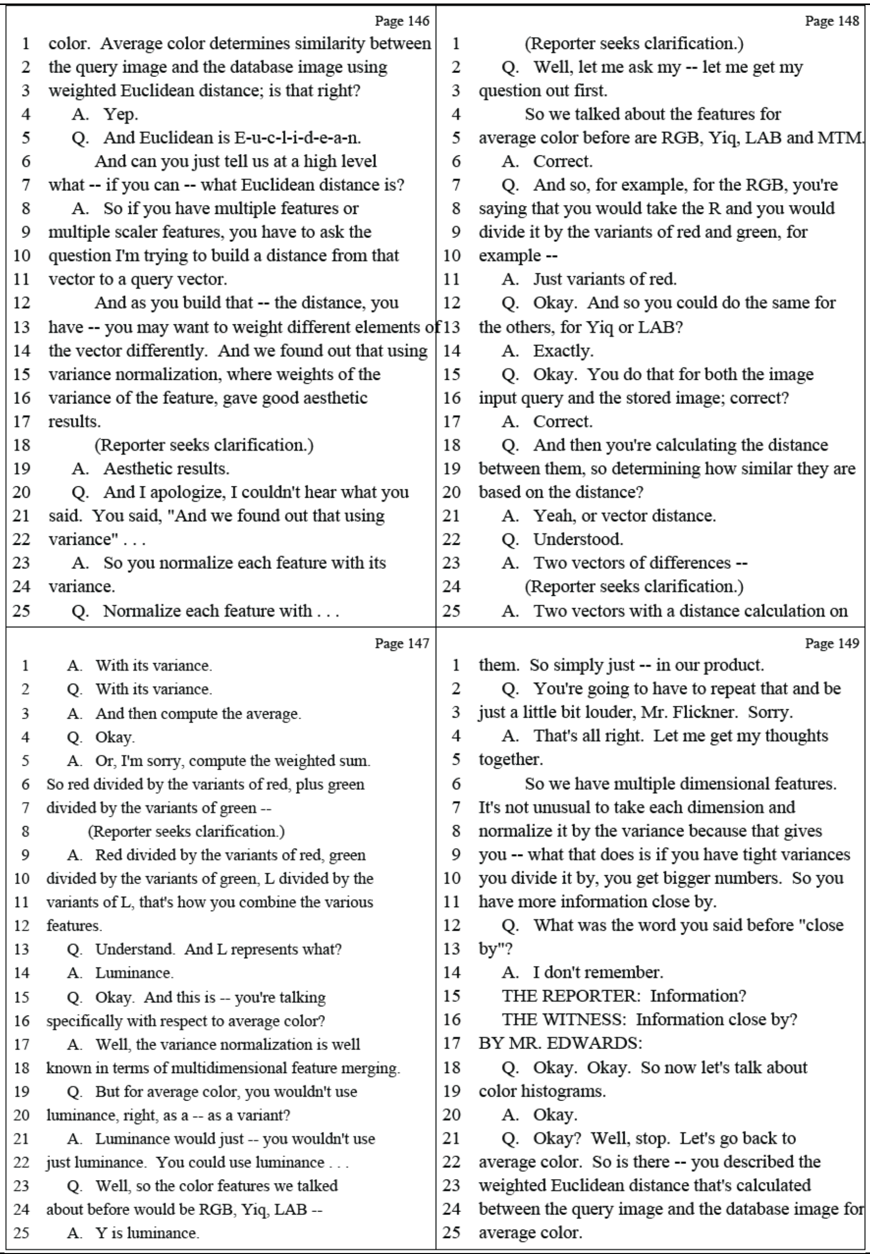

Page 146

color. Average color determines similarity between the query image and the database image using weighted Euclidean distance; is that right?
A. Yep.
Q. And Euclidean is E-u-c-l-i-d-e-a-n.
And can you just tell us at a high level what -- if you can -- what Euclidean distance is?
A. So if you have multiple features or multiple scaler features, you have to ask the question I'm trying to build a distance from that vector to a query vector.
And as you build that -- the distance, you have -- you may want to weight different elements of the vector differently. And we found out that using variance normalization, where weights of the variance of the feature, gave good aesthetic results.
(Reporter seeks clarification.)
A. Aesthetic results.
Q. And I apologize, I couldn't hear what you said. You said, "And we found out that using variance" . . .
A. So you normalize each feature with its variance.
Q. Normalize each feature with . . .

Page 147

A. With its variance.
Q. With its variance.
A. And then compute the average.
Q. Okay.
A. Or, I'm sorry, compute the weighted sum. So red divided by the variants of red, plus green divided by the variants of green --
(Reporter seeks clarification.)
A. Red divided by the variants of red, green divided by the variants of green, L divided by the variants of L, that's how you combine the various features.
Q. Understand. And L represents what?
A. Luminance.
Q. Okay. And this is -- you're talking specifically with respect to average color?
A. Well, the variance normalization is well known in terms of multidimensional feature merging.
Q. But for average color, you wouldn't use luminance, right, as a -- as a variant?
A. Luminance would just -- you wouldn't use just luminance. You could use luminance . . .
Q. Well, so the color features we talked about before would be RGB, Yiq, LAB --
A. Y is luminance.

Page 148

(Reporter seeks clarification.)
Q. Well, let me ask my -- let me get my question out first.
So we talked about the features for average color before are RGB, Yiq, LAB and MTM.
A. Correct.
Q. And so, for example, for the RGB, you're saying that you would take the R and you would divide it by the variants of red and green, for example --
A. Just variants of red.
Q. Okay. And so you could do the same for the others, for Yiq or LAB?
A. Exactly.
Q. Okay. You do that for both the image input query and the stored image; correct?
A. Correct.
Q. And then you're calculating the distance between them, so determining how similar they are based on the distance?
A. Yeah, or vector distance.
Q. Understood.
A. Two vectors of differences --
(Reporter seeks clarification.)
A. Two vectors with a distance calculation on

Page 149

them. So simply just -- in our product.
Q. You're going to have to repeat that and be just a little bit louder, Mr. Flickner. Sorry.
A. That's all right. Let me get my thoughts together.
So we have multiple dimensional features. It's not unusual to take each dimension and normalize it by the variance because that gives you -- what that does is if you have tight variances you divide it by, you get bigger numbers. So you have more information close by.
Q. What was the word you said before "close by"?
A. I don't remember.
THE REPORTER: Information?
THE WITNESS: Information close by?
BY MR. EDWARDS:
Q. Okay. Okay. So now let's talk about color histograms.
A. Okay.
Q. Okay? Well, stop. Let's go back to average color. So is there -- you described the weighted Euclidean distance that's calculated between the query image and the database image for average color.

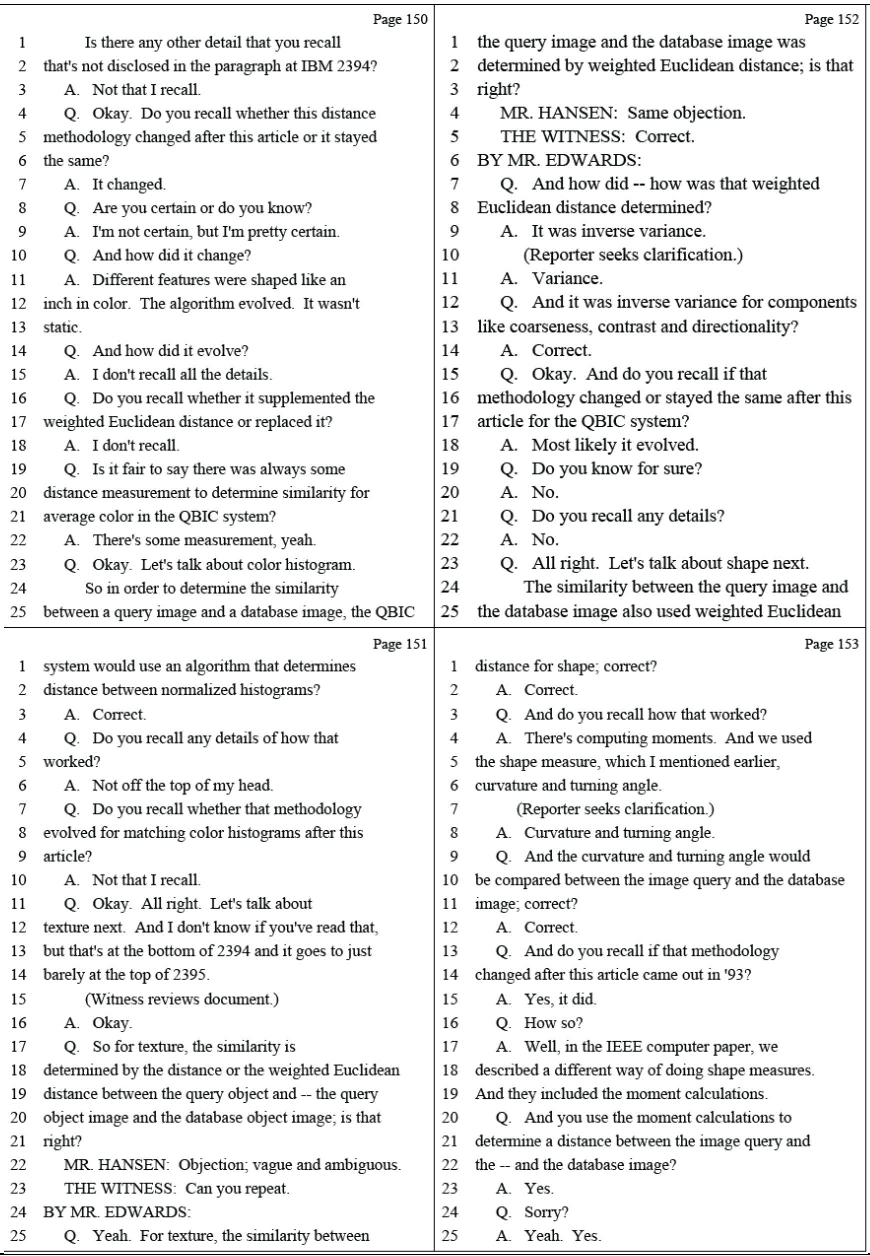

Page 150

Is there any other detail that you recall that's not disclosed in the paragraph at IBM 2394?
A. Not that I recall.
Q. Okay. Do you recall whether this distance methodology changed after this article or it stayed the same?
A. It changed.
Q. Are you certain or do you know?
A. I'm not certain, but I'm pretty certain.
Q. And how did it change?
A. Different features were shaped like an inch in color. The algorithm evolved. It wasn't static.
Q. And how did it evolve?
A. I don't recall all the details.
Q. Do you recall whether it supplemented the weighted Euclidean distance or replaced it?
A. I don't recall.
Q. Is it fair to say there was always some distance measurement to determine similarity for average color in the QBIC system?
A. There's some measurement, yeah.
Q. Okay. Let's talk about color histogram.
So in order to determine the similarity between a query image and a database image, the QBIC

Page 151

system would use an algorithm that determines distance between normalized histograms?
A. Correct.
Q. Do you recall any details of how that worked?
A. Not off the top of my head.
Q. Do you recall whether that methodology evolved for matching color histograms after this article?
A. Not that I recall.
Q. Okay. All right. Let's talk about texture next. And I don't know if you've read that, but that's at the bottom of 2394 and it goes to just barely at the top of 2395.
(Witness reviews document.)
A. Okay.
Q. So for texture, the similarity is determined by the distance or the weighted Euclidean distance between the query object and -- the query object image and the database object image; is that right?
MR. HANSEN: Objection; vague and ambiguous.
THE WITNESS: Can you repeat.
BY MR. EDWARDS:
Q. Yeah. For texture, the similarity between

Page 152

the query image and the database image was determined by weighted Euclidean distance; is that right?
MR. HANSEN: Same objection.
THE WITNESS: Correct.
BY MR. EDWARDS:
Q. And how did -- how was that weighted Euclidean distance determined?
A. It was inverse variance.
(Reporter seeks clarification.)
A. Variance.
Q. And it was inverse variance for components like coarseness, contrast and directionality?
A. Correct.
Q. Okay. And do you recall if that methodology changed or stayed the same after this article for the QBIC system?
A. Most likely it evolved.
Q. Do you know for sure?
A. No.
Q. Do you recall any details?
A. No.
Q. All right. Let's talk about shape next.
The similarity between the query image and the database image also used weighted Euclidean

Page 153

distance for shape; correct?
A. Correct.
Q. And do you recall how that worked?
A. There's computing moments. And we used the shape measure, which I mentioned earlier, curvature and turning angle.
(Reporter seeks clarification.)
A. Curvature and turning angle.
Q. And the curvature and turning angle would be compared between the image query and the database image; correct?
A. Correct.
Q. And do you recall if that methodology changed after this article came out in '93?
A. Yes, it did.
Q. How so?
A. Well, in the IEEE computer paper, we described a different way of doing shape measures. And they included the moment calculations.
Q. And you use the moment calculations to determine a distance between the image query and the -- and the database image?
A. Yes.
Q. Sorry?
A. Yeah. Yes.

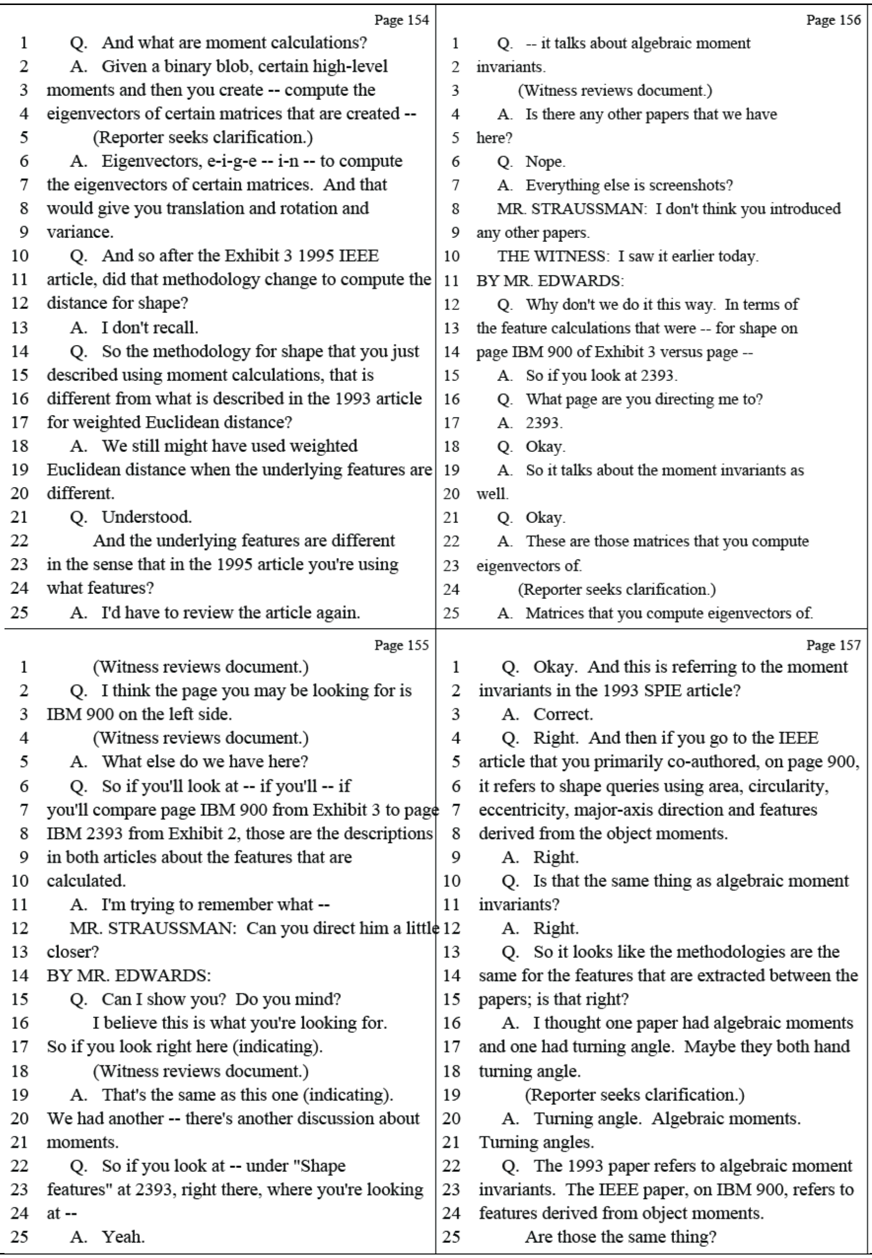

Page 154

Q. And what are moment calculations?
A. Given a binary blob, certain high-level moments and then you create -- compute the eigenvectors of certain matrices that are created --
(Reporter seeks clarification.)
A. Eigenvectors, e-i-g-e -- i-n -- to compute the eigenvectors of certain matrices. And that would give you translation and rotation and variance.
Q. And so after the Exhibit 3 1995 IEEE article, did that methodology change to compute the distance for shape?
A. I don't recall.
Q. So the methodology for shape that you just described using moment calculations, that is different from what is described in the 1993 article for weighted Euclidean distance?
A. We still might have used weighted Euclidean distance when the underlying features are different.
Q. Understood.
And the underlying features are different in the sense that in the 1995 article you're using what features?
A. I'd have to review the article again.

Page 155

(Witness reviews document.)
Q. I think the page you may be looking for is IBM 900 on the left side.
(Witness reviews document.)
A. What else do we have here?
Q. So if you'll look at -- if you'll -- if you'll compare page IBM 900 from Exhibit 3 to page IBM 2393 from Exhibit 2, those are the descriptions in both articles about the features that are calculated.
A. I'm trying to remember what --
MR. STRAUSSMAN: Can you direct him a little closer?
BY MR. EDWARDS:
Q. Can I show you? Do you mind?
I believe this is what you're looking for. So if you look right here (indicating).
(Witness reviews document.)
A. That's the same as this one (indicating). We had another -- there's another discussion about moments.
Q. So if you look at -- under "Shape features" at 2393, right there, where you're looking at --
A. Yeah.

Page 156

Q. -- it talks about algebraic moment invariants.
(Witness reviews document.)
A. Is there any other papers that we have here?
Q. Nope.
A. Everything else is screenshots?
MR. STRAUSSMAN: I don't think you introduced any other papers.
THE WITNESS: I saw it earlier today.
BY MR. EDWARDS:
Q. Why don't we do it this way. In terms of the feature calculations that were -- for shape on page IBM 900 of Exhibit 3 versus page --
A. So if you look at 2393.
Q. What page are you directing me to?
A. 2393.
Q. Okay.
A. So it talks about the moment invariants as well.
Q. Okay.
A. These are those matrices that you compute eigenvectors of.
(Reporter seeks clarification.)
A. Matrices that you compute eigenvectors of.

Page 157

Q. Okay. And this is referring to the moment invariants in the 1993 SPIE article?
A. Correct.
Q. Right. And then if you go to the IEEE article that you primarily co-authored, on page 900, it refers to shape queries using area, circularity, eccentricity, major-axis direction and features derived from the object moments.
A. Right.
Q. Is that the same thing as algebraic moment invariants?
A. Right.
Q. So it looks like the methodologies are the same for the features that are extracted between the papers; is that right?
A. I thought one paper had algebraic moments and one had turning angle. Maybe they both hand turning angle.
(Reporter seeks clarification.)
A. Turning angle. Algebraic moments. Turning angles.
Q. The 1993 paper refers to algebraic moment invariants. The IEEE paper, on IBM 900, refers to features derived from object moments.
Are those the same thing?

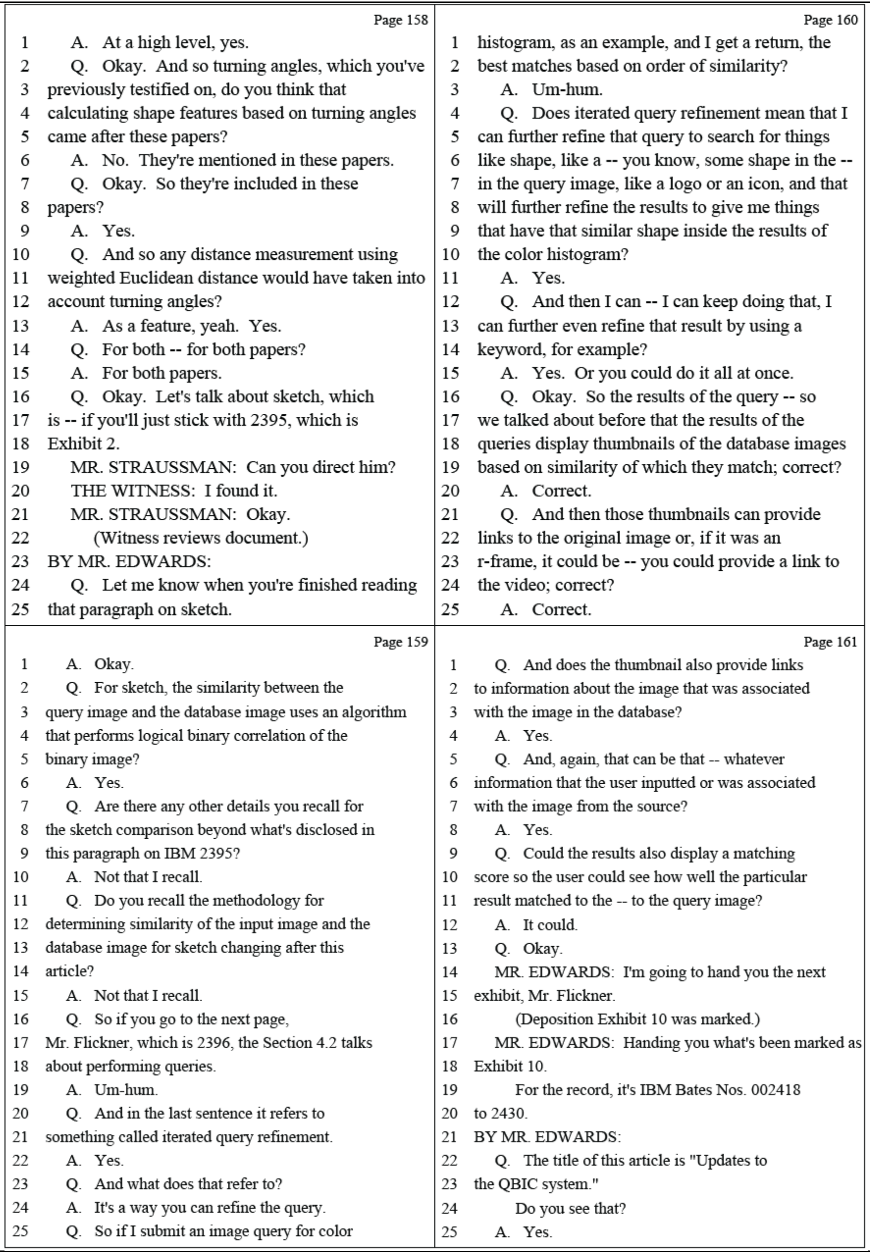

Page 158

A. At a high level, yes.
Q. Okay. And so turning angles, which you've previously testified on, do you think that calculating shape features based on turning angles came after these papers?
A. No. They're mentioned in these papers.
Q. Okay. So they're included in these papers?
A. Yes.
Q. And so any distance measurement using weighted Euclidean distance would have taken into account turning angles?
A. As a feature, yeah. Yes.
Q. For both -- for both papers?
A. For both papers.
Q. Okay. Let's talk about sketch, which is -- if you'll just stick with 2395, which is Exhibit 2.
MR. STRAUSSMAN: Can you direct him?
THE WITNESS: I found it.
MR. STRAUSSMAN: Okay.
(Witness reviews document.)
BY MR. EDWARDS:
Q. Let me know when you're finished reading that paragraph on sketch.

