



US 20130155474A1

(19) **United States**

(12) **Patent Application Publication**
Roach et al.

(10) **Pub. No.: US 2013/0155474 A1**

(43) **Pub. Date: Jun. 20, 2013**

(54) **SYSTEMS AND METHODS FOR AUTOMATIC IMAGE CAPTURE ON A MOBILE DEVICE**

(60) Provisional application No. 61/561,772, filed on Nov. 18, 2011, provisional application No. 61/022,279, filed on Jan. 18, 2008.

(71) Applicant: **MITEK SYSTEMS**, San Diego, CA (US)

(72) Inventors: **John J. Roach**, San Diego, CA (US); **Grigori Nepomniachtchi**, San Diego, CA (US); **Robert Couch**, Poway, CA (US); **Oleg Rybakov**, San Diego, CA (US); **Michael Gillen**, San Diego, CA (US); **Kevin Andrew Bell**, Seattle, WA (US)

(73) Assignee: **MITEK SYSTEMS**, San Diego, CA (US)

(21) Appl. No.: **13/769,833**

(22) Filed: **Feb. 19, 2013**

Related U.S. Application Data

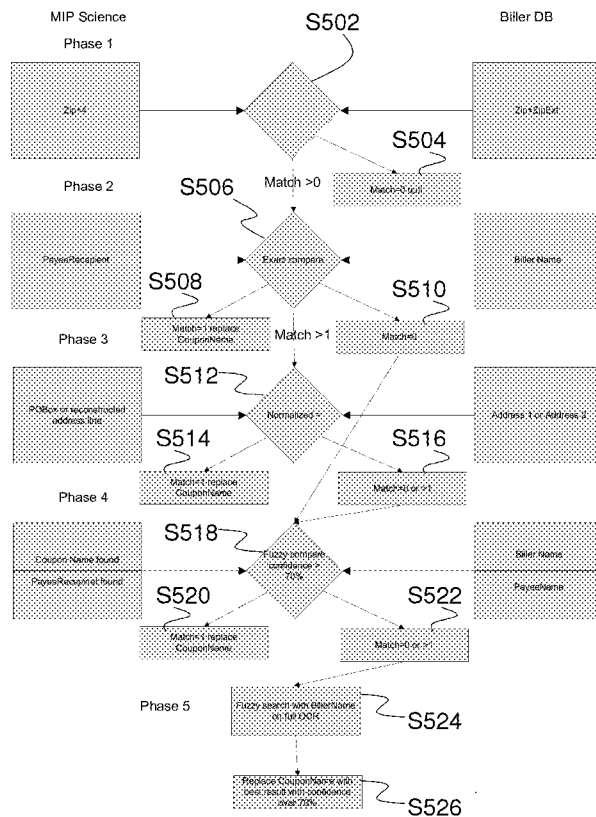
(63) Continuation of application No. 13/461,726, filed on May 1, 2012, Continuation-in-part of application No. 12/906,036, filed on Oct. 15, 2010, which is a continuation-in-part of application No. 12/778,943, filed on May 12, 2010, Continuation-in-part of application No. 12/346,026, filed on Dec. 30, 2008, now Pat. No. 7,978,900.

Publication Classification

(51) **Int. Cl.**
H04N 1/00 (2006.01)
(52) **U.S. Cl.**
CPC **H04N 1/00129** (2013.01)
USPC **358/505; 358/473**

(57) **ABSTRACT**

Real-time evaluation and enhancement of image quality prior to capturing an image of a document on a mobile device is provided. An image capture process is initiated on a mobile device during which a user of the mobile device prepares to capture the image of the document, utilizing hardware and software on the mobile device to measure and achieve optimal parameters for image capture. Feedback may be provided to a user of the mobile device to instruct the user on how to manually optimize certain parameters relating to image quality, such as the angle, motion and distance of the mobile device from the document. When the optimal parameters for image capture of the document are achieved, at least one image of the document is automatically captured by the mobile device.