Page 159

A. Okay.
Q. For sketch, the similarity between the query image and the database image uses an algorithm that performs logical binary correlation of the binary image?
A. Yes.
Q. Are there any other details you recall for the sketch comparison beyond what's disclosed in this paragraph on IBM 2395?
A. Not that I recall.
Q. Do you recall the methodology for determining similarity of the input image and the database image for sketch changing after this article?
A. Not that I recall.
Q. So if you go to the next page, Mr. Flickner, which is 2396, the Section 4.2 talks about performing queries.
A. Um-hum.
Q. And in the last sentence it refers to something called iterated query refinement.
A. Yes.
Q. And what does that refer to?
A. It's a way you can refine the query.
Q. So if I submit an image query for color

Page 160

histogram, as an example, and I get a return, the best matches based on order of similarity?
A. Um-hum.
Q. Does iterated query refinement mean that I can further refine that query to search for things like shape, like a -- you know, some shape in the -- in the query image, like a logo or an icon, and that will further refine the results to give me things that have that similar shape inside the results of the color histogram?
A. Yes.
Q. And then I can -- I can keep doing that, I can further even refine that result by using a keyword, for example?
A. Yes. Or you could do it all at once.
Q. Okay. So the results of the query -- so we talked about before that the results of the queries display thumbnails of the database images based on similarity of which they match; correct?
A. Correct.
Q. And then those thumbnails can provide links to the original image or, if it was an r-frame, it could be -- you could provide a link to the video; correct?
A. Correct.

Page 161

Q. And does the thumbnail also provide links to information about the image that was associated with the image in the database?
A. Yes.
Q. And, again, that can be that -- whatever information that the user inputted or was associated with the image from the source?
A. Yes.
Q. Could the results also display a matching score so the user could see how well the particular result matched to the -- to the query image?
A. It could.
Q. Okay.
MR. EDWARDS: I'm going to hand you the next exhibit, Mr. Flickner.
(Deposition Exhibit 10 was marked.)
MR. EDWARDS: Handing you what's been marked as Exhibit 10.
For the record, it's IBM Bates Nos. 002418 to 2430.
BY MR. EDWARDS:
Q. The title of this article is "Updates to the QBIC system."
Do you see that?
A. Yes.

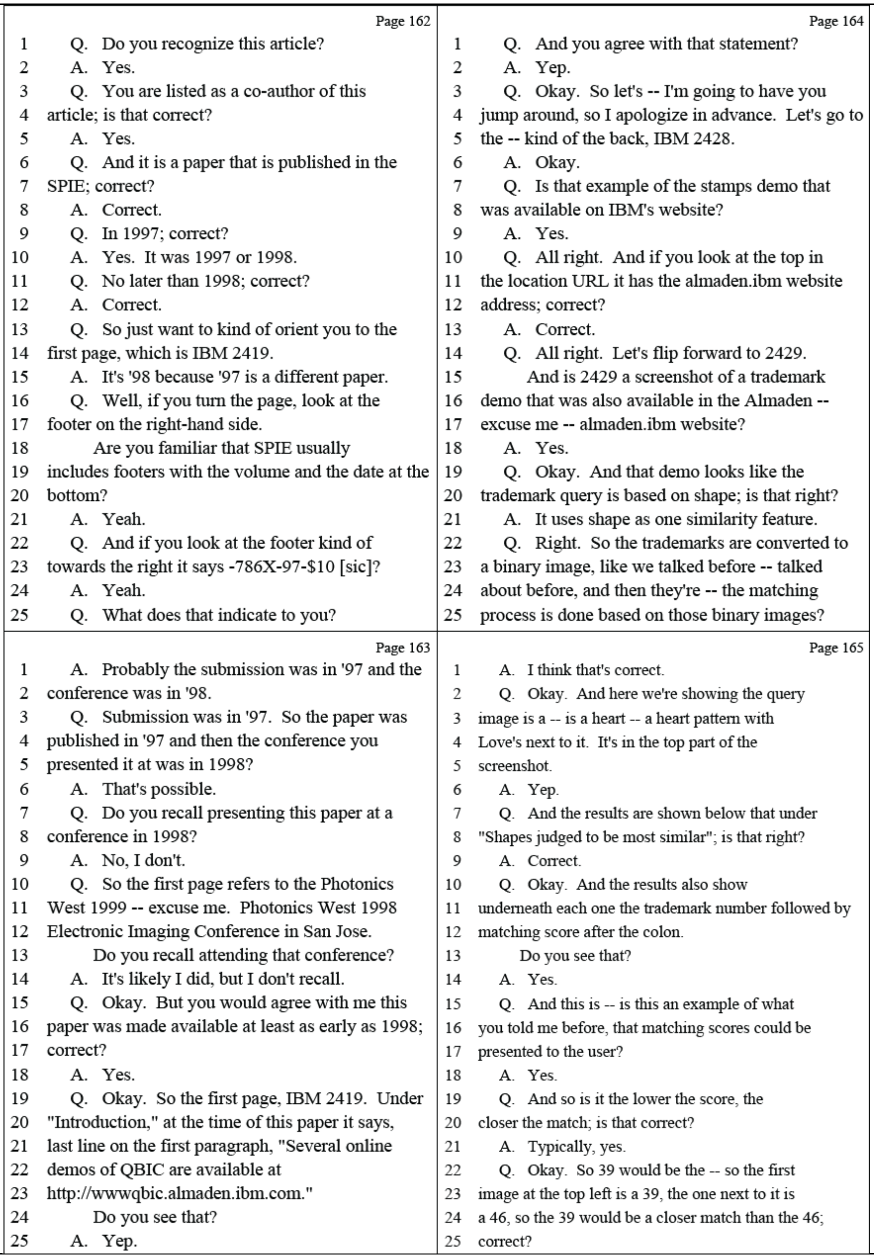

Page 162

Q. Do you recognize this article?
A. Yes.
Q. You are listed as a co-author of this article; is that correct?
A. Yes.
Q. And it is a paper that is published in the SPIE; correct?
A. Correct.
Q. In 1997; correct?
A. Yes. It was 1997 or 1998.
Q. No later than 1998; correct?
A. Correct.
Q. So just want to kind of orient you to the first page, which is IBM 2419.
A. It's '98 because '97 is a different paper.
Q. Well, if you turn the page, look at the footer on the right-hand side.
Are you familiar that SPIE usually includes footers with the volume and the date at the bottom?
A. Yeah.
Q. And if you look at the footer kind of towards the right it says -786X-97-$10 [sic]?
A. Yeah.
Q. What does that indicate to you?

Page 163

A. Probably the submission was in '97 and the conference was in '98.
Q. Submission was in '97. So the paper was published in '97 and then the conference you presented it at was in 1998?
A. That's possible.
Q. Do you recall presenting this paper at a conference in 1998?
A. No, I don't.
Q. So the first page refers to the Photonics West 1999 -- excuse me. Photonics West 1998 Electronic Imaging Conference in San Jose.
Do you recall attending that conference?
A. It's likely I did, but I don't recall.
Q. Okay. But you would agree with me this paper was made available at least as early as 1998; correct?
A. Yes.
Q. Okay. So the first page, IBM 2419. Under "Introduction," at the time of this paper it says, last line on the first paragraph, "Several online demos of QBIC are available at http://wwwqbic.almaden.ibm.com."
Do you see that?
A. Yep.

Page 164

Q. And you agree with that statement?
A. Yep.
Q. Okay. So let's -- I'm going to have you jump around, so I apologize in advance. Let's go to the -- kind of the back, IBM 2428.
A. Okay.
Q. Is that example of the stamps demo that was available on IBM's website?
A. Yes.
Q. All right. And if you look at the top in the location URL it has the almaden.ibm website address; correct?
A. Correct.
Q. All right. Let's flip forward to 2429.
And is 2429 a screenshot of a trademark demo that was also available in the Almaden -- excuse me -- almaden.ibm website?
A. Yes.
Q. Okay. And that demo looks like the trademark query is based on shape; is that right?
A. It uses shape as one similarity feature.
Q. Right. So the trademarks are converted to a binary image, like we talked before -- talked about before, and then they're -- the matching process is done based on those binary images?

Page 165

A. I think that's correct.
Q. Okay. And here we're showing the query image is a -- is a heart -- a heart pattern with Love's next to it. It's in the top part of the screenshot.
A. Yep.
Q. And the results are shown below that under "Shapes judged to be most similar"; is that right?
A. Correct.
Q. Okay. And the results also show underneath each one the trademark number followed by matching score after the colon.
Do you see that?
A. Yes.
Q. And this is -- is this an example of what you told me before, that matching scores could be presented to the user?
A. Yes.
Q. And so is it the lower the score, the closer the match; is that correct?
A. Typically, yes.
Q. Okay. So 39 would be the -- so the first image at the top left is a 39, the one next to it is a 46, so the 39 would be a closer match than the 46; correct?

Element	'897 Claim Recitation	Exemplary Citations in Reference
		Id. at 174–182 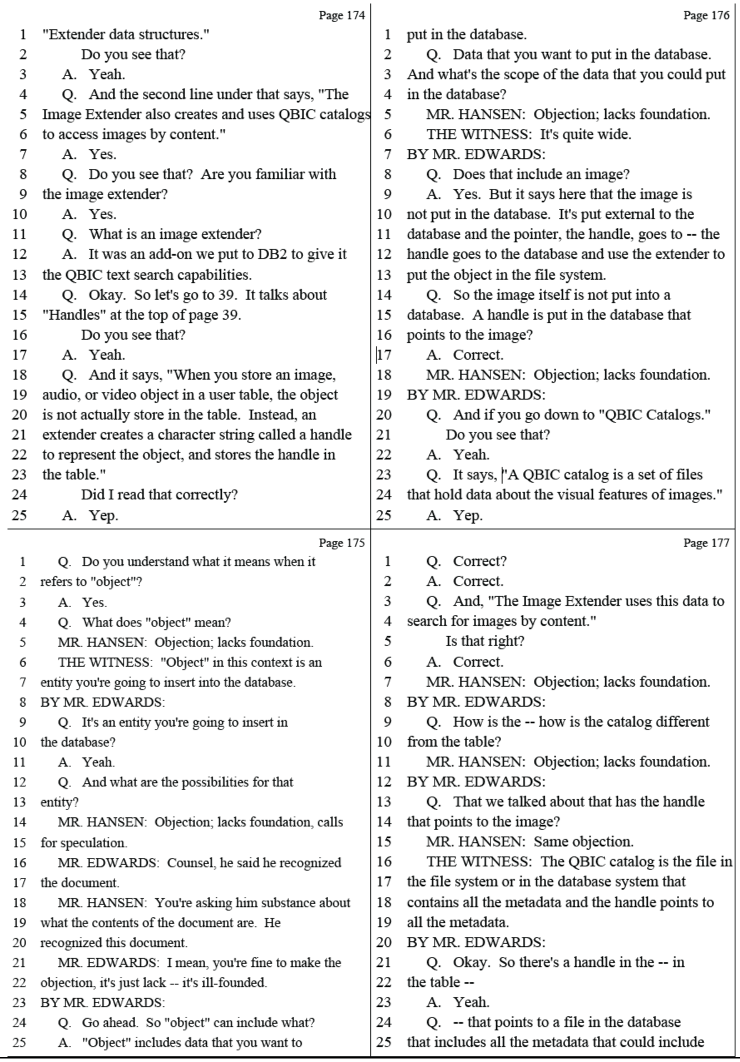

Page 174

"Extender data structures."
Do you see that?
A. Yeah.
Q. And the second line under that says, "The Image Extender also creates and uses QBIC catalogs to access images by content."
A. Yes.
Q. Do you see that? Are you familiar with the image extender?
A. Yes.
Q. What is an image extender?
A. It was an add-on we put to DB2 to give it the QBIC text search capabilities.
Q. Okay. So let's go to 39. It talks about "Handles" at the top of page 39.
Do you see that?
A. Yeah.
Q. And it says, "When you store an image, audio, or video object in a user table, the object is not actually store in the table. Instead, an extender creates a character string called a handle to represent the object, and stores the handle in the table."
Did I read that correctly?
A. Yep.

Page 175

Q. Do you understand what it means when it refers to "object"?
A. Yes.
Q. What does "object" mean?
MR. HANSEN: Objection; lacks foundation.
THE WITNESS: "Object" in this context is an entity you're going to insert into the database.
BY MR. EDWARDS:
Q. It's an entity you're going to insert in the database?
A. Yeah.
Q. And what are the possibilities for that entity?
MR. HANSEN: Objection; lacks foundation, calls for speculation.
MR. EDWARDS: Counsel, he said he recognized the document.
MR. HANSEN: You're asking him substance about what the contents of the document are. He recognized this document.
MR. EDWARDS: I mean, you're fine to make the objection, it's just lack -- it's ill-founded.
BY MR. EDWARDS:
Q. Go ahead. So "object" can include what?
A. "Object" includes data that you want to

Page 176

put in the database.
Q. Data that you want to put in the database. And what's the scope of the data that you could put in the database?
MR. HANSEN: Objection; lacks foundation.
THE WITNESS: It's quite wide.
BY MR. EDWARDS:
Q. Does that include an image?
A. Yes. But it says here that the image is not put in the database. It's put external to the database and the pointer, the handle, goes to -- the handle goes to the database and use the extender to put the object in the file system.
Q. So the image itself is not put into a database. A handle is put in the database that points to the image?
A. Correct.
MR. HANSEN: Objection; lacks foundation.
BY MR. EDWARDS:
Q. And if you go down to "QBIC Catalogs."
Do you see that?
A. Yeah.
Q. It says, "A QBIC catalog is a set of files that hold data about the visual features of images."
A. Yep.

Page 177

Q. Correct?
A. Correct.
Q. And, "The Image Extender uses this data to search for images by content."
Is that right?
A. Correct.
MR. HANSEN: Objection; lacks foundation.
BY MR. EDWARDS:
Q. How is the -- how is the catalog different from the table?
MR. HANSEN: Objection; lacks foundation.
BY MR. EDWARDS:
Q. That we talked about that has the handle that points to the image?
MR. HANSEN: Same objection.
THE WITNESS: The QBIC catalog is the file in the file system or in the database system that contains all the metadata and the handle points to all the metadata.
BY MR. EDWARDS:
Q. Okay. So there's a handle in the -- in the table --
A. Yeah.
Q. -- that points to a file in the database that includes all the metadata that could include

			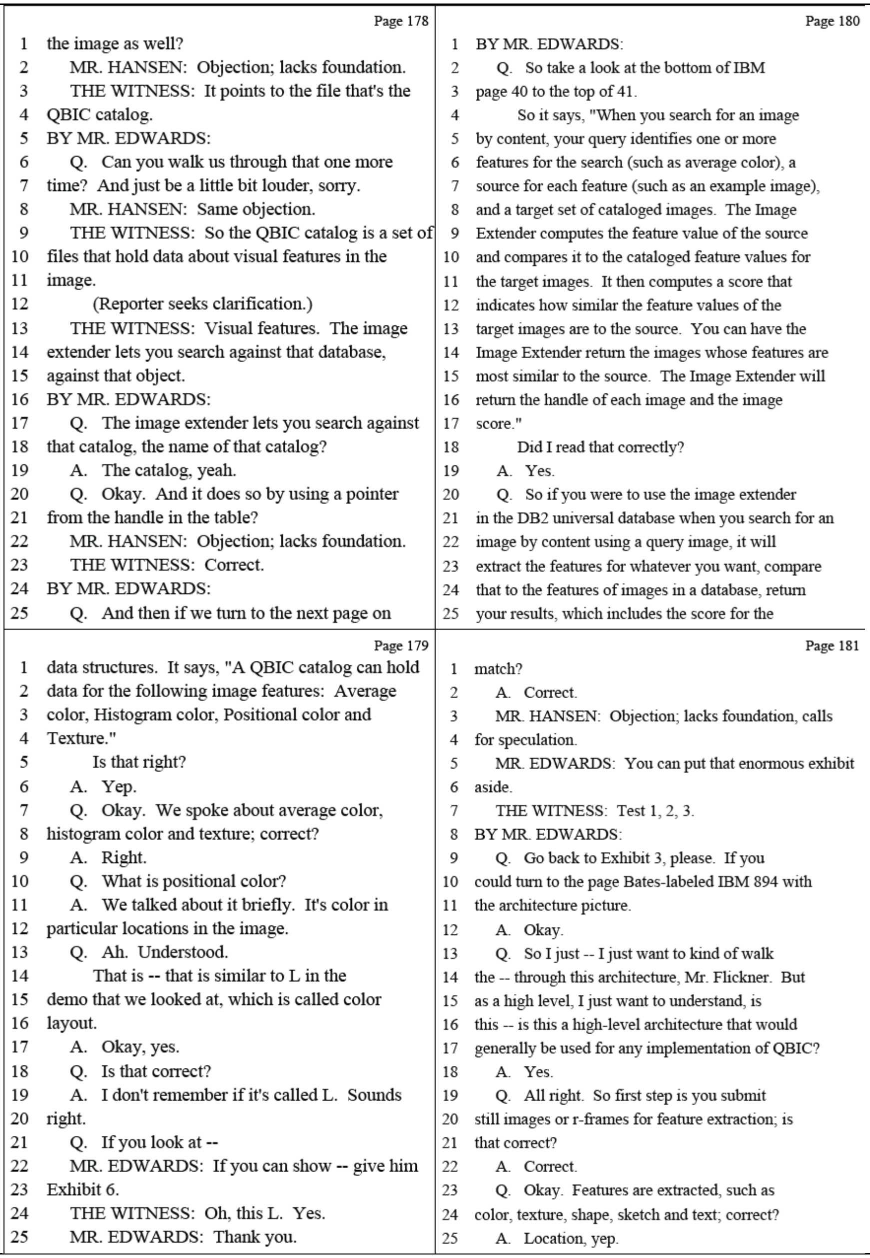

Page 178

the image as well?

MR. HANSEN: Objection; lacks foundation.

THE WITNESS: It points to the file that's the QBIC catalog.

BY MR. EDWARDS:

Q. Can you walk us through that one more time? And just be a little bit louder, sorry.

MR. HANSEN: Same objection.

THE WITNESS: So the QBIC catalog is a set of files that hold data about visual features in the image.

(Reporter seeks clarification.)

THE WITNESS: Visual features. The image extender lets you search against that database, against that object.

BY MR. EDWARDS:

Q. The image extender lets you search against that catalog, the name of that catalog?

A. The catalog, yeah.

Q. Okay. And it does so by using a pointer from the handle in the table?

MR. HANSEN: Objection; lacks foundation.

THE WITNESS: Correct.

BY MR. EDWARDS:

Q. And then if we turn to the next page on

Page 179

data structures. It says, "A QBIC catalog can hold data for the following image features: Average color, Histogram color, Positional color and Texture."

Is that right?

A. Yep.

Q. Okay. We spoke about average color, histogram color and texture; correct?

A. Right.

Q. What is positional color?

A. We talked about it briefly. It's color in particular locations in the image.

Q. Ah. Understood.

That is -- that is similar to L in the demo that we looked at, which is called color layout.

A. Okay, yes.

Q. Is that correct?

A. I don't remember if it's called L. Sounds right.

Q. If you look at --

MR. EDWARDS: If you can show -- give him Exhibit 6.

THE WITNESS: Oh, this L. Yes.

MR. EDWARDS: Thank you.

Page 180

BY MR. EDWARDS:

Q. So take a look at the bottom of IBM page 40 to the top of 41.

So it says, "When you search for an image by content, your query identifies one or more features for the search (such as average color), a source for each feature (such as an example image), and a target set of cataloged images. The Image Extender computes the feature value of the source and compares it to the cataloged feature values for the target images. It then computes a score that indicates how similar the feature values of the target images are to the source. You can have the Image Extender return the images whose features are most similar to the source. The Image Extender will return the handle of each image and the image score."

Did I read that correctly?

A. Yes.

Q. So if you were to use the image extender in the DB2 universal database when you search for an image by content using a query image, it will extract the features for whatever you want, compare that to the features of images in a database, return your results, which includes the score for the

Page 181

match?

A. Correct.

MR. HANSEN: Objection; lacks foundation, calls for speculation.

MR. EDWARDS: You can put that enormous exhibit aside.

THE WITNESS: Test 1, 2, 3.

BY MR. EDWARDS:

Q. Go back to Exhibit 3, please. If you could turn to the page Bates-labeled IBM 894 with the architecture picture.

A. Okay.

Q. So I just -- I just want to kind of walk the -- through this architecture, Mr. Flickner. But as a high level, I just want to understand, is this -- is this a high-level architecture that would generally be used for any implementation of QBIC?

A. Yes.

Q. All right. So first step is you submit still images or r-frames for feature extraction; is that correct?

A. Correct.

Q. Okay. Features are extracted, such as color, texture, shape, sketch and text; correct?

A. Location, yep.

Element	'897 Claim Recitation	Exemplary Citations in Reference
		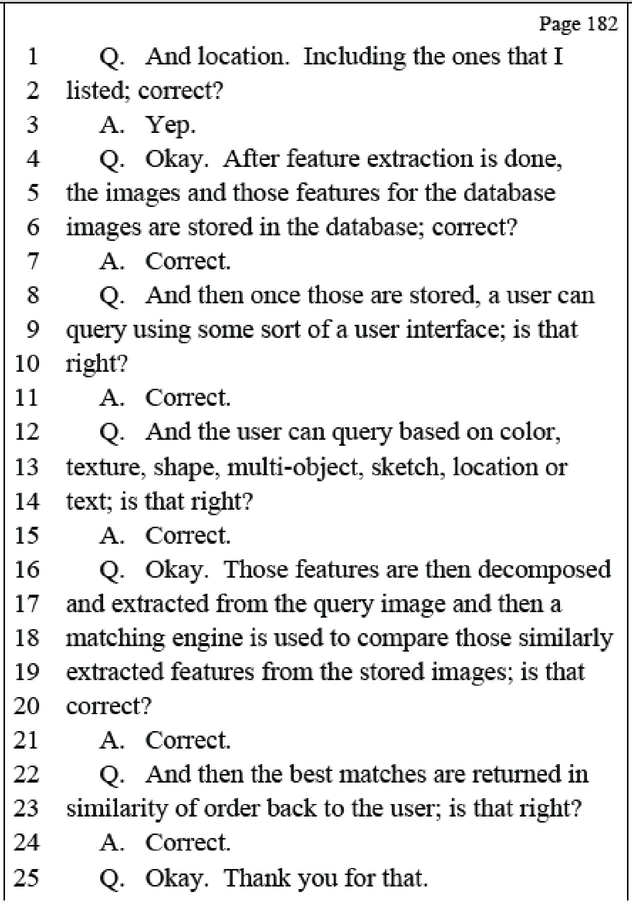 Page 182 Q. And location. Including the ones that I listed; correct? A. Yep. Q. Okay. After feature extraction is done, the images and those features for the database images are stored in the database; correct? A. Correct. Q. And then once those are stored, a user can query using some sort of a user interface; is that right? A. Correct. Q. And the user can query based on color, texture, shape, multi-object, sketch, location or text; is that right? A. Correct. Q. Okay. Those features are then decomposed and extracted from the query image and then a matching engine is used to compare those similarly extracted features from the stored images; is that correct? A. Correct. Q. And then the best matches are returned in similarity of order back to the user; is that right? A. Correct. Q. Okay. Thank you for that. *Id.* at 190–191 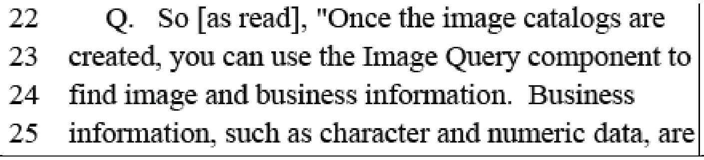 Q. So [as read], "Once the image catalogs are created, you can use the Image Query component to find image and business information. Business information, such as character and numeric data, are

		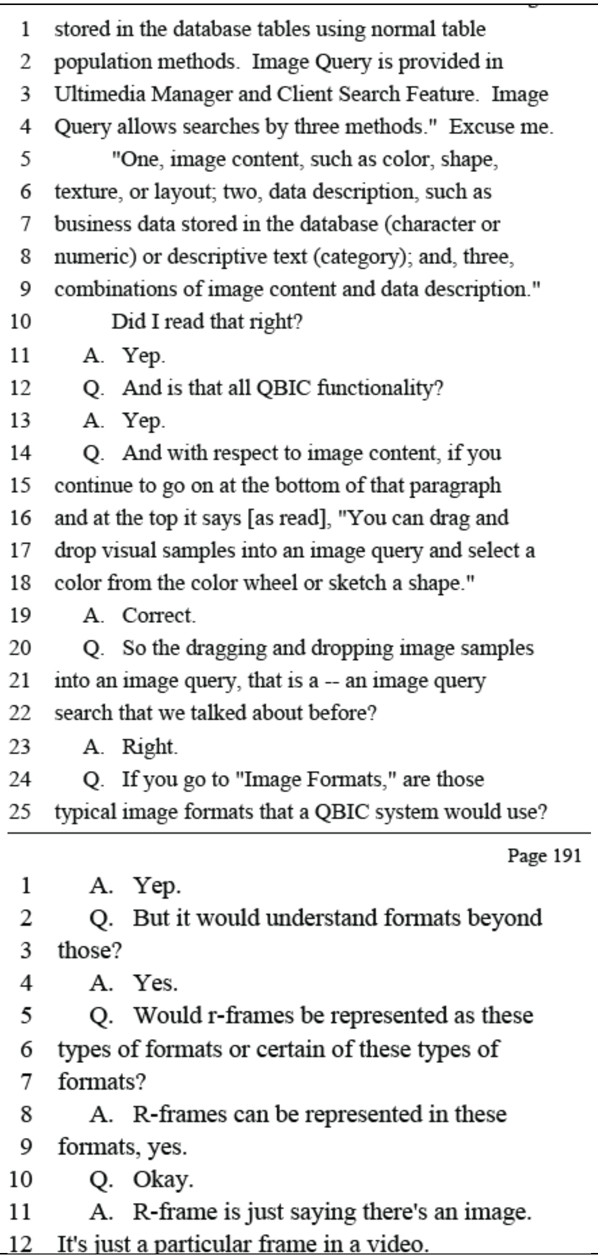stored in the database tables using normal table population methods. Image Query is provided in Ultimedia Manager and Client Search Feature. Image Query allows searches by three methods." Excuse me. "One, image content, such as color, shape, texture, or layout; two, data description, such as business data stored in the database (character or numeric) or descriptive text (category); and, three, combinations of image content and data description." Did I read that right? A. Yep. Q. And is that all QBIC functionality? A. Yep. Q. And with respect to image content, if you continue to go on at the bottom of that paragraph and at the top it says [as read], "You can drag and drop visual samples into an image query and select a color from the color wheel or sketch a shape." A. Correct. Q. So the dragging and dropping image samples into an image query, that is a -- an image query search that we talked about before? A. Right. Q. If you go to "Image Formats," are those typical image formats that a QBIC system would use? Page 191 A. Yep. Q. But it would understand formats beyond those? A. Yes. Q. Would r-frames be represented as these types of formats or certain of these types of formats? A. R-frames can be represented in these formats, yes. Q. Okay. A. R-frame is just saying there's an image. It's just a particular frame in a video.

Element	'897 Claim Recitation	Exemplary Citations in Reference
		Id. at 198 – 199 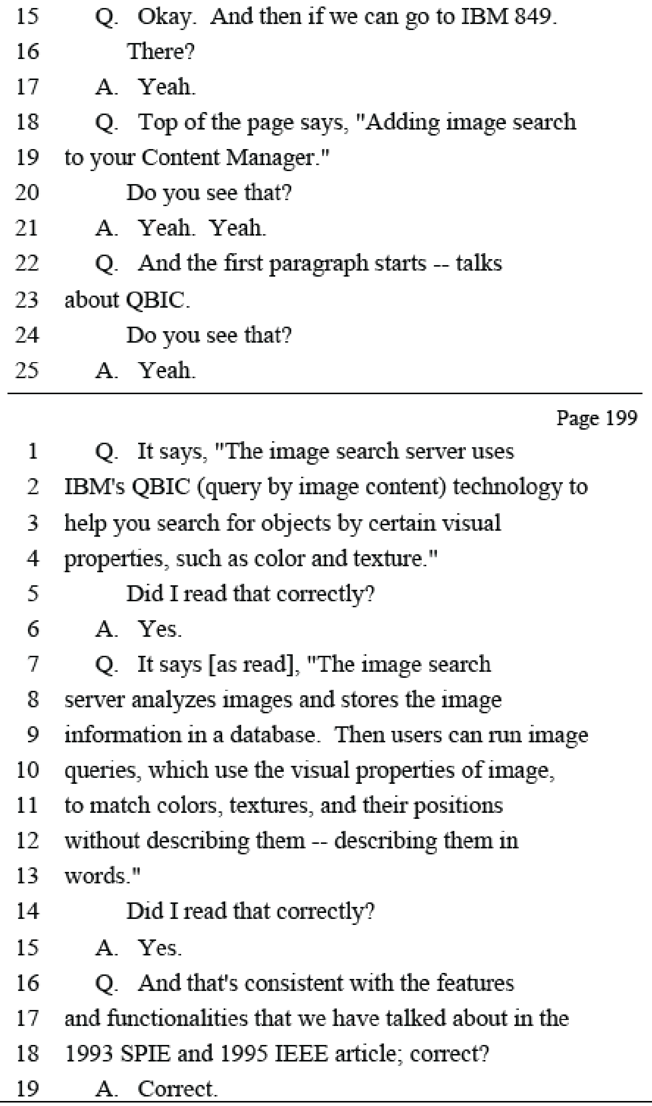 Q. Okay. And then if we can go to IBM 849. There? A. Yeah. Q. Top of the page says, "Adding image search to your Content Manager." Do you see that? A. Yeah. Yeah. Q. And the first paragraph starts -- talks about QBIC. Do you see that? A. Yeah. Page 199 Q. It says, "The image search server uses IBM's QBIC (query by image content) technology to help you search for objects by certain visual properties, such as color and texture." Did I read that correctly? A. Yes. Q. It says [as read], "The image search server analyzes images and stores the image information in a database. Then users can run image queries, which use the visual properties of image, to match colors, textures, and their positions without describing them -- describing them in words." Did I read that correctly? A. Yes. Q. And that's consistent with the features and functionalities that we have talked about in the 1993 SPIE and 1995 IEEE article; correct? A. Correct.

Element	'897 Claim Recitation	Exemplary Citations in Reference
		Id. at 312 – 313 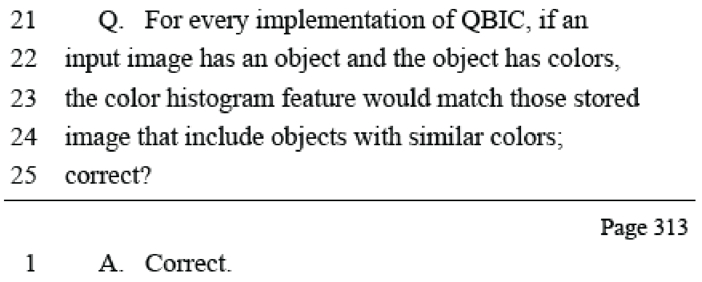 Q. For every implementation of QBIC, if an input image has an object and the object has colors, the color histogram feature would match those stored image that include objects with similar colors; correct? Page 313 A. Correct. *Id.* at 317–322 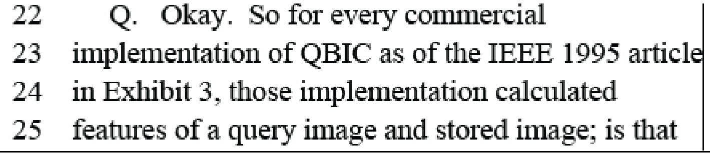 Q. Okay. So for every commercial implementation of QBIC as of the IEEE 1995 article in Exhibit 3, those implementation calculated features of a query image and stored image; is that

Page 318

correct?
MR. HANSEN: Objection; lacks foundation.
THE WITNESS: Correct.
BY MR. EDWARDS:
Q. And those calculations of features included color, shape and/or texture in all implementations that we looked at today, correct?
MR. HANSEN: Same objections.
THE WITNESS: I don't remember if all versions supported all features.
BY MR. EDWARDS:
Q. If the 1995 article references calculating color, shape and texture, you would agree with me that the commercial embodiments implementing QBIC at the time would have been able to calculate color, shape and texture as part of feature extraction; correct?
MR. HANSEN: Objection; lacks foundation.
THE WITNESS: I don't remember exactly which products incorporated any subsets of the features.
BY MR. EDWARDS:
Q. Do you recall that all the products used color?
MR. HANSEN: Objection; lacks foundation.
THE WITNESS: I'm pretty sure they all did use

Page 319

color.
BY MR. EDWARDS:
Q. And color histograms specifically?
MR. HANSEN: Same objection.
(Reporter seeks clarification.)
BY MR. EDWARDS:
Q. Histograms specifically.
(Reporter seeks clarification.)
THE WITNESS: Most likely, yes.
BY MR. EDWARDS:
Q. Well, pick up Exhibit 12, please.
And go to IBM 2264.
A. Okay.
Q. So the Ultimedia Manager 1.1 was able to perform an image query using color, shape and texture; correct?
MR. HANSEN: Objection; lacks foundation.
(Witness reviews document.)
BY MR. EDWARDS:
Q. Let me orient you a little more. If you go to the bottom of IBM 2264, where it says "Image Query."
You see that?
A. Yeah.
Q. Where the bullet points start, it says,

Page 320

[as read] "Image Query allows searches by three methods: Image content, such as color, shape and texture."
You see that?
A. Yeah.
Q. So Ultimedia Manager 1.1 was able to perform an image query using color, shape and texture, correct?
MR. HANSEN: Objection; lacks foundation.
THE WITNESS: And layout.
BY MR. EDWARDS:
Q. And layout. But also color, shape and texture, correct?
A. Yes.
MR. HANSEN: Same objection.
BY MR. EDWARDS:
Q. And then if you go to the last sentence it goes on -- and the bottom of the left column on that same page, going to the top of the right column, it says, "You can drag and drop visual samples into an image query."
You see that?
A. Yeah.
Q. And that would include images, correct?
A. Correct.

Page 321

MR. HANSEN: Objection; lacks foundation.
BY MR. EDWARDS:
Q. Okay?
A. Correct.
Q. Mr. Flickner, every commercial implementation of QBIC as of the Exhibit 3, IEEE 1995 article, returned results of the matching process to the user in the form of a stored image and information associated with that image; is that fair?
MR. HANSEN: Objection; lacks foundation.
THE WITNESS: Yes.
BY MR. EDWARDS:
Q. Okay. And a thumbnail is an image; correct?
A. Thumbnail is an image.
Q. And for every commercial embodiment of QBIC as of the 1995 IEEE article in Exhibit 3, those implementations calculated a distance metric as a measure of similarity for each feature between an input image and a stored image; correct?
MR. HANSEN: Objection; lacks foundation.
THE WITNESS: Correct.
BY MR. EDWARDS:
Q. Okay. And a distance measure value

Element	'897 Claim Recitation	Exemplary Citations in Reference
		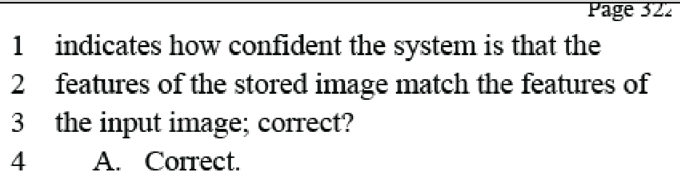Page 322 indicates how confident the system is that the features of the stored image match the features of the input image; correct? A. Correct. *Id.* at 53–56 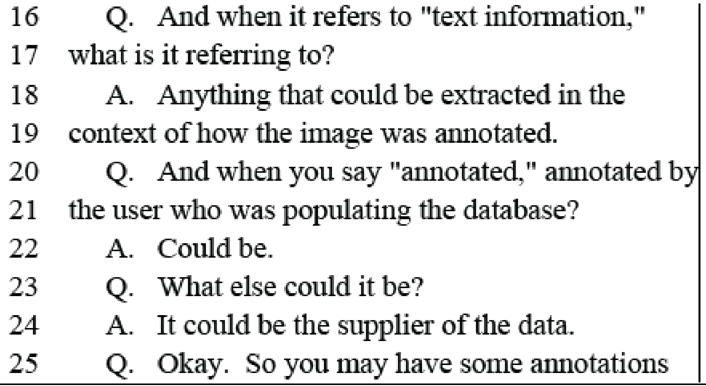Q. And when it refers to "text information," what is it referring to? A. Anything that could be extracted in the context of how the image was annotated. Q. And when you say "annotated," annotated by the user who was populating the database? A. Could be. Q. What else could it be? A. It could be the supplier of the data. Q. Okay. So you may have some annotations

Element	'897 Claim Recitation	Exemplary Citations in Reference
		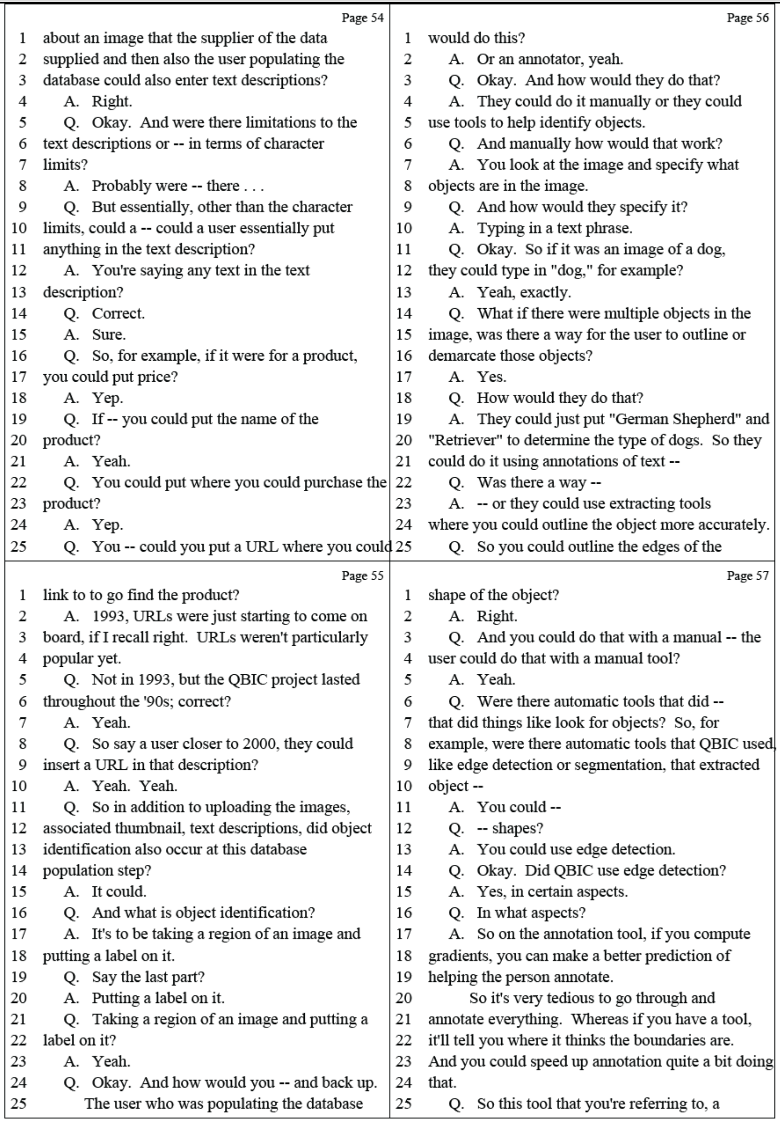 Page 54 about an image that the supplier of the data supplied and then also the user populating the database could also enter text descriptions? A. Right. Q. Okay. And were there limitations to the text descriptions or -- in terms of character limits? A. Probably were -- there . . . Q. But essentially, other than the character limits, could a -- could a user essentially put anything in the text description? A. You're saying any text in the text description? Q. Correct. A. Sure. Q. So, for example, if it were for a product, you could put price? A. Yep. Q. If -- you could put the name of the product? A. Yeah. Q. You could put where you could purchase the product? A. Yep. Q. You -- could you put a URL where you could Page 55 link to to go find the product? A. 1993, URLs were just starting to come on board, if I recall right. URLs weren't particularly popular yet. Q. Not in 1993, but the QBIC project lasted throughout the '90s; correct? A. Yeah. Q. So say a user closer to 2000, they could insert a URL in that description? A. Yeah. Yeah. Q. So in addition to uploading the images, associated thumbnail, text descriptions, did object identification also occur at this database population step? A. It could. Q. And what is object identification? A. It's to be taking a region of an image and putting a label on it. Q. Say the last part? A. Putting a label on it. Q. Taking a region of an image and putting a label on it? A. Yeah. Q. Okay. And how would you -- and back up. The user who was populating the database Page 56 would do this? A. Or an annotator, yeah. Q. Okay. And how would they do that? A. They could do it manually or they could use tools to help identify objects. Q. And manually how would that work? A. You look at the image and specify what objects are in the image. Q. And how would they specify it? A. Typing in a text phrase. Q. Okay. So if it was an image of a dog, they could type in "dog," for example? A. Yeah, exactly. Q. What if there were multiple objects in the image, was there a way for the user to outline or demarcate those objects? A. Yes. Q. How would they do that? A. They could just put "German Shepherd" and "Retriever" to determine the type of dogs. So they could do it using annotations of text -- Q. Was there a way -- A. -- or they could use extracting tools where you could outline the object more accurately. Q. So you could outline the edges of the Page 57 shape of the object? A. Right. Q. And you could do that with a manual -- the user could do that with a manual tool? A. Yeah. Q. Were there automatic tools that did -- that did things like look for objects? So, for example, were there automatic tools that QBIC used, like edge detection or segmentation, that extracted object -- A. You could -- Q. -- shapes? A. You could use edge detection. Q. Okay. Did QBIC use edge detection? A. Yes, in certain aspects. Q. In what aspects? A. So on the annotation tool, if you compute gradients, you can make a better prediction of helping the person annotate. So it's very tedious to go through and annotate everything. Whereas if you have a tool, it'll tell you where it thinks the boundaries are. And you could speed up annotation quite a bit doing that. Q. So this tool that you're referring to, a

Element	'897 Claim Recitation	Exemplary Citations in Reference
		Id. at 60 – 62 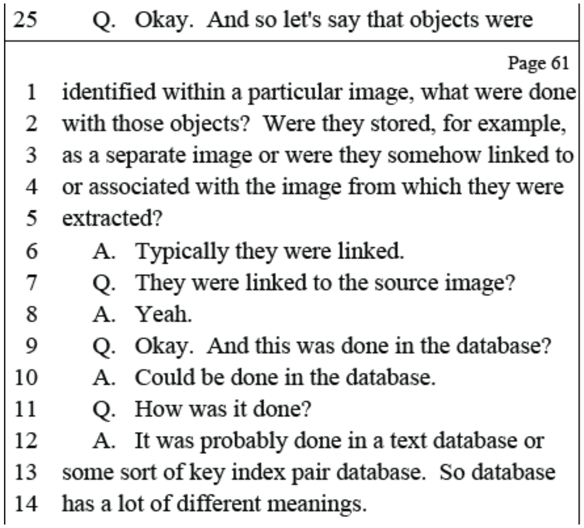 Q. Okay. And so let's say that objects were Page 61 identified within a particular image, what were done with those objects? Were they stored, for example, as a separate image or were they somehow linked to or associated with the image from which they were extracted? A. Typically they were linked. Q. They were linked to the source image? A. Yeah. Q. Okay. And this was done in the database? A. Could be done in the database. Q. How was it done? A. It was probably done in a text database or some sort of key index pair database. So database has a lot of different meanings. 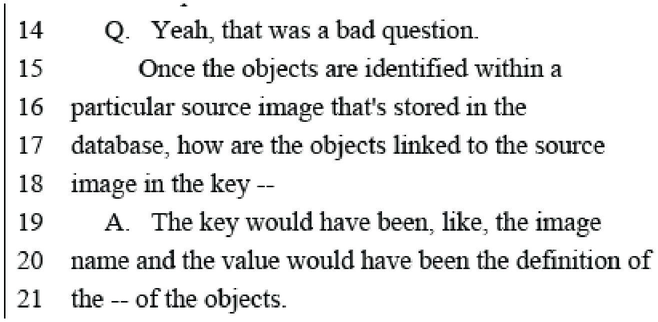 Q. Yeah, that was a bad question. Once the objects are identified within a particular source image that's stored in the database, how are the objects linked to the source image in the key -- A. The key would have been, like, the image name and the value would have been the definition of the -- of the objects. *Id.* at 103

Element	'897 Claim Recitation	Exemplary Citations in Reference
		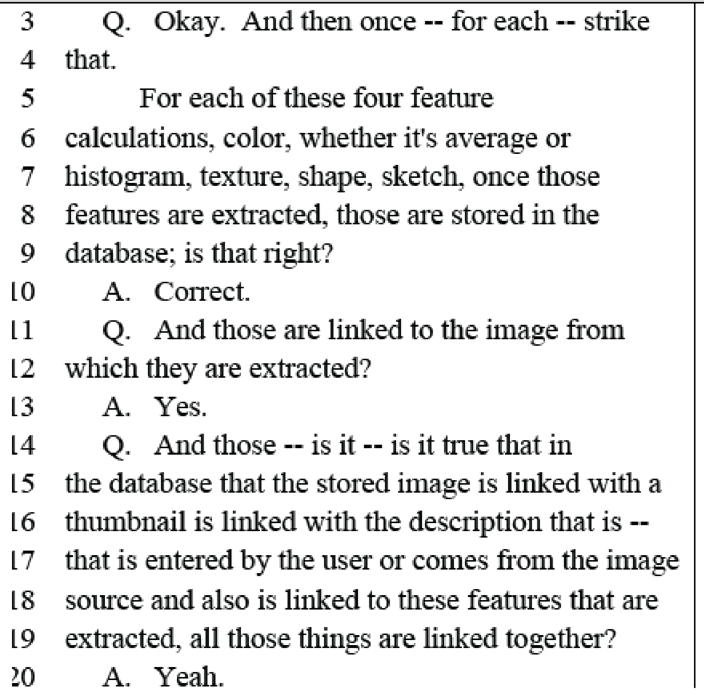 Q. Okay. And then once -- for each -- strike that. For each of these four feature calculations, color, whether it's average or histogram, texture, shape, sketch, once those features are extracted, those are stored in the database; is that right? A. Correct. Q. And those are linked to the image from which they are extracted? A. Yes. Q. And those -- is it -- is it true that in the database that the stored image is linked with a thumbnail is linked with the description that is -- that is entered by the user or comes from the image source and also is linked to these features that are extracted, all those things are linked together? A. Yeah. *Id.* at 129 – 131

Element	'897 Claim Recitation	Exemplary Citations in Reference
		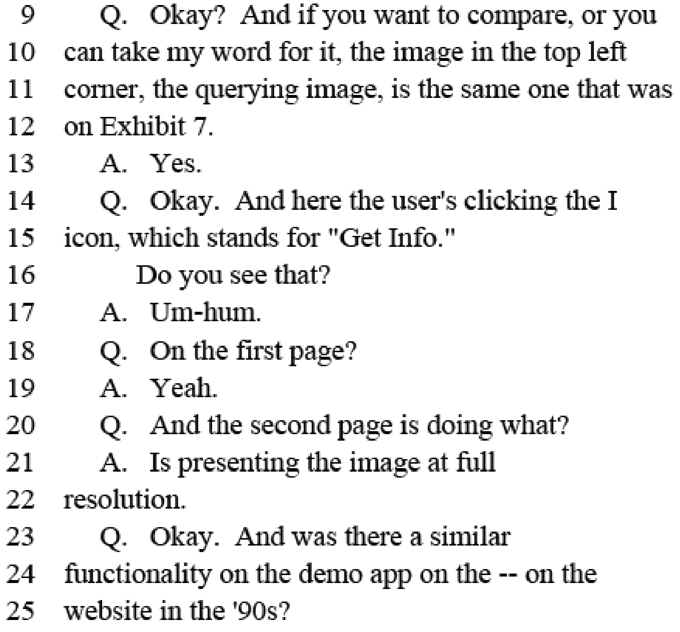 Q. Okay? And if you want to compare, or you can take my word for it, the image in the top left corner, the querying image, is the same one that was on Exhibit 7. A. Yes. Q. Okay. And here the user's clicking the I icon, which stands for "Get Info." Do you see that? A. Um-hum. Q. On the first page? A. Yeah. Q. And the second page is doing what? A. Is presenting the image at full resolution. Q. Okay. And was there a similar functionality on the demo app on the -- on the website in the '90s?

		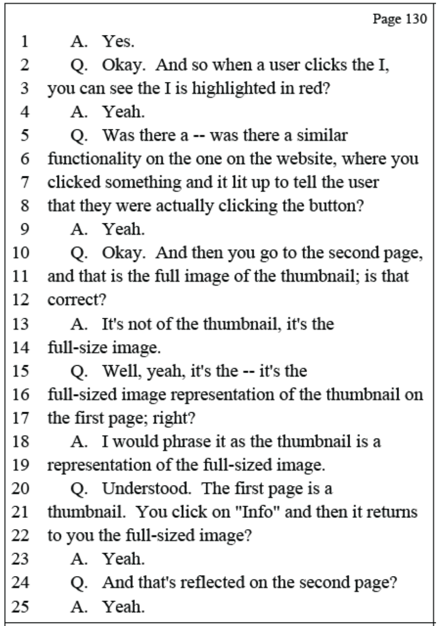Page 130 A. Yes. Q. Okay. And so when a user clicks the I, you can see the I is highlighted in red? A. Yeah. Q. Was there a -- was there a similar functionality on the one on the website, where you clicked something and it lit up to tell the user that they were actually clicking the button? A. Yeah. Q. Okay. And then you go to the second page, and that is the full image of the thumbnail; is that correct? A. It's not of the thumbnail, it's the full-size image. Q. Well, yeah, it's the -- it's the full-sized image representation of the thumbnail on the first page; right? A. I would phrase it as the thumbnail is a representation of the full-sized image. Q. Understood. The first page is a thumbnail. You click on "Info" and then it returns to you the full-sized image? A. Yeah. Q. And that's reflected on the second page? A. Yeah. 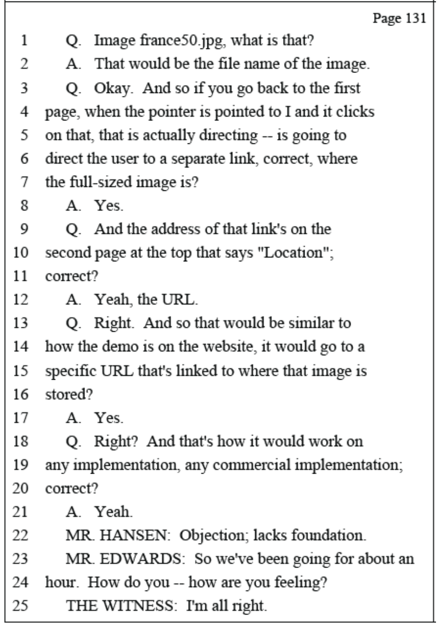Page 131 Q. Image france50.jpg, what is that? A. That would be the file name of the image. Q. Okay. And so if you go back to the first page, when the pointer is pointed to I and it clicks on that, that is actually directing -- is going to direct the user to a separate link, correct, where the full-sized image is? A. Yes. Q. And the address of that link's on the second page at the top that says "Location"; correct? A. Yeah, the URL. Q. Right. And so that would be similar to how the demo is on the website, it would go to a specific URL that's linked to where that image is stored? A. Yes. Q. Right? And that's how it would work on any implementation, any commercial implementation; correct? A. Yeah. MR. HANSEN: Objection; lacks foundation. MR. EDWARDS: So we've been going for about an hour. How do you -- how are you feeling? THE WITNESS: I'm all right.

Element	'897 Claim Recitation	Exemplary Citations in Reference
		Id. at 161 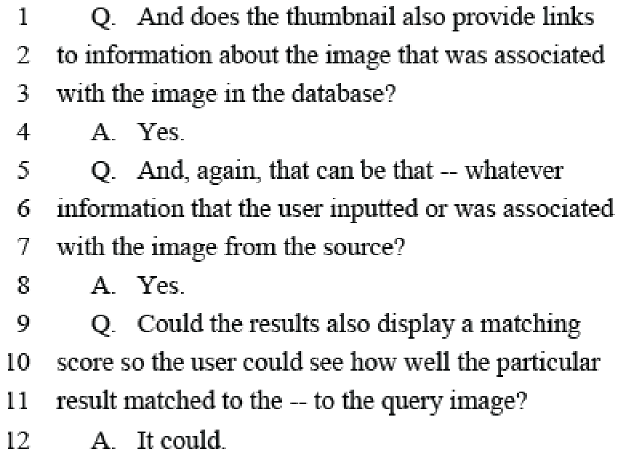 Q. And does the thumbnail also provide links to information about the image that was associated with the image in the database? A. Yes. Q. And, again, that can be that -- whatever information that the user inputted or was associated with the image from the source? A. Yes. Q. Could the results also display a matching score so the user could see how well the particular result matched to the -- to the query image? A. It could. *Id.* at 321 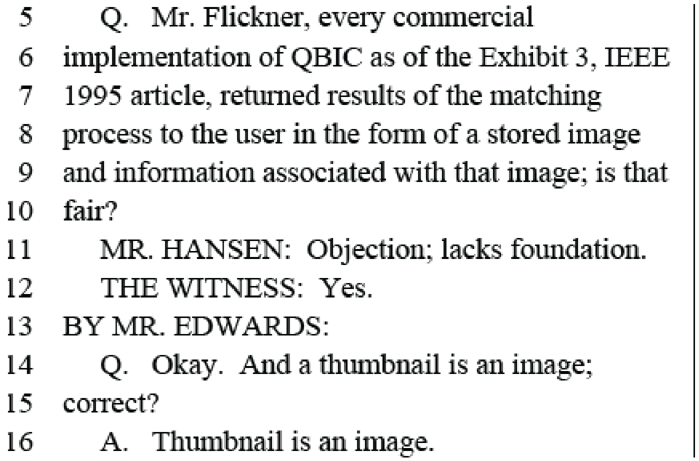 Q. Mr. Flickner, every commercial implementation of QBIC as of the Exhibit 3, IEEE 1995 article, returned results of the matching process to the user in the form of a stored image and information associated with that image; is that fair? MR. HANSEN: Objection; lacks foundation. THE WITNESS: Yes. BY MR. EDWARDS: Q. Okay. And a thumbnail is an image; correct? A. Thumbnail is an image. *Id.* at 118–119

Element	'897 Claim Recitation	Exemplary Citations in Reference
		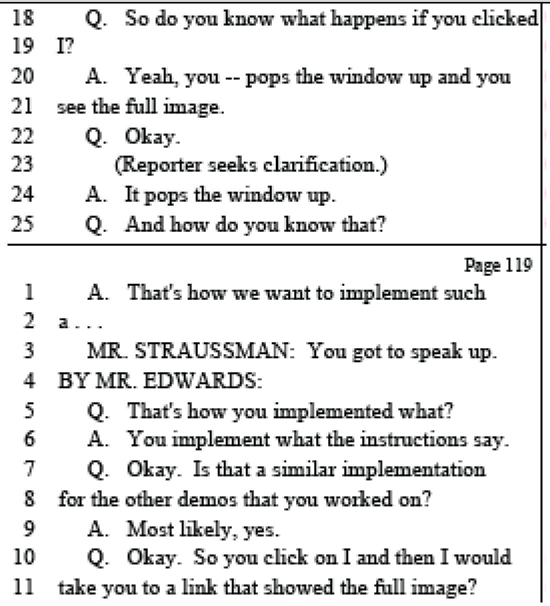 Q. So do you know what happens if you clicked I? A. Yeah, you -- pops the window up and you see the full image. Q. Okay. (Reporter seeks clarification.) A. It pops the window up. Q. And how do you know that? Page 119 A. That's how we want to implement such a . . . MR. STRAUSSMAN: You got to speak up. BY MR. EDWARDS: Q. That's how you implemented what? A. You implement what the instructions say. Q. Okay. Is that a similar implementation for the other demos that you worked on? A. Most likely, yes. Q. Okay. So you click on I and then I would take you to a link that showed the full image? *Id.* at 129 – 131

Element	'897 Claim Recitation	Exemplary Citations in Reference
		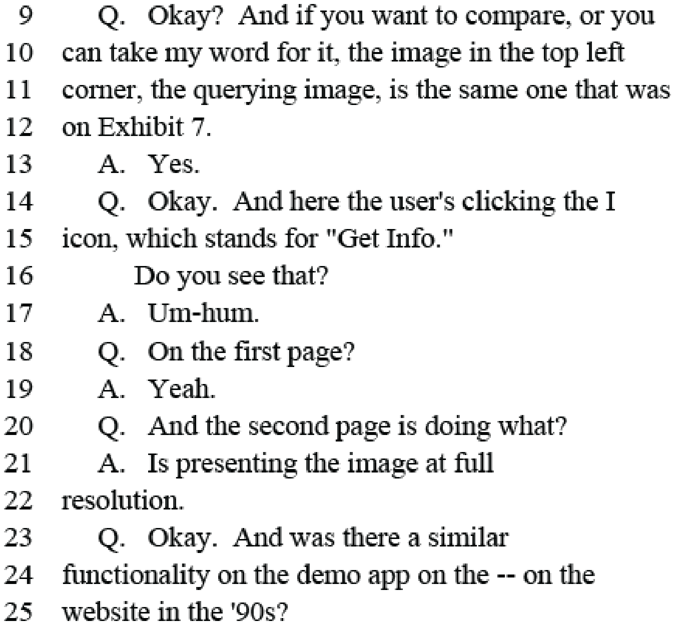 Q. Okay? And if you want to compare, or you can take my word for it, the image in the top left corner, the querying image, is the same one that was on Exhibit 7. A. Yes. Q. Okay. And here the user's clicking the I icon, which stands for "Get Info." Do you see that? A. Um-hum. Q. On the first page? A. Yeah. Q. And the second page is doing what? A. Is presenting the image at full resolution. Q. Okay. And was there a similar functionality on the demo app on the -- on the website in the '90s?

		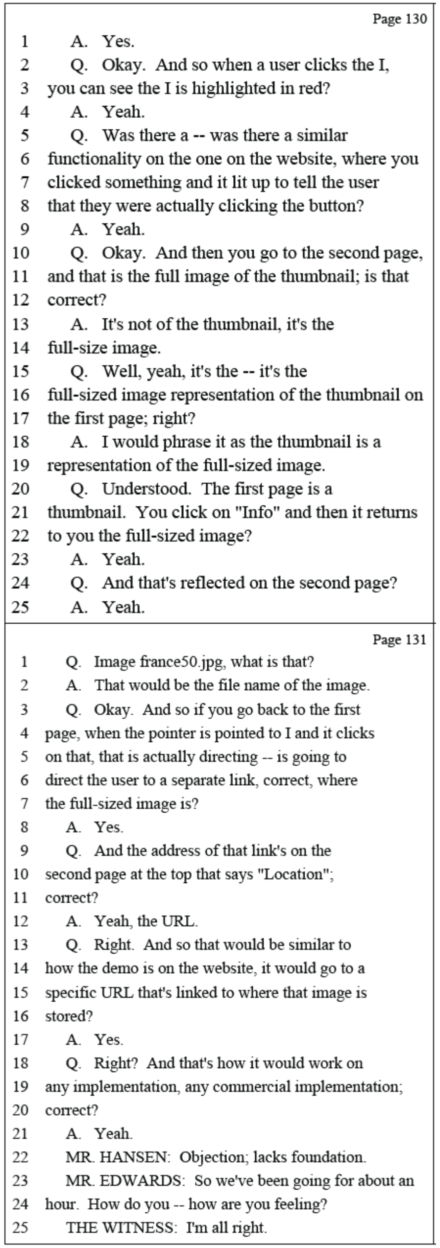

Page 130

A. Yes.
Q. Okay. And so when a user clicks the I, you can see the I is highlighted in red?
A. Yeah.
Q. Was there a -- was there a similar functionality on the one on the website, where you clicked something and it lit up to tell the user that they were actually clicking the button?
A. Yeah.
Q. Okay. And then you go to the second page, and that is the full image of the thumbnail; is that correct?
A. It's not of the thumbnail, it's the full-size image.
Q. Well, yeah, it's the -- it's the full-sized image representation of the thumbnail on the first page; right?
A. I would phrase it as the thumbnail is a representation of the full-sized image.
Q. Understood. The first page is a thumbnail. You click on "Info" and then it returns to you the full-sized image?
A. Yeah.
Q. And that's reflected on the second page?
A. Yeah.

Page 131

Q. Image france50.jpg, what is that?
A. That would be the file name of the image.
Q. Okay. And so if you go back to the first page, when the pointer is pointed to I and it clicks on that, that is actually directing -- is going to direct the user to a separate link, correct, where the full-sized image is?
A. Yes.
Q. And the address of that link's on the second page at the top that says "Location"; correct?
A. Yeah, the URL.
Q. Right. And so that would be similar to how the demo is on the website, it would go to a specific URL that's linked to where that image is stored?
A. Yes.
Q. Right? And that's how it would work on any implementation, any commercial implementation; correct?
A. Yeah.
MR. HANSEN: Objection; lacks foundation.
MR. EDWARDS: So we've been going for about an hour. How do you -- how are you feeling?
THE WITNESS: I'm all right.

Element	'897 Claim Recitation	Exemplary Citations in Reference
		Id. at 197–198 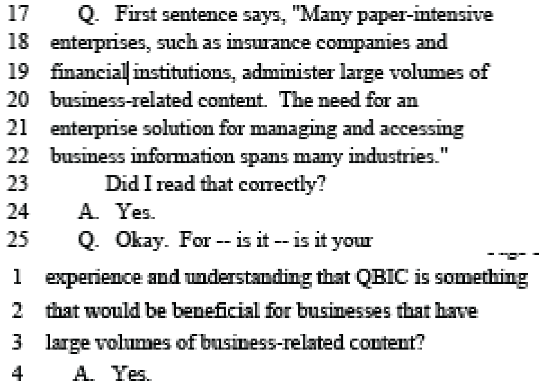 Q. First sentence says, "Many paper-intensive enterprises, such as insurance companies and financial institutions, administer large volumes of business-related content. The need for an enterprise solution for managing and accessing business information spans many industries." Did I read that correctly? A. Yes. Q. Okay. For -- is it -- is it your experience and understanding that QBIC is something that would be beneficial for businesses that have large volumes of business-related content? A. Yes. *Id.* at 307 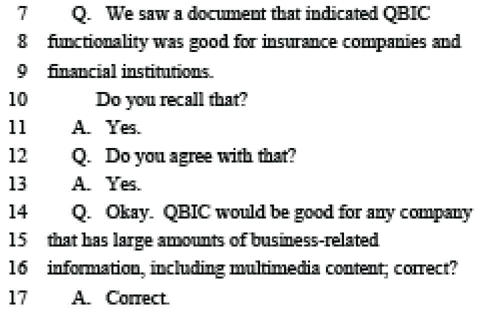 Q. We saw a document that indicated QBIC functionality was good for insurance companies and financial institutions. Do you recall that? A. Yes. Q. Do you agree with that? A. Yes. Q. Okay. QBIC would be good for any company that has large amounts of business-related information, including multimedia content; correct? A. Correct. To the extent it is found that QBIC does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that QBIC does not anticipate this claim, QBIC renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been

Element	'897 Claim Recitation	Exemplary Citations in Reference
		motivated to combine QBIC with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results. For example, a person of ordinary skill in the art would have been motivated to improve, modify, or combine QBIC with prior art reference teachings to arrive at this limitation at least because QBIC was designed for applications that included large amounts of business-related information, including multi-media content, including for financial institutions. *See, e.g.*, IBM 000789; Flickner Depo Rough Tr. at 191:16-192:8, 293:14-294:8.

2. Claim 30

Claim	'897 Claim Recitation	Exemplary Citations in Reference
30	The method of claim 25, wherein the step of displaying the displayed image includes displaying a time varying image.	QBIC discloses the method of claim 25, wherein the step of displaying the displayed image includes displaying a time varying image. Query by Image and Video Content: The QBIC System (IBM 0000893–902): "We have developed the QBIC (Query by Image Content) system to explore content-based retrieval methods. QBIC allows queries on large image and video databases based on • example images, • user-constructed sketches and drawings, • selected color and texture patterns, • camera and object motion, and • other graphical information." *Id*. at pp. 1–3. "To achieve this functionality, QBIC has two main components: database population (the process of creating an image database) and database query. During the population, images and videos are processed to extract features describing their content-colors, textures, shapes, and camera and object motion-and the features are stored in a database. During the query, the user composes a query graphically. Features are generated from the graphical query and then input to a matching engine that finds images or videos from the database with similar features. Figure 2 shows the system architecture." *Id*. at p. 3. "For both population and query, the QBIC data model has • still images or scenes (full images) that contain objects (subsets of an image), and • video shots that consist of sets of contiguous frames and contain motion objects. For still images, the QBIC data model distinguishes between 'scenes' (or images) and 'objects.' A scene is an image or single representative frame of video. An object is a part of a scene-for example, the fox in Figure 3-or a moving entity in a video. For still image database population, features are extracted from images and objects and stored in a database as shown in the top left part of Figure 2.

		Videos are broken into clips called shots. Representative frames, or r-frames, are generated for each extracted shot. R-frames are treated as still images, and features are extracted and stored in the database. Further processing of shots generates motion objects-for example, a car moving across the screen. Queries are allowed on objects ('Find images with a red, round object'), scenes ('Find images that have approximately 30-percent red and IS-percent blue colors'), shots ('Find all shots panning from left to right'), or any combination ('Find images that have 30 percent red and contain a blue textured object')." *Id.* at p. 3. "In still image database population, the images are reduced to a standard-sized icon called a thumbnail and annotated with any available text information. Object identification is an optional but key part of this step. It lets users manually, semiautomatically, or fully automatically identify interesting regions-which we call objects-in the images. . . . For video data, database population has three major components: • shot detection, • representative frame creation for each shot, and • derivation of a layered representation of coherently moving structures/ objects. . . . Once the shot boundaries have been detected, each shot is represented using an r-frame. R-frames are used for several purposes. First, during database population, r-frames are treated as still images in which objects can be identified by using the previously described methods." *Id.* at pp. 3–6. Image, Audio, and Video Extenders Administration and Programming (IBM 000001–617): "Your applications can even search for images by content. Imagine an application that uses visual examples to search for images. With such an application, users could select an example image and have the application find other images that have colors or textures similar to those in the example. With DB2 extenders' Query by Image Content (QBIC) capability, you can create applications that search for images in this visual way." *Id.* at p. 4.

		"To make an image available for searching by content, you catalog the image. When you catalog an image, the Image Extender analyzes the image, by computing the feature values for the image, and stores the values in a QBIC catalog. When you search for an image by content, your query identifies one or more features for the search (such as average color), a source for each feature (such as an example image), and a target set of cataloged images. The Image Extender computes the feature value of the source and compares it to the cataloged feature values for the target images. It then computes a score that indicates how similar the feature values of the target images are to the source. You can have the Image Extender return the images whose features are most similar to the source. The Image Extender will return the handle of each image and the image score. You can also have the Image Extender return only the score of a single image." *Id.* at pp. 20–21.

		Flickner Depo at 214-215 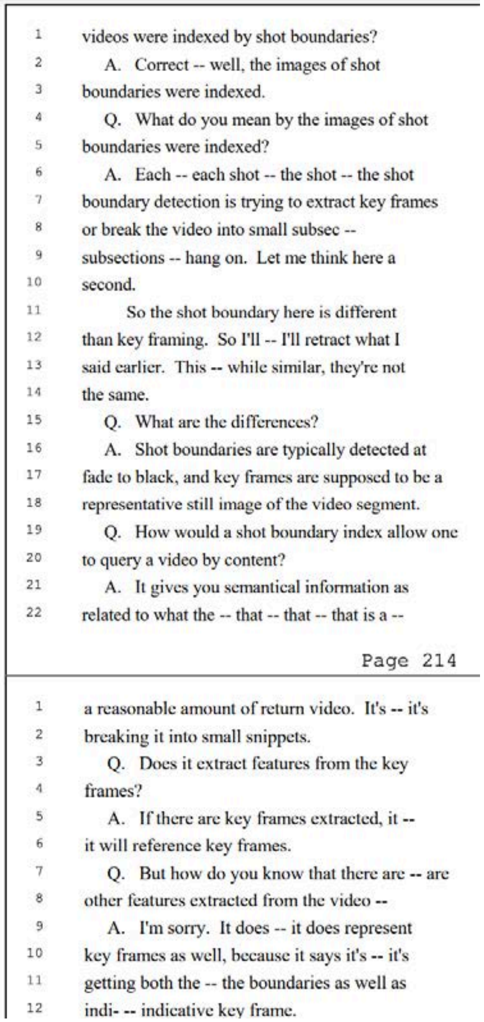 videos were indexed by shot boundaries? A. Correct -- well, the images of shot boundaries were indexed. Q. What do you mean by the images of shot boundaries were indexed? A. Each -- each shot -- the shot -- the shot boundary detection is trying to extract key frames or break the video into small subsec -- subsections -- hang on. Let me think here a second. So the shot boundary here is different than key framing. So I'll -- I'll retract what I said earlier. This -- while similar, they're not the same. Q. What are the differences? A. Shot boundaries are typically detected at fade to black, and key frames are supposed to be a representative still image of the video segment. Q. How would a shot boundary index allow one to query a video by content? A. It gives you semantical information as related to what the -- that -- that -- that is a -- Page 214 a reasonable amount of return video. It's -- it's breaking it into small snippets. Q. Does it extract features from the key frames? A. If there are key frames extracted, it -- it will reference key frames. Q. But how do you know that there are -- are other features extracted from the video -- A. I'm sorry. It does -- it does represent key frames as well, because it says it's -- it's getting both the -- the boundaries as well as indi- -- indicative key frame.

		Flickner 3/29/23 Depo at 32–33 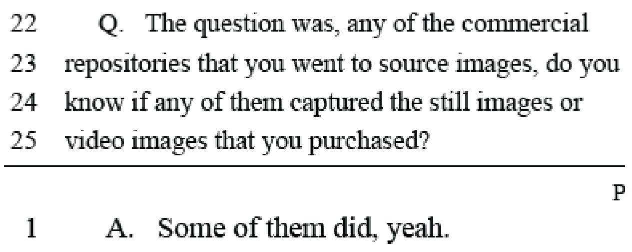 Q. The question was, any of the commercial repositories that you went to source images, do you know if any of them captured the still images or video images that you purchased? P A. Some of them did, yeah. *Id.* at 34 – 35

		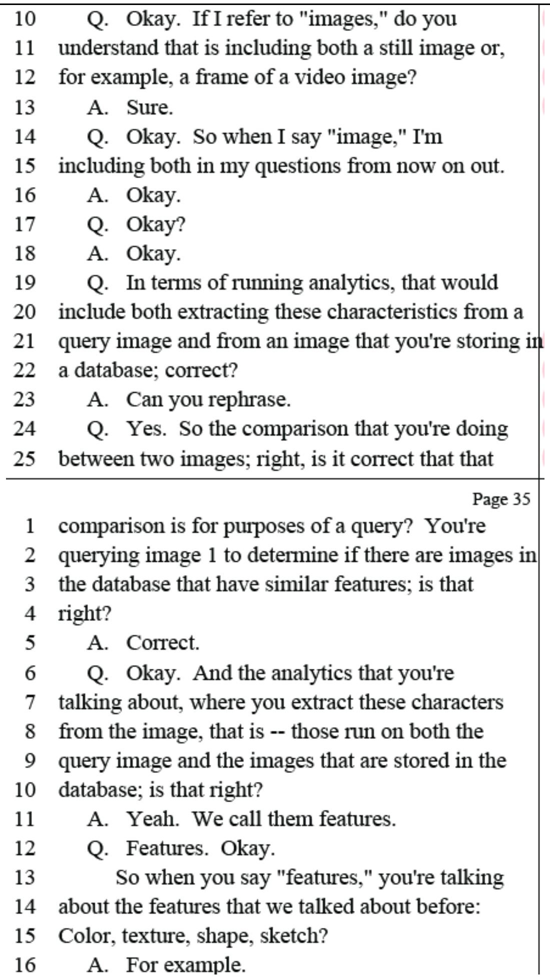 Q. Okay. If I refer to "images," do you understand that is including both a still image or, for example, a frame of a video image? A. Sure. Q. Okay. So when I say "image," I'm including both in my questions from now on out. A. Okay. Q. Okay? A. Okay. Q. In terms of running analytics, that would include both extracting these characteristics from a query image and from an image that you're storing in a database; correct? A. Can you rephrase. Q. Yes. So the comparison that you're doing between two images; right, is it correct that that Page 35 comparison is for purposes of a query? You're querying image 1 to determine if there are images in the database that have similar features; is that right? A. Correct. Q. Okay. And the analytics that you're talking about, where you extract these characters from the image, that is -- those run on both the query image and the images that are stored in the database; is that right? A. Yeah. We call them features. Q. Features. Okay. So when you say "features," you're talking about the features that we talked about before: Color, texture, shape, sketch? A. For example. *Id.* at 68 – 72

		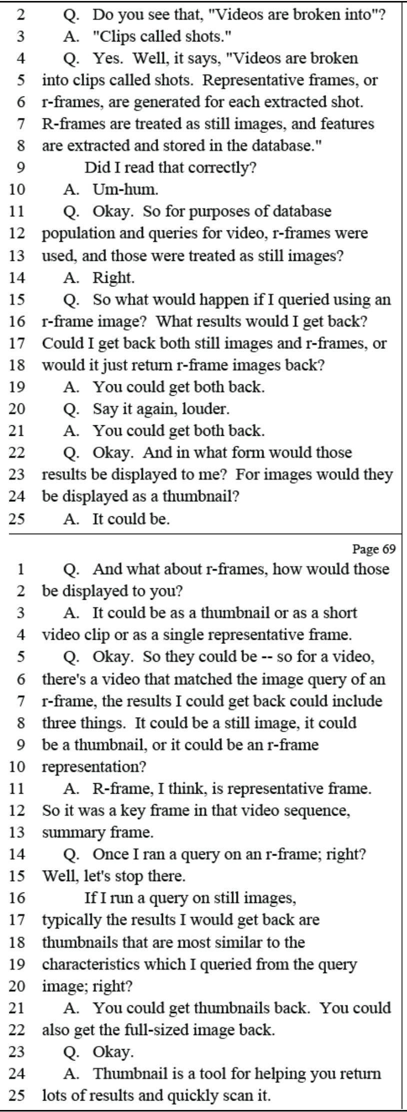 Q. Do you see that, "Videos are broken into"? A. "Clips called shots." Q. Yes. Well, it says, "Videos are broken into clips called shots. Representative frames, or r-frames, are generated for each extracted shot. R-frames are treated as still images, and features are extracted and stored in the database." Did I read that correctly? A. Um-hum. Q. Okay. So for purposes of database population and queries for video, r-frames were used, and those were treated as still images? A. Right. Q. So what would happen if I queried using an r-frame image? What results would I get back? Could I get back both still images and r-frames, or would it just return r-frame images back? A. You could get both back. Q. Say it again, louder. A. You could get both back. Q. Okay. And in what form would those results be displayed to me? For images would they be displayed as a thumbnail? A. It could be. Page 69 Q. And what about r-frames, how would those be displayed to you? A. It could be as a thumbnail or as a short video clip or as a single representative frame. Q. Okay. So they could be -- so for a video, there's a video that matched the image query of an r-frame, the results I could get back could include three things. It could be a still image, it could be a thumbnail, or it could be an r-frame representation? A. R-frame, I think, is representative frame. So it was a key frame in that video sequence, summary frame. Q. Once I ran a query on an r-frame; right? Well, let's stop there. If I run a query on still images, typically the results I would get back are thumbnails that are most similar to the characteristics which I queried from the query image; right? A. You could get thumbnails back. You could also get the full-sized image back. Q. Okay. A. Thumbnail is a tool for helping you return lots of results and quickly scan it.	

		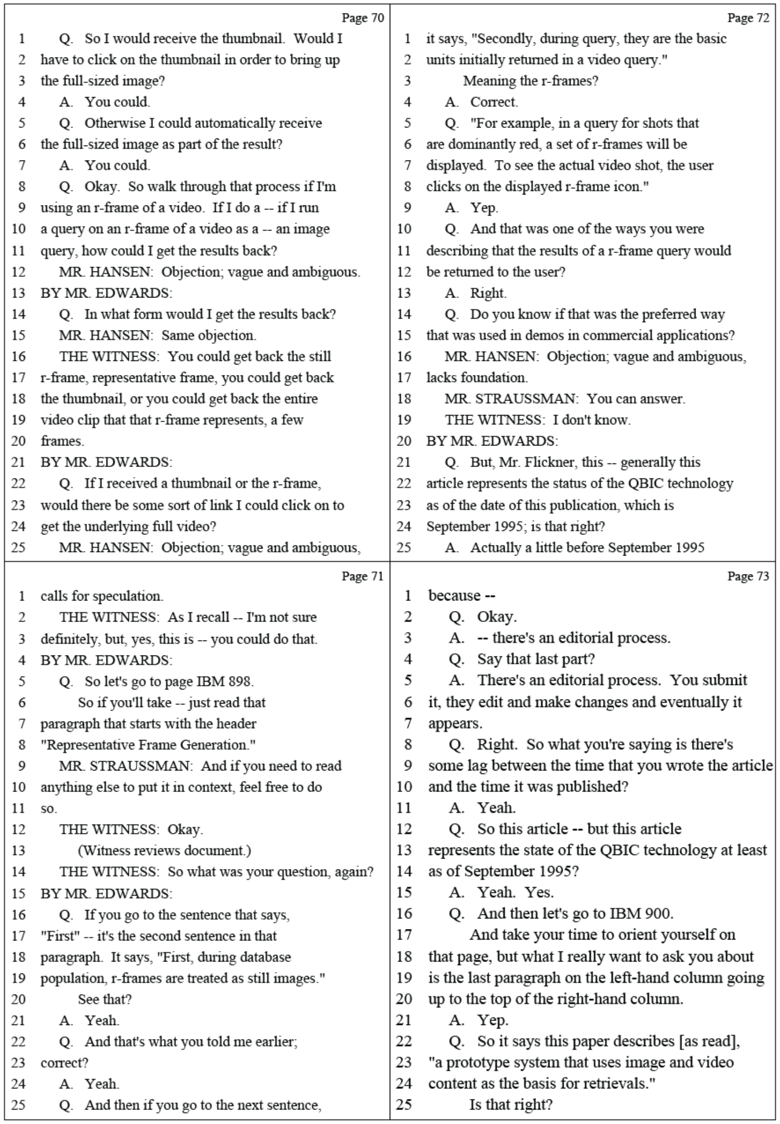 Page 70 Q. So I would receive the thumbnail. Would I have to click on the thumbnail in order to bring up the full-sized image? A. You could. Q. Otherwise I could automatically receive the full-sized image as part of the result? A. You could. Q. Okay. So walk through that process if I'm using an r-frame of a video. If I do a -- if I run a query on an r-frame of a video as a -- an image query, how could I get the results back? MR. HANSEN: Objection; vague and ambiguous. BY MR. EDWARDS: Q. In what form would I get the results back? MR. HANSEN: Same objection. THE WITNESS: You could get back the still r-frame, representative frame, you could get back the thumbnail, or you could get back the entire video clip that that r-frame represents, a few frames. BY MR. EDWARDS: Q. If I received a thumbnail or the r-frame, would there be some sort of link I could click on to get the underlying full video? MR. HANSEN: Objection; vague and ambiguous, Page 71 calls for speculation. THE WITNESS: As I recall -- I'm not sure definitely, but, yes, this is -- you could do that. BY MR. EDWARDS: Q. So let's go to page IBM 898. So if you'll take -- just read that paragraph that starts with the header "Representative Frame Generation." MR. STRAUSSMAN: And if you need to read anything else to put it in context, feel free to do so. THE WITNESS: Okay. (Witness reviews document.) THE WITNESS: So what was your question, again? BY MR. EDWARDS: Q. If you go to the sentence that says, "First" -- it's the second sentence in that paragraph. It says, "First, during database population, r-frames are treated as still images." See that? A. Yeah. Q. And that's what you told me earlier; correct? A. Yeah. Q. And then if you go to the next sentence, Page 72 it says, "Secondly, during query, they are the basic units initially returned in a video query." Meaning the r-frames? A. Correct. Q. "For example, in a query for shots that are dominantly red, a set of r-frames will be displayed. To see the actual video shot, the user clicks on the displayed r-frame icon." A. Yep. Q. And that was one of the ways you were describing that the results of a r-frame query would be returned to the user? A. Right. Q. Do you know if that was the preferred way that was used in demos in commercial applications? MR. HANSEN: Objection; vague and ambiguous, lacks foundation. MR. STRAUSSMAN: You can answer. THE WITNESS: I don't know. BY MR. EDWARDS: Q. But, Mr. Flickner, this -- generally this article represents the status of the QBIC technology as of the date of this publication, which is September 1995; is that right? A. Actually a little before September 1995 Page 73 because -- Q. Okay. A. -- there's an editorial process. Q. Say that last part? A. There's an editorial process. You submit it, they edit and make changes and eventually it appears. Q. Right. So what you're saying is there's some lag between the time that you wrote the article and the time it was published? A. Yeah. Q. So this article -- but this article represents the state of the QBIC technology at least as of September 1995? A. Yeah. Yes. Q. And then let's go to IBM 900. And take your time to orient yourself on that page, but what I really want to ask you about is the last paragraph on the left-hand column going up to the top of the right-hand column. A. Yep. Q. So it says this paper describes [as read], "a prototype system that uses image and video content as the basis for retrievals." Is that right? *Id.* at 75–76

		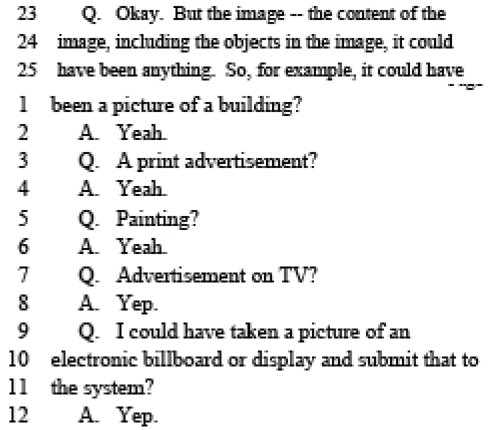 Q. Okay. But the image -- the content of the image, including the objects in the image, it could have been anything. So, for example, it could have been a picture of a building? A. Yeah. Q. A print advertisement? A. Yeah. Q. Painting? A. Yeah. Q. Advertisement on TV? A. Yep. Q. I could have taken a picture of an electronic billboard or display and submit that to the system? A. Yep. *See also* Flickner Depo at 38-41, 201, 209-210. To the extent it is found that QBIC does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that QBIC does not anticipate this claim, QBIC renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine QBIC with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

3. Claim 33

Claim	'897 Claim Recitation	Exemplary Citations in Reference
33	The method of claim 25, wherein the step of enabling performance of the transaction includes initiating a verification.	QBIC discloses the method of claim 25, wherein the step of enabling performance of the transaction includes initiating a verification. Query by Image and Video Content: The QBIC System (IBM 0000893–902): "INTEGRATION WITH TEXT AND PARAMETRIC ANNOTATION. Query by visual content complements and extends existing query methods. Systems must be able to integrate queries combining date, subject matter, price, and availability with content properties such as color, texture, and shape." *Id.* at p. 8. Image, Audio, and Video Extenders Administration and Programming (IBM 000001–617): "DB2 (DB2) Universal Database (UDB) is a powerful, object-relational database manager. It stores and protects traditional numeric and character data, as well as large, complex objects (LOBs). DB2 extenders help you exploit DB2 's object-relational features. The extenders define distinct data types and special functions for image, audio, video, and text objects. … Because of that, the extenders give your applications a single point of access to any or all of these types of data, along with traditional numeric and character data." *Id.* at p. 3. "A trigger defines a set of actions that are activated by a change to a table. Triggers can be used to perform actions such as validating input data, automatically generating a value for a newly inserted row, reading from other tables for cross-referencing purposes, or writing to other tables for auditing purposes. Triggers are often used for integrity checking or to enforce business rules." *Id.* at 16–17. Managing Enterprise Information Portal (IBM 000777–880):

"Many paper-intensive enterprises, such as insurance companies and financial institutions, administer large volumes of business-related content. The need for an enterprise solution for managing and accessing business information spans many industries. A content server is a software system that stores multimedia, business forms, documents, and related data, along with metadata that allows employees to process and work with the content." *Id*. at p. 3.

"The image search server uses IBM's QBIC (query by image content) technology to help you search for objects by certain visual properties, such as color and texture. The image search server analyzes images and stores the image information in a database. Then users can run image queries, which use the visual properties of images, to match colors, textures, and their positions without describing them in words. You can combine content-based queries with text and keyword searches for more powerful retrieval of image and multimedia data." *Id*. at p. 63.

QBIC Demo (IBM000747)

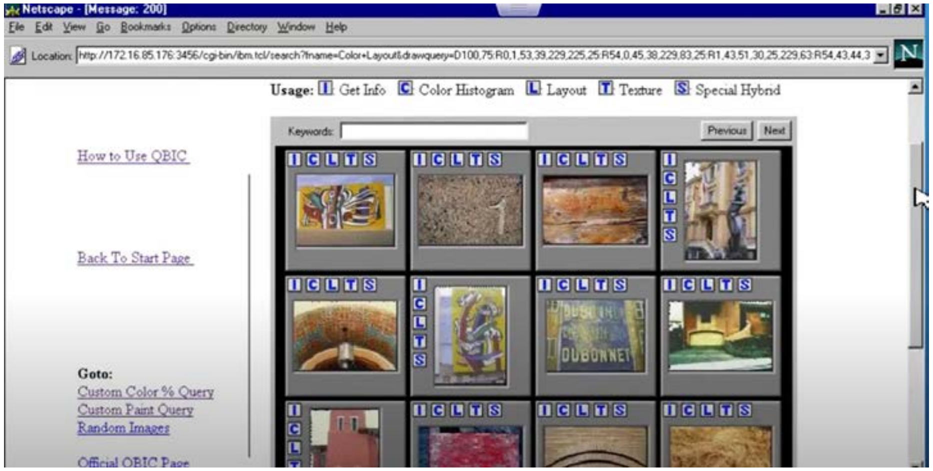

Location: http://172.16.85.176:3456/cgi-bin/ibm.tcl/search?fname=Color+Layout&drawquery=D100,75:R0,1,53,39,229,225,25:R54,0,45,38,229,83,25:R1,43,51,30,25,229,63:R54,43,44,3

		Flickner Depo at 131-133 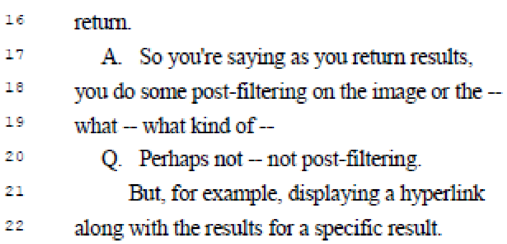 Q. An action that the QBIC system would automatically take based on the results it would return. A. So you're saying as you return results, you do some post-filtering on the image or the -- what -- what kind of -- Q. Perhaps not -- not post-filtering. But, for example, displaying a hyperlink along with the results for a specific result. 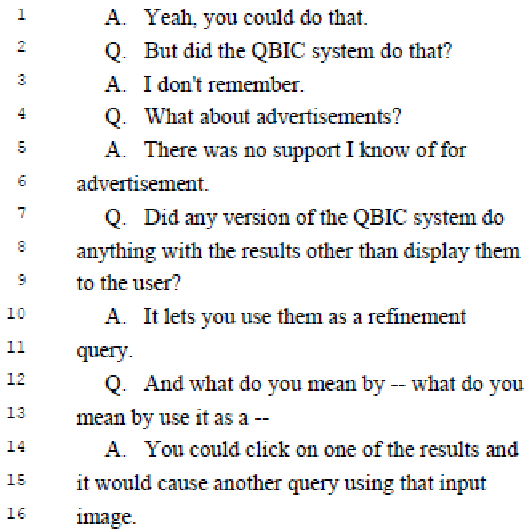 A. Yeah, you could do that. Q. But did the QBIC system do that? A. I don't remember. Q. What about advertisements? A. There was no support I know of for advertisement. Q. Did any version of the QBIC system do anything with the results other than display them to the user? A. It lets you use them as a refinement query. Q. And what do you mean by -- what do you mean by use it as a -- A. You could click on one of the results and it would cause another query using that input image.

		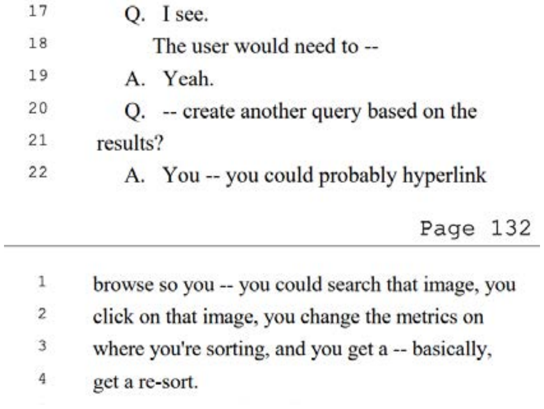 Q. I see. The user would need to -- A. Yeah. Q. -- create another query based on the results? A. You -- you could probably hyperlink Page 132 browse so you -- you could search that image, you click on that image, you change the metrics on where you're sorting, and you get a -- basically, get a re-sort. Flickner 8/29/23 Depo at 53–56 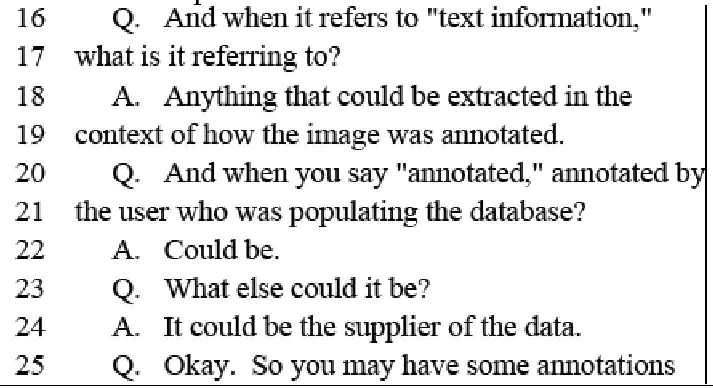 Q. And when it refers to "text information," what is it referring to? A. Anything that could be extracted in the context of how the image was annotated. Q. And when you say "annotated," annotated by the user who was populating the database? A. Could be. Q. What else could it be? A. It could be the supplier of the data. Q. Okay. So you may have some annotations

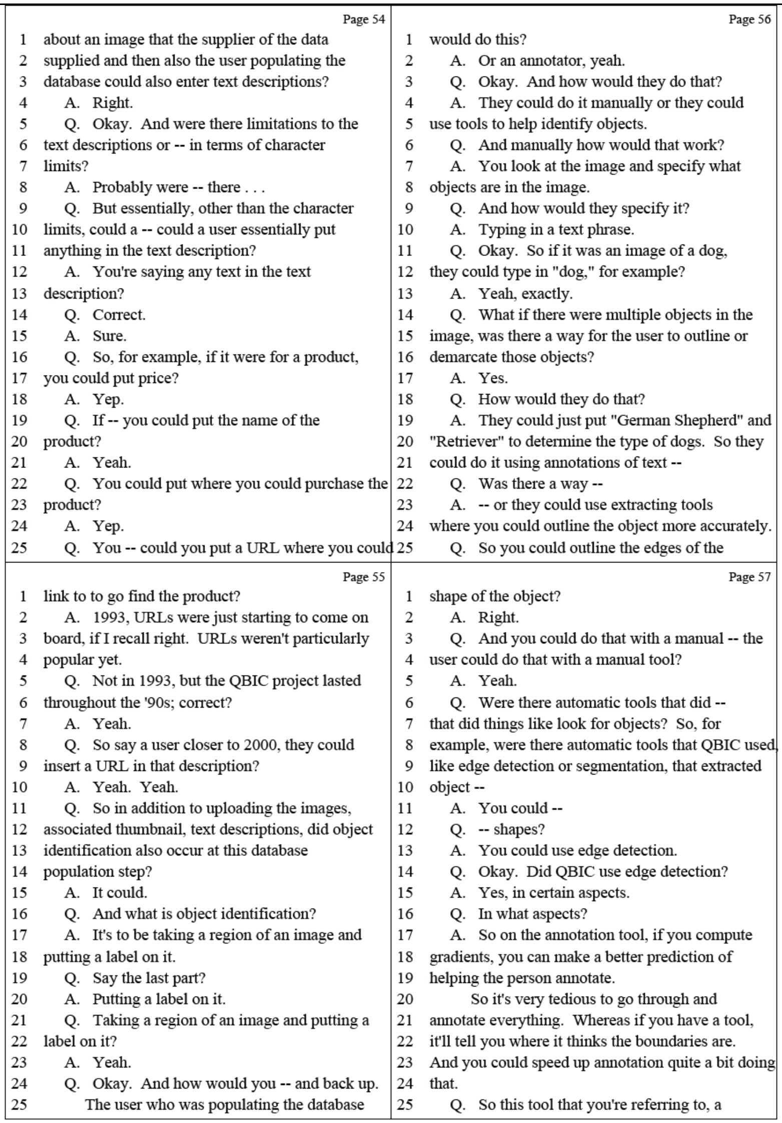

Page 54

about an image that the supplier of the data supplied and then also the user populating the database could also enter text descriptions?
A. Right.
Q. Okay. And were there limitations to the text descriptions or -- in terms of character limits?
A. Probably were -- there . . .
Q. But essentially, other than the character limits, could a -- could a user essentially put anything in the text description?
A. You're saying any text in the text description?
Q. Correct.
A. Sure.
Q. So, for example, if it were for a product, you could put price?
A. Yep.
Q. If -- you could put the name of the product?
A. Yeah.
Q. You could put where you could purchase the product?
A. Yep.
Q. You -- could you put a URL where you could

Page 55

link to to go find the product?
A. 1993, URLs were just starting to come on board, if I recall right. URLs weren't particularly popular yet.
Q. Not in 1993, but the QBIC project lasted throughout the '90s; correct?
A. Yeah.
Q. So say a user closer to 2000, they could insert a URL in that description?
A. Yeah. Yeah.
Q. So in addition to uploading the images, associated thumbnail, text descriptions, did object identification also occur at this database population step?
A. It could.
Q. And what is object identification?
A. It's to be taking a region of an image and putting a label on it.
Q. Say the last part?
A. Putting a label on it.
Q. Taking a region of an image and putting a label on it?
A. Yeah.
Q. Okay. And how would you -- and back up.
The user who was populating the database

Page 56

would do this?
A. Or an annotator, yeah.
Q. Okay. And how would they do that?
A. They could do it manually or they could use tools to help identify objects.
Q. And manually how would that work?
A. You look at the image and specify what objects are in the image.
Q. And how would they specify it?
A. Typing in a text phrase.
Q. Okay. So if it was an image of a dog, they could type in "dog," for example?
A. Yeah, exactly.
Q. What if there were multiple objects in the image, was there a way for the user to outline or demarcate those objects?
A. Yes.
Q. How would they do that?
A. They could just put "German Shepherd" and "Retriever" to determine the type of dogs. So they could do it using annotations of text --
Q. Was there a way --
A. -- or they could use extracting tools where you could outline the object more accurately.
Q. So you could outline the edges of the

Page 57

shape of the object?
A. Right.
Q. And you could do that with a manual -- the user could do that with a manual tool?
A. Yeah.
Q. Were there automatic tools that did -- that did things like look for objects? So, for example, were there automatic tools that QBIC used, like edge detection or segmentation, that extracted object --
A. You could --
Q. -- shapes?
A. You could use edge detection.
Q. Okay. Did QBIC use edge detection?
A. Yes, in certain aspects.
Q. In what aspects?
A. So on the annotation tool, if you compute gradients, you can make a better prediction of helping the person annotate.
So it's very tedious to go through and annotate everything. Whereas if you have a tool, it'll tell you where it thinks the boundaries are. And you could speed up annotation quite a bit doing that.
Q. So this tool that you're referring to, a

Id. at 68–72

		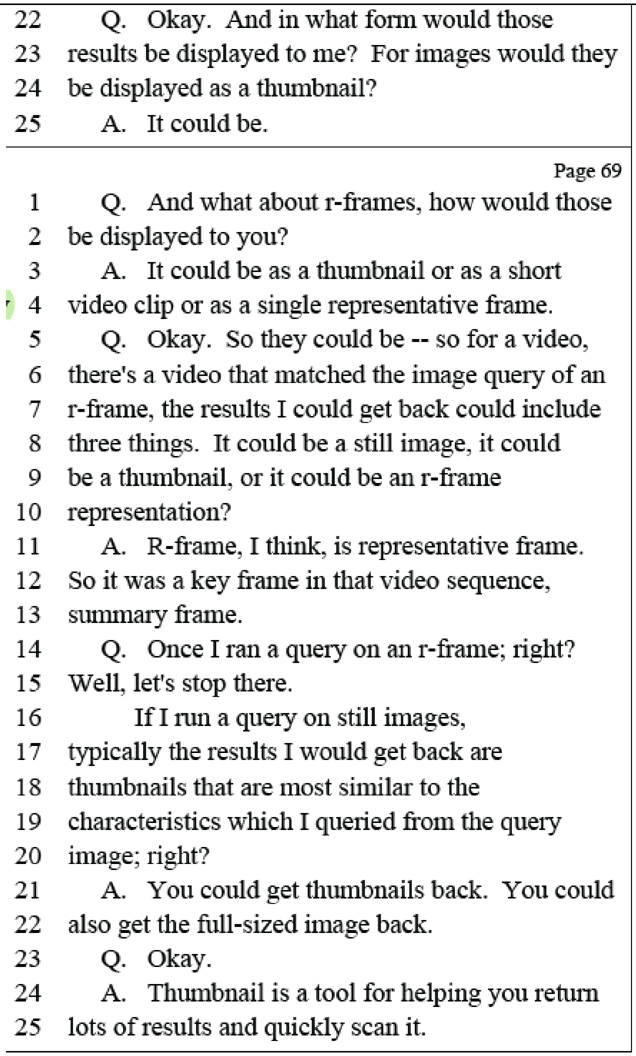 Q. Okay. And in what form would those results be displayed to me? For images would they be displayed as a thumbnail? A. It could be. Page 69 Q. And what about r-frames, how would those be displayed to you? A. It could be as a thumbnail or as a short video clip or as a single representative frame. Q. Okay. So they could be -- so for a video, there's a video that matched the image query of an r-frame, the results I could get back could include three things. It could be a still image, it could be a thumbnail, or it could be an r-frame representation? A. R-frame, I think, is representative frame. So it was a key frame in that video sequence, summary frame. Q. Once I ran a query on an r-frame; right? Well, let's stop there. If I run a query on still images, typically the results I would get back are thumbnails that are most similar to the characteristics which I queried from the query image; right? A. You could get thumbnails back. You could also get the full-sized image back. Q. Okay. A. Thumbnail is a tool for helping you return lots of results and quickly scan it.	

		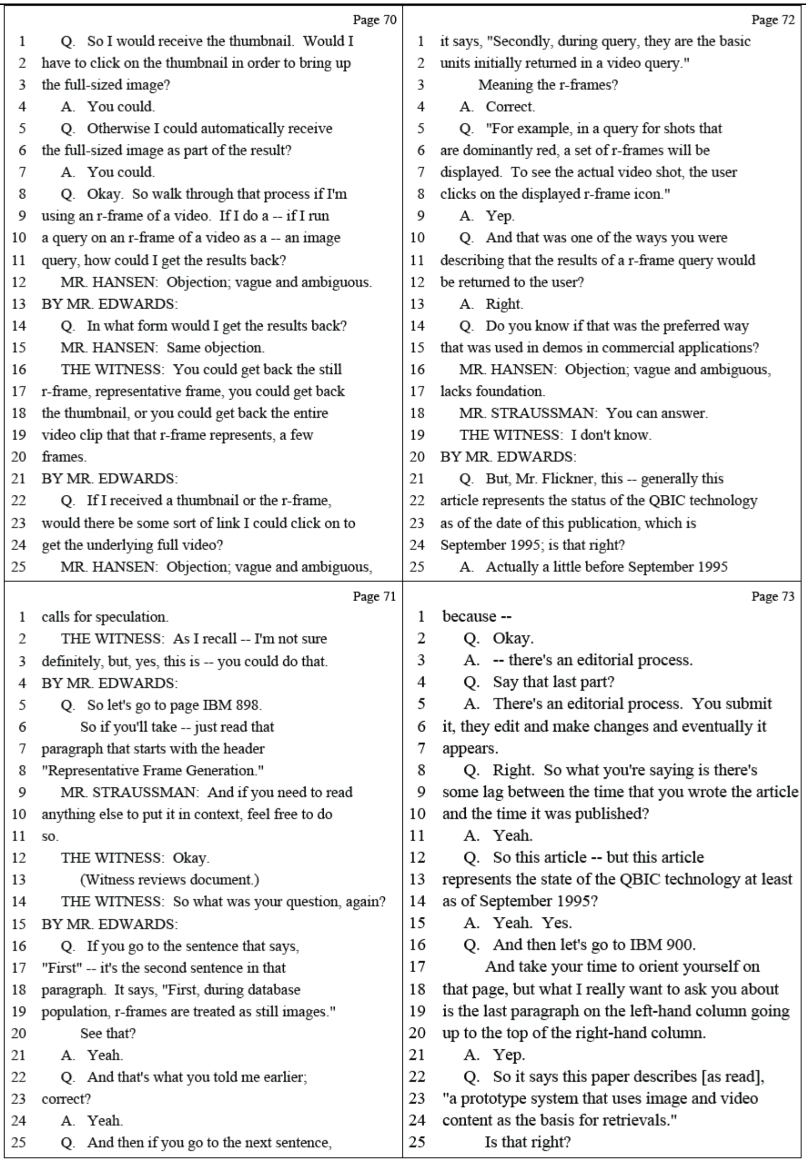 Page 70 Q. So I would receive the thumbnail. Would I have to click on the thumbnail in order to bring up the full-sized image? A. You could. Q. Otherwise I could automatically receive the full-sized image as part of the result? A. You could. Q. Okay. So walk through that process if I'm using an r-frame of a video. If I do a -- if I run a query on an r-frame of a video as a -- an image query, how could I get the results back? MR. HANSEN: Objection; vague and ambiguous. BY MR. EDWARDS: Q. In what form would I get the results back? MR. HANSEN: Same objection. THE WITNESS: You could get back the still r-frame, representative frame, you could get back the thumbnail, or you could get back the entire video clip that that r-frame represents, a few frames. BY MR. EDWARDS: Q. If I received a thumbnail or the r-frame, would there be some sort of link I could click on to get the underlying full video? MR. HANSEN: Objection; vague and ambiguous, Page 71 calls for speculation. THE WITNESS: As I recall -- I'm not sure definitely, but, yes, this is -- you could do that. BY MR. EDWARDS: Q. So let's go to page IBM 898. So if you'll take -- just read that paragraph that starts with the header "Representative Frame Generation." MR. STRAUSSMAN: And if you need to read anything else to put it in context, feel free to do so. THE WITNESS: Okay. (Witness reviews document.) THE WITNESS: So what was your question, again? BY MR. EDWARDS: Q. If you go to the sentence that says, "First" -- it's the second sentence in that paragraph. It says, "First, during database population, r-frames are treated as still images." See that? A. Yeah. Q. And that's what you told me earlier; correct? A. Yeah. Q. And then if you go to the next sentence, Page 72 it says, "Secondly, during query, they are the basic units initially returned in a video query." Meaning the r-frames? A. Correct. Q. "For example, in a query for shots that are dominantly red, a set of r-frames will be displayed. To see the actual video shot, the user clicks on the displayed r-frame icon." A. Yep. Q. And that was one of the ways you were describing that the results of a r-frame query would be returned to the user? A. Right. Q. Do you know if that was the preferred way that was used in demos in commercial applications? MR. HANSEN: Objection; vague and ambiguous, lacks foundation. MR. STRAUSSMAN: You can answer. THE WITNESS: I don't know. BY MR. EDWARDS: Q. But, Mr. Flickner, this -- generally this article represents the status of the QBIC technology as of the date of this publication, which is September 1995; is that right? A. Actually a little before September 1995 Page 73 because -- Q. Okay. A. -- there's an editorial process. Q. Say that last part? A. There's an editorial process. You submit it, they edit and make changes and eventually it appears. Q. Right. So what you're saying is there's some lag between the time that you wrote the article and the time it was published? A. Yeah. Q. So this article -- but this article represents the state of the QBIC technology at least as of September 1995? A. Yeah. Yes. Q. And then let's go to IBM 900. And take your time to orient yourself on that page, but what I really want to ask you about is the last paragraph on the left-hand column going up to the top of the right-hand column. A. Yep. Q. So it says this paper describes [as read], "a prototype system that uses image and video content as the basis for retrievals." Is that right? *Id.* at 118–119

		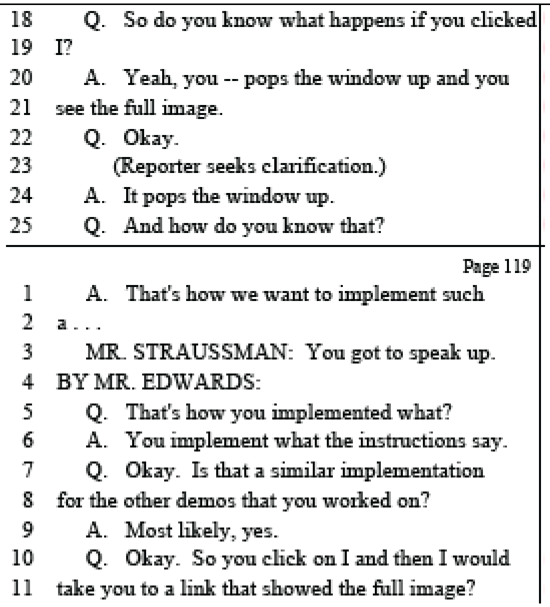 Q. So do you know what happens if you clicked I? A. Yeah, you -- pops the window up and you see the full image. Q. Okay. (Reporter seeks clarification.) A. It pops the window up. Q. And how do you know that? Page 119 A. That's how we want to implement such a . . . MR. STRAUSSMAN: You got to speak up. BY MR. EDWARDS: Q. That's how you implemented what? A. You implement what the instructions say. Q. Okay. Is that a similar implementation for the other demos that you worked on? A. Most likely, yes. Q. Okay. So you click on I and then I would take you to a link that showed the full image? *Id.* at 129 – 131

		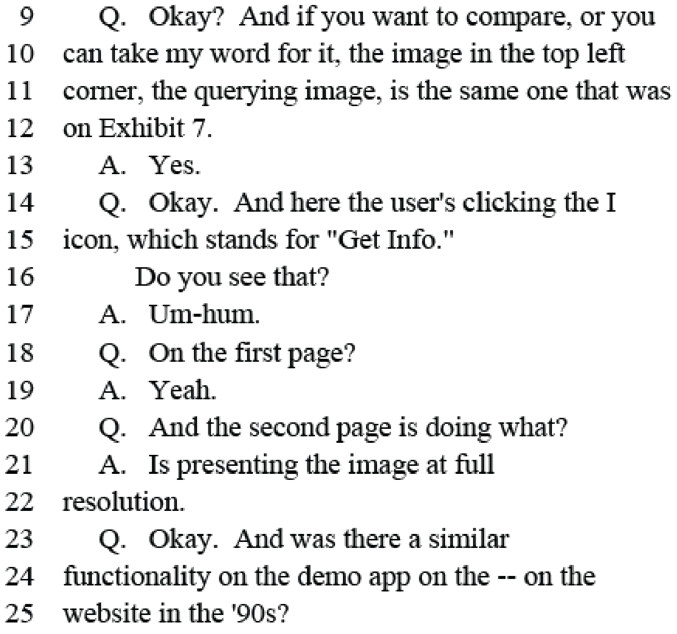 Q. Okay? And if you want to compare, or you can take my word for it, the image in the top left corner, the querying image, is the same one that was on Exhibit 7. A. Yes. Q. Okay. And here the user's clicking the I icon, which stands for "Get Info." Do you see that? A. Um-hum. Q. On the first page? A. Yeah. Q. And the second page is doing what? A. Is presenting the image at full resolution. Q. Okay. And was there a similar functionality on the demo app on the -- on the website in the '90s?

			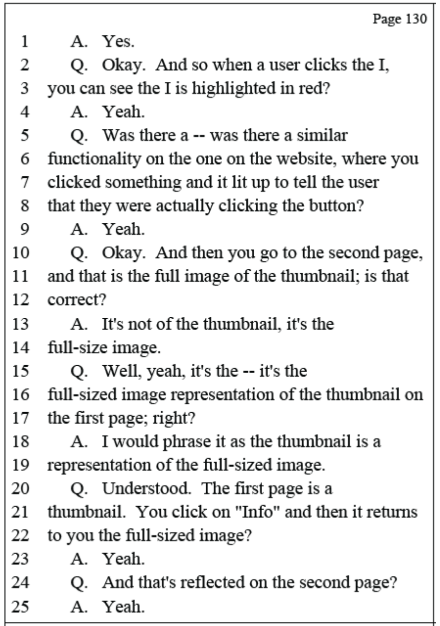 Page 130 A. Yes. Q. Okay. And so when a user clicks the I, you can see the I is highlighted in red? A. Yeah. Q. Was there a -- was there a similar functionality on the one on the website, where you clicked something and it lit up to tell the user that they were actually clicking the button? A. Yeah. Q. Okay. And then you go to the second page, and that is the full image of the thumbnail; is that correct? A. It's not of the thumbnail, it's the full-size image. Q. Well, yeah, it's the -- it's the full-sized image representation of the thumbnail on the first page; right? A. I would phrase it as the thumbnail is a representation of the full-sized image. Q. Understood. The first page is a thumbnail. You click on "Info" and then it returns to you the full-sized image? A. Yeah. Q. And that's reflected on the second page? A. Yeah. 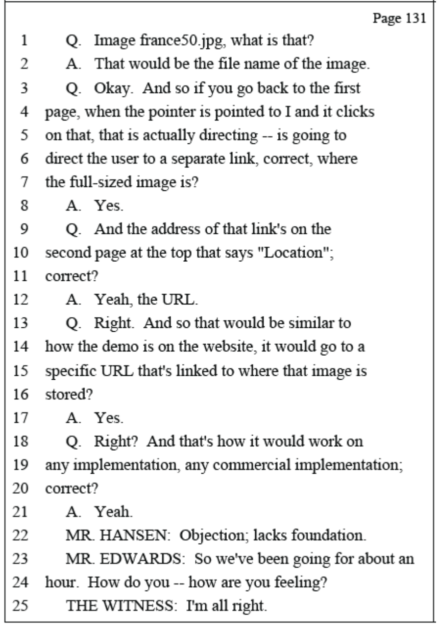 Page 131 Q. Image france50.jpg, what is that? A. That would be the file name of the image. Q. Okay. And so if you go back to the first page, when the pointer is pointed to I and it clicks on that, that is actually directing -- is going to direct the user to a separate link, correct, where the full-sized image is? A. Yes. Q. And the address of that link's on the second page at the top that says "Location"; correct? A. Yeah, the URL. Q. Right. And so that would be similar to how the demo is on the website, it would go to a specific URL that's linked to where that image is stored? A. Yes. Q. Right? And that's how it would work on any implementation, any commercial implementation; correct? A. Yeah. MR. HANSEN: Objection; lacks foundation. MR. EDWARDS: So we've been going for about an hour. How do you -- how are you feeling? THE WITNESS: I'm all right.

		Id. at 159–160 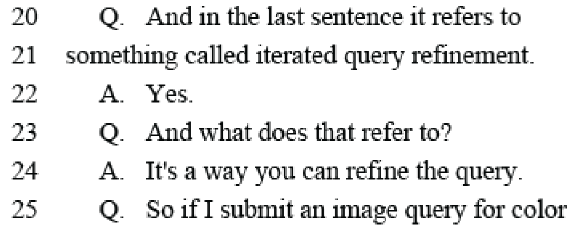 Q. And in the last sentence it refers to something called iterated query refinement. A. Yes. Q. And what does that refer to? A. It's a way you can refine the query. Q. So if I submit an image query for color 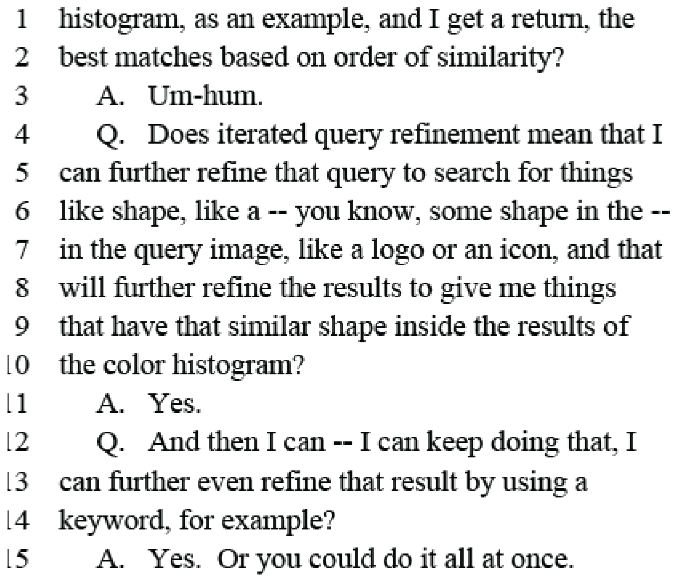 histogram, as an example, and I get a return, the best matches based on order of similarity? A. Um-hum. Q. Does iterated query refinement mean that I can further refine that query to search for things like shape, like a -- you know, some shape in the -- in the query image, like a logo or an icon, and that will further refine the results to give me things that have that similar shape inside the results of the color histogram? A. Yes. Q. And then I can -- I can keep doing that, I can further even refine that result by using a keyword, for example? A. Yes. Or you could do it all at once. *Id.* at 160–161

		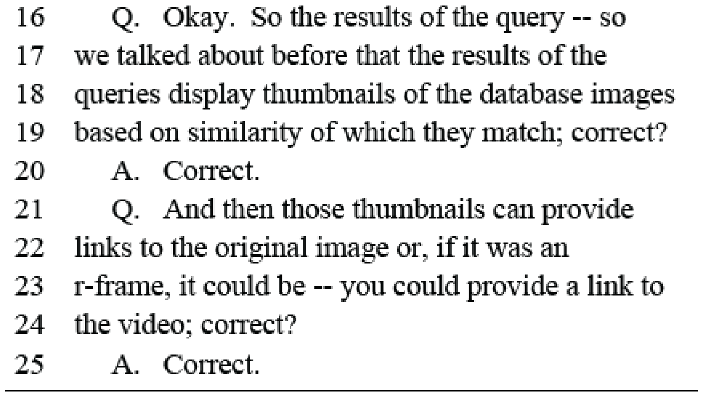 Q. Okay. So the results of the query -- so we talked about before that the results of the queries display thumbnails of the database images based on similarity of which they match; correct? A. Correct. Q. And then those thumbnails can provide links to the original image or, if it was an r-frame, it could be -- you could provide a link to the video; correct? A. Correct. 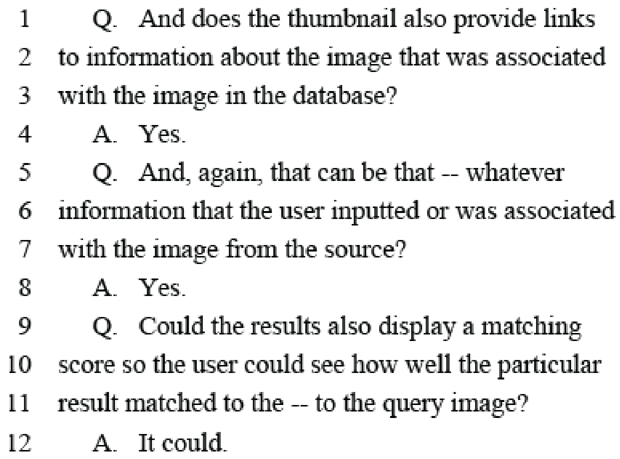 Q. And does the thumbnail also provide links to information about the image that was associated with the image in the database? A. Yes. Q. And, again, that can be that -- whatever information that the user inputted or was associated with the image from the source? A. Yes. Q. Could the results also display a matching score so the user could see how well the particular result matched to the -- to the query image? A. It could. *Id.* at 197–198 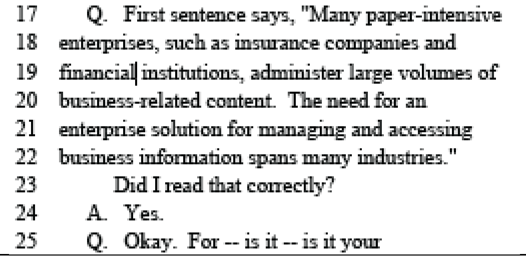 Q. First sentence says, "Many paper-intensive enterprises, such as insurance companies and financial institutions, administer large volumes of business-related content. The need for an enterprise solution for managing and accessing business information spans many industries." Did I read that correctly? A. Yes. Q. Okay. For -- is it -- is it your

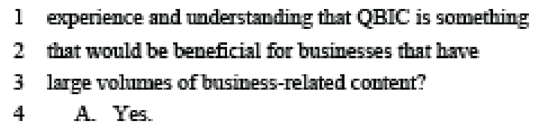

experience and understanding that QBIC is something that would be beneficial for businesses that have large volumes of business-related content?
A. Yes.

Id. at 307

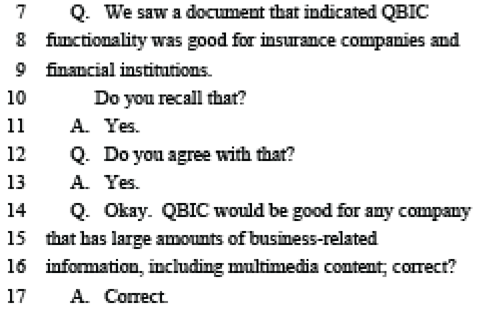

Q. We saw a document that indicated QBIC functionality was good for insurance companies and financial institutions.
Do you recall that?
A. Yes.
Q. Do you agree with that?
A. Yes.
Q. Okay. QBIC would be good for any company that has large amounts of business-related information, including multimedia content; correct?
A. Correct.

To the extent it is found that QBIC does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that QBIC does not anticipate this claim, QBIC renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine QBIC with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results. For example, a person of ordinary skill in the art would have been motivated to improve, modify, or combine QBIC with prior art reference teachings to arrive at this limitation at least because QBIC was designed for applications that included large amounts of business-related information, including multi-media content, including for financial

		institutions. *See, e.g.*, IBM 000789; Flickner Depo Rough Tr. at 191:16-192:8, 293:14-294:8.

4. Claim 34

Claim	'897 Claim Recitation	Exemplary Citations in Reference
34	The method of claim 25 wherein the step of enabling performance of the transaction includes initiating an authorization.	QBIC discloses the method of claim 25 wherein the step of enabling performance of the transaction includes initiating an authorization. Query by Image and Video Content: The QBIC System (IBM 0000893–902): "INTEGRATION WITH TEXT AND PARAMETRIC ANNOTATION. Query by visual content complements and extends existing query methods. Systems must be able to integrate queries combining date, subject matter, price, and availability with content properties such as color, texture, and shape." *Id.* at p. 8. Image, Audio, and Video Extenders Administration and Programming (IBM 000001–617): "DB2 (DB2) Universal Database (UDB) is a powerful, object-relational database manager. It stores and protects traditional numeric and character data, as well as large, complex objects (LOBs). DB2 extenders help you exploit DB2 's object-relational features. The extenders define distinct data types and special functions for image, audio, video, and text objects. … Because of that, the extenders give your applications a single point of access to any or all of these types of data, along with traditional numeric and character data." *Id.* at p. 3. "A trigger defines a set of actions that are activated by a change to a table. Triggers can be used to perform actions such as validating input data, automatically generating a value for a newly inserted row, reading from other tables for cross-referencing purposes, or writing to other tables for auditing purposes. Triggers are often used for integrity checking or to enforce business rules." *Id.* at 16–17. Managing Enterprise Information Portal (IBM 000777–880):

> "Many paper-intensive enterprises, such as insurance companies and financial institutions, administer large volumes of business-related content. The need for an enterprise solution for managing and accessing business information spans many industries. A content server is a software system that stores multimedia, business forms, documents, and related data, along with metadata that allows employees to process and work with the content." *Id*. at p. 3.
>
> "The image search server uses IBM's QBIC (query by image content) technology to help you search for objects by certain visual properties, such as color and texture. The image search server analyzes images and stores the image information in a database. Then users can run image queries, which use the visual properties of images, to match colors, textures, and their positions without describing them in words. You can combine content-based queries with text and keyword searches for more powerful retrieval of image and multimedia data." *Id*. at p. 63.

QBIC Demo (IBM000747)

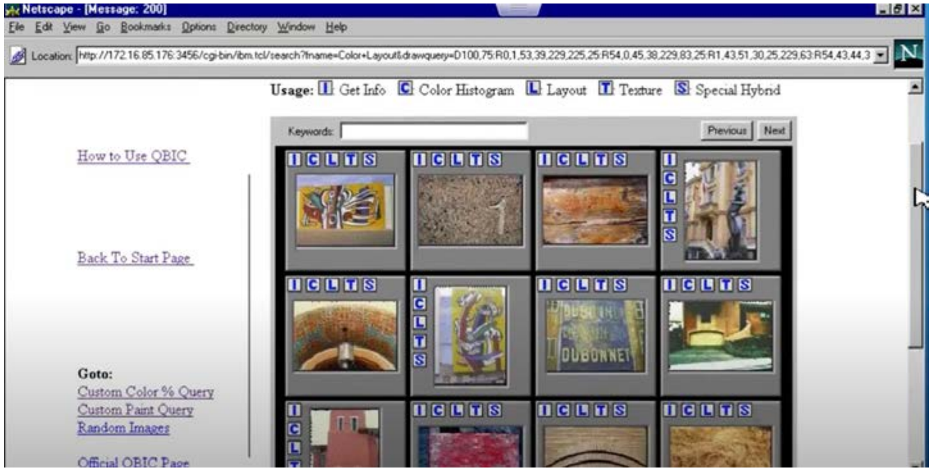

		Flickner Depo at 131-133 Flickner 8/31/23 Depo at 197–198 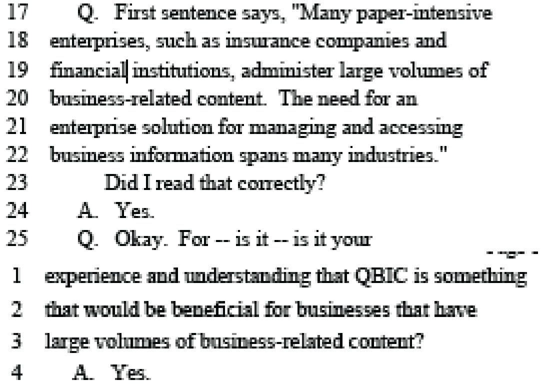 Q. First sentence says, "Many paper-intensive enterprises, such as insurance companies and financial institutions, administer large volumes of business-related content. The need for an enterprise solution for managing and accessing business information spans many industries." Did I read that correctly? A. Yes. Q. Okay. For -- is it -- is it your experience and understanding that QBIC is something that would be beneficial for businesses that have large volumes of business-related content? A. Yes. *Id.* at 307 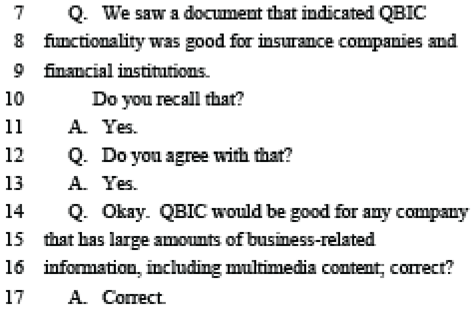 Q. We saw a document that indicated QBIC functionality was good for insurance companies and financial institutions. Do you recall that? A. Yes. Q. Do you agree with that? A. Yes. Q. Okay. QBIC would be good for any company that has large amounts of business-related information, including multimedia content; correct? A. Correct. To the extent it is found that QBIC does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that QBIC does not anticipate this claim, QBIC renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Invalidity Contentions. A person of ordinary skill in the art would have been

		motivated to combine QBIC with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results. For example, a person of ordinary skill in the art would have been motivated to improve, modify, or combine QBIC with prior art reference teachings to arrive at this limitation at least because QBIC was designed for applications that included large amounts of business-related information, including multi-media content, including for financial institutions. *See, e.g.*, IBM 000789; Flickner Depo Rough Tr. at 191:16-192:8, 293:14-294:8.

5. Claim 38

Claim	'897 Claim Recitation	Exemplary Citations in Reference
38	The method of claim 25, wherein the step of enabling performance of the interaction includes initiating a transaction with an account.	To the extent it is found that QBIC does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that QBIC does not anticipate this claim, QBIC renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's Preliminary Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine QBIC with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results.

6. Claim 39

Claim	'897 Claim Recitation	Exemplary Citations in Reference
39	The method of claim 25, wherein the step of enabling performance of the interaction includes initiating a transaction with an account.	QBIC discloses the method of claim 25, wherein the step of enabling performance of the interaction includes initiating a transaction with an account. Query by Image and Video Content: The QBIC System (IBM 0000893–902): "QBIC has two main components: database population (the process of creating an image database) and database query. During the population, images and videos are processed to extract features describing their content-colors, textures, shapes, and camera and object motion-and the features are stored in a database. During the query, the user composes a query graphically. Features are generated from the graphical query and then input to a matching engine that finds images or videos from the database with similar features. Figure 2 shows the system architecture." *Id.* at p. 3. 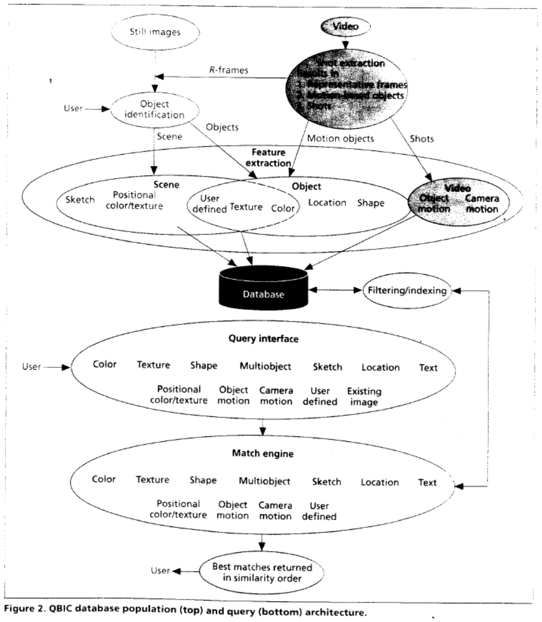 Figure 2. QBIC database population (top) and query (bottom) architecture. *Id.* at Fig. 2.

		"INTEGRATION WITH TEXT AND PARAMETRIC ANNOTATION. Query by visual content complements and extends existing query methods. Systems must be able to integrate queries combining date, subject matter, price, and availability with content properties such as color, texture, and shape." *Id.* at p. 8. Image, Audio, and Video Extenders Administration and Programming (IBM 000001–617): "To make an image available for searching by content, you catalog the image. When you catalog an image, the Image Extender analyzes the image, by computing the feature values for the image, and stores the values in a QBIC catalog. When you search for an image by content, your query identifies one or more features for the search (such as average color), a source for each feature (such as an example image), and a target set of cataloged images. The Image Extender computes the feature value of the source and compares it to the cataloged feature values for the target images. It then computes a score that indicates how similar the feature values of the target images are to the source. You can have the Image Extender return the images whose features are most similar to the source. The Image Extender will return the handle of each image and the image score. You can also have the Image Extender return only the score of a single image." *Id.* at pp. 20–21. Managing Enterprise Information Portal (IBM 000777–880): "Many paper-intensive enterprises, such as insurance companies and financial institutions, administer large volumes of business-related content. The need for an enterprise solution for managing and accessing business information spans many industries. A content server is a software system that stores multimedia, business forms, documents, and related data, along with metadata that allows employees to process and work with the content." *Id.* at p. 3.

		"The image search server uses IBM's QBIC (query by image content) technology to help you search for objects by certain visual properties, such as color and texture. The image search server analyzes images and stores the image information in a database. Then users can run image queries, which use the visual properties of images, to match colors, textures, and their positions without describing them in words. You can combine content-based queries with text and keyword searches for more powerful retrieval of image and multimedia data." *Id.* at p. 63. Flickner 8/31/23 Depo at 197–198 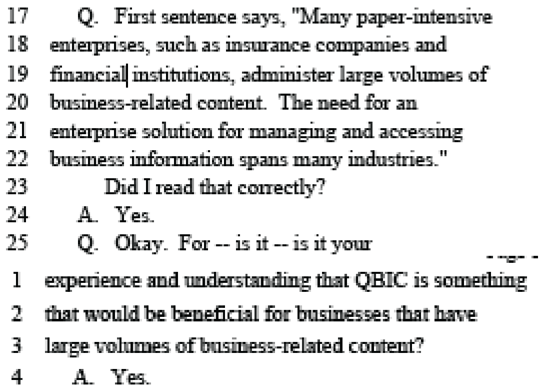 Q. First sentence says, "Many paper-intensive enterprises, such as insurance companies and financial institutions, administer large volumes of business-related content. The need for an enterprise solution for managing and accessing business information spans many industries." Did I read that correctly? A. Yes. Q. Okay. For -- is it -- is it your experience and understanding that QBIC is something that would be beneficial for businesses that have large volumes of business-related content? A. Yes. *Id.* at 307

		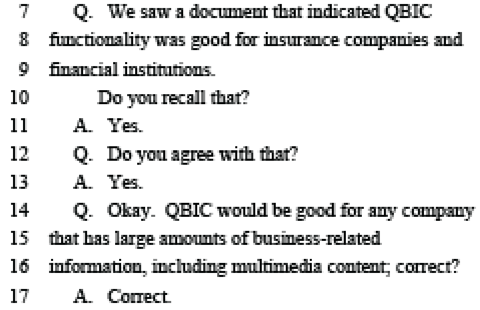 Q. We saw a document that indicated QBIC functionality was good for insurance companies and financial institutions. Do you recall that? A. Yes. Q. Do you agree with that? A. Yes. Q. Okay. QBIC would be good for any company that has large amounts of business-related information, including multimedia content; correct? A. Correct. IBM000001 at 39 – 41

		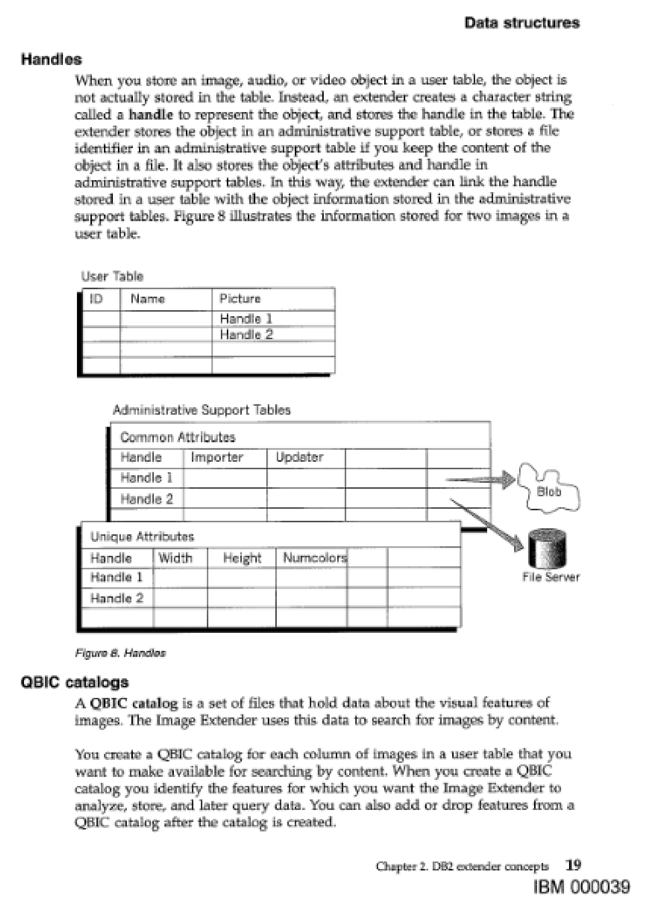 Data structures **Handles** When you store an image, audio, or video object in a user table, the object is not actually stored in the table. Instead, an extender creates a character string called a **handle** to represent the object, and stores the handle in the table. The extender stores the object in an administrative support table, or stores a file identifier in an administrative support table if you keep the content of the object in a file. It also stores the object's attributes and handle in administrative support tables. In this way, the extender can link the handle stored in a user table with the object information stored in the administrative support tables. Figure 8 illustrates the information stored for two images in a user table. *Figure 8. Handles* **QBIC catalogs** A **QBIC catalog** is a set of files that hold data about the visual features of images. The Image Extender uses this data to search for images by content. You create a QBIC catalog for each column of images in a user table that you want to make available for searching by content. When you create a QBIC catalog you identify the features for which you want the Image Extender to analyze, store, and later query data. You can also add or drop features from a QBIC catalog after the catalog is created. Chapter 2. DB2 extender concepts 19 IBM 000039	

		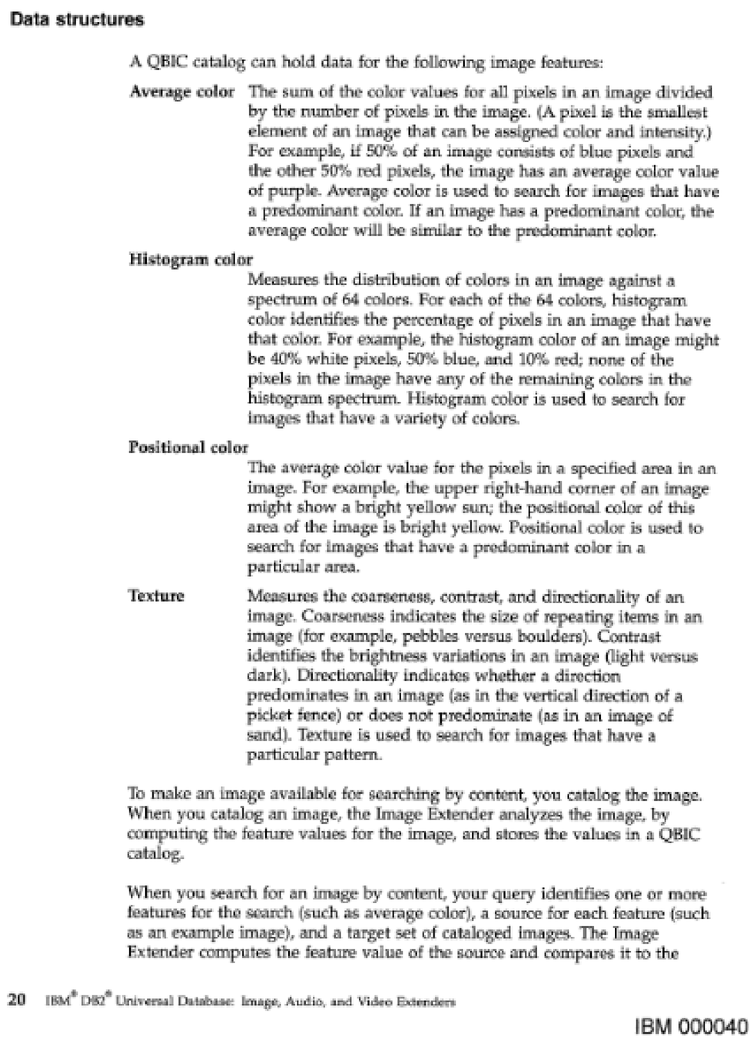 **Data structures** A QBIC catalog can hold data for the following image features: **Average color** The sum of the color values for all pixels in an image divided by the number of pixels in the image. (A pixel is the smallest element of an image that can be assigned color and intensity.) For example, if 50% of an image consists of blue pixels and the other 50% red pixels, the image has an average color value of purple. Average color is used to search for images that have a predominant color. If an image has a predominant color, the average color will be similar to the predominant color. **Histogram color** Measures the distribution of colors in an image against a spectrum of 64 colors. For each of the 64 colors, histogram color identifies the percentage of pixels in an image that have that color. For example, the histogram color of an image might be 40% white pixels, 50% blue, and 10% red; none of the pixels in the image have any of the remaining colors in the histogram spectrum. Histogram color is used to search for images that have a variety of colors. **Positional color** The average color value for the pixels in a specified area in an image. For example, the upper right-hand corner of an image might show a bright yellow sun; the positional color of this area of the image is bright yellow. Positional color is used to search for images that have a predominant color in a particular area. **Texture** Measures the coarseness, contrast, and directionality of an image. Coarseness indicates the size of repeating items in an image (for example, pebbles versus boulders). Contrast identifies the brightness variations in an image (light versus dark). Directionality indicates whether a direction predominates in an image (as in the vertical direction of a picket fence) or does not predominate (as in an image of sand). Texture is used to search for images that have a particular pattern. To make an image available for searching by content, you catalog the image. When you catalog an image, the Image Extender analyzes the image, by computing the feature values for the image, and stores the values in a QBIC catalog. When you search for an image by content, your query identifies one or more features for the search (such as average color), a source for each feature (such as an example image), and a target set of cataloged images. The Image Extender computes the feature value of the source and compares it to the 20 IBM® DB2® Universal Database: Image, Audio, and Video Extenders IBM 000040

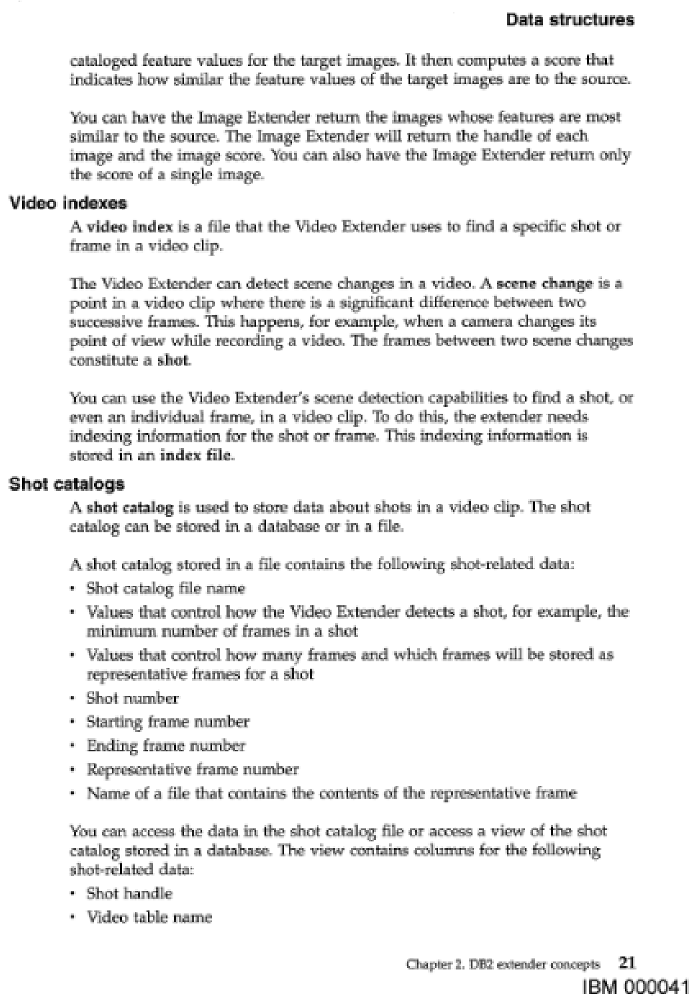

cataloged feature values for the target images. It then computes a score that indicates how similar the feature values of the target images are to the source.

You can have the Image Extender return the images whose features are most similar to the source. The Image Extender will return the handle of each image and the image score. You can also have the Image Extender return only the score of a single image.

Video indexes

A **video index** is a file that the Video Extender uses to find a specific shot or frame in a video clip.

The Video Extender can detect scene changes in a video. A **scene change** is a point in a video clip where there is a significant difference between two successive frames. This happens, for example, when a camera changes its point of view while recording a video. The frames between two scene changes constitute a **shot**.

You can use the Video Extender's scene detection capabilities to find a shot, or even an individual frame, in a video clip. To do this, the extender needs indexing information for the shot or frame. This indexing information is stored in an **index file**.

Shot catalogs

A **shot catalog** is used to store data about shots in a video clip. The shot catalog can be stored in a database or in a file.

A shot catalog stored in a file contains the following shot-related data:

- Shot catalog file name
- Values that control how the Video Extender detects a shot, for example, the minimum number of frames in a shot
- Values that control how many frames and which frames will be stored as representative frames for a shot
- Shot number
- Starting frame number
- Ending frame number
- Representative frame number
- Name of a file that contains the contents of the representative frame

You can access the data in the shot catalog file or access a view of the shot catalog stored in a database. The view contains columns for the following shot-related data:

- Shot handle
- Video table name

IBM0002263 at 226

Description

Ultimedia Manager 1.1 opens new dimensions in visual queries, helping you manage your image database more efficiently. With its easy-to-use interface, even entry-level users can perform sophisticated searches. You can search by image content, data description, or a combination of both. Once the images are found, you can view the images in full resolution, use the images in another application, or read the associated textual comments.

You can use Ultimedia Manager 1.1 either as a stand-alone product or in a LAN environment. As a stand-alone product, Ultimedia Manager 1.1 employs full image classification and image query capabilities. Within the LAN environment, the Ultimedia Manager image catalogs can be sorted on a DB2 for OS/2 or DB2 for AIX database server. The image catalogs can be accessed by Windows or OS/2 clients using the Client Search Feature or the Visualizer Ultimedia Query.

Ultimedia Manager runs as an OS/2 application to a DB2 for OS/2 or DB2 for AIX database server in a LAN, or as a stand-alone application with a local DB2 for OS/2 database. The Client Search Feature runs on DOS or Windows, and accesses files in dBASE flat file format. A database manager is not required to access catalogs in dBASE file format.

Image Classification

With the Image Classifier, you can create visual indexes based on color, shape, category, and layout elements, which are areas of the image containing the patterns you want to use for searches. You can store these elements in catalogs in DB2 for OS/2 databases. The catalogs can be used locally or, through the use of utility programs, be transferred to DB2 for OS/2 or AIX servers. In addition, you can convert the catalog data into dBASE flat file format for access by Windows client applications.

The Image Classifier and Analyzer allows automatic or interactive classification of images for later image queries, flexibility in storage of catalog information, and access to multiple users. You can organize large image catalogs so that later searches for images can be performed easily. You can classify automatically using a batch program, or interactively by specifying areas of interest in an image and indicating the visual cue (color, shape, or texture) to be saved for later queries.

Image Query

Once the image catalogs are created, you can use the Image Query component to find image and business information. Business information, such as character and numeric data, are stored in the database tables using normal table population methods. Image Query is provided in Ultimedia Manager and Client Search Feature.

Image Query allows searches by three methods:

- Image content, such as color, shape, texture, or layout
- Data description, such as business data stored in the database (character or numeric) or descriptive text (category)
- Combinations of image content and data description.

Image Content: To search by image content, specify the attributes (color, shape, texture, or layout). You can drag and drop visual samples into an image query, select a color from the color wheel, or sketch a shape.

Data Description: To search by data description, enter character and numeric data values in the database tables using the simple fill-in-the-blanks query panel. You can enter data such as name, number, date, or a range of dates. You can integrate character or numeric data as part of the image query, enabling a more effective search for better results. The simple fill-in-the-blanks method for specifying desired conditions is so easy even non-programmers can use database table information as part of their searches.

Combinations: Image queries using either image content or data description, or combinations of both, allow you to search through thousands of images more quickly and effectively than the traditional text and keyword search methods.

Image Formats

Ultimedia Manager 1.1 understands these image formats:

- Audio Visual Connection (AVC) images
- OS/2 bitmap
- Windows 3.x bitmap
- PCX
- Intel and Motorola TIFF
- TARGA TGA
- TIFF (FAX)
- GIF
- Kodak Photo-CD

Product Positioning

The advanced image search technology in Ultimedia Manager, query by image content, was developed in conjunction with the IBM Almaden Research Center. It provides function not available with conventional, text-only image management systems.

Ultimedia Manager fulfills image search and retrieval needs of business users and creative professionals through use of online image libraries or collections. It is intended for users who have large image libraries and need fast, flexible access to their image data, users who may also want to add character or numeric data to their image queries, and users who feel limited by text-based systems.

Ultimedia Manager 1.1 works with the DB2 Family of products and the Visualizer Ultimedia Query product. Catalogs created by Ultimedia Manager can be stored in DB2 for OS/2 database tables. Image catalogs also can be stored in DB2 for OS/2 or DB2 for AIX database servers in client/server environments.

OS/2 and Windows clients in a LAN environment using the Client Search Feature and the Visualizer Ultimedia Query can access catalogs created with Ultimedia Manager 1.1. Visualizer Ultimedia Query provides users with the capability to:

- Add multimedia columns (such as image, video, audio, and formatted text) to DB2 for OS/2 database tables
- Query and display images
- View or listen to video and audio information
- Display formatted text

Users of Ultimedia Manager 1.1 and Client Search Feature are focused on image management and query applications. Visualizer Ultimedia Query users are

295-042 -2-

IBM 0002264

IBM000777 at 794

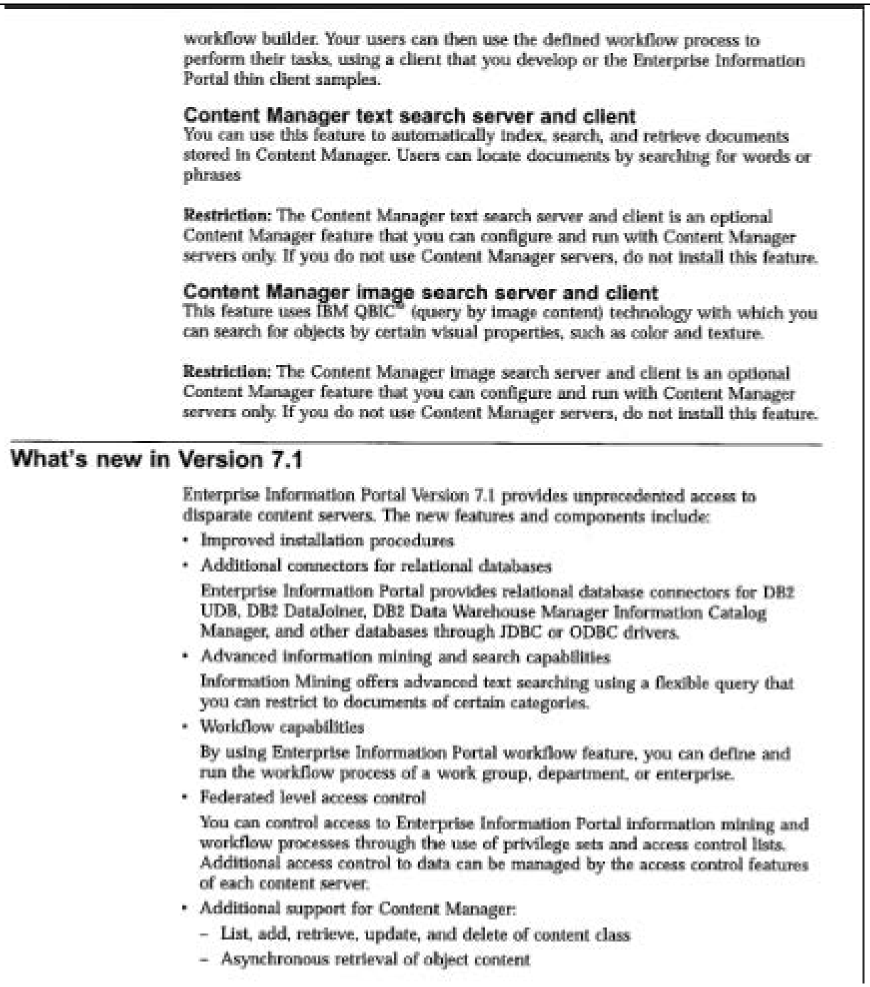

workflow builder. Your users can then use the defined workflow process to perform their tasks, using a client that you develop or the Enterprise Information Portal thin client samples.

Content Manager text search server and client
You can use this feature to automatically index, search, and retrieve documents stored in Content Manager. Users can locate documents by searching for words or phrases

Restriction: The Content Manager text search server and client is an optional Content Manager feature that you can configure and run with Content Manager servers only. If you do not use Content Manager servers, do not install this feature.

Content Manager image search server and client
This feature uses IBM QBIC™ (query by image content) technology with which you can search for objects by certain visual properties, such as color and texture.

Restriction: The Content Manager image search server and client is an optional Content Manager feature that you can configure and run with Content Manager servers only. If you do not use Content Manager servers, do not install this feature.

What's new in Version 7.1

Enterprise Information Portal Version 7.1 provides unprecedented access to disparate content servers. The new features and components include:

- Improved installation procedures
- Additional connectors for relational databases
  Enterprise Information Portal provides relational database connectors for DB2 UDB, DB2 DataJoiner, DB2 Data Warehouse Manager Information Catalog Manager, and other databases through JDBC or ODBC drivers.
- Advanced information mining and search capabilities
  Information Mining offers advanced text searching using a flexible query that you can restrict to documents of certain categories.
- Workflow capabilities
  By using Enterprise Information Portal workflow feature, you can define and run the workflow process of a work group, department, or enterprise.
- Federated level access control
  You can control access to Enterprise Information Portal information mining and workflow processes through the use of privilege sets and access control lists. Additional access control to data can be managed by the access control features of each content server.
- Additional support for Content Manager:
  - List, add, retrieve, update, and delete of content class
  - Asynchronous retrieval of object content

To the extent it is found that QBIC does not expressly disclose this limitation, the limitation is inherent. And to the extent it is found that QBIC does not anticipate this claim, QBIC renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified herein or in the cover pleading of BofA's

		Invalidity Contentions. A person of ordinary skill in the art would have been motivated to combine QBIC with any of the other prior art references identified in the cover pleading or herein with a reasonable expectation of success for the reasons identified in the cover pleading, including because each of these references contains familiar elements within the similar fields of endeavor to address similar known problems, according to known methods which would yield predictable results. For example, a person of ordinary skill in the art would have been motivated to improve, modify, or combine QBIC with prior art reference teachings to arrive at this limitation at least because QBIC was designed for applications that included large amounts of business-related information, including multi-media content, including for financial institutions. *See, e.g.*, IBM 000789; Flickner Depo Tr. at 191:16-192:8, 293:14-294:8.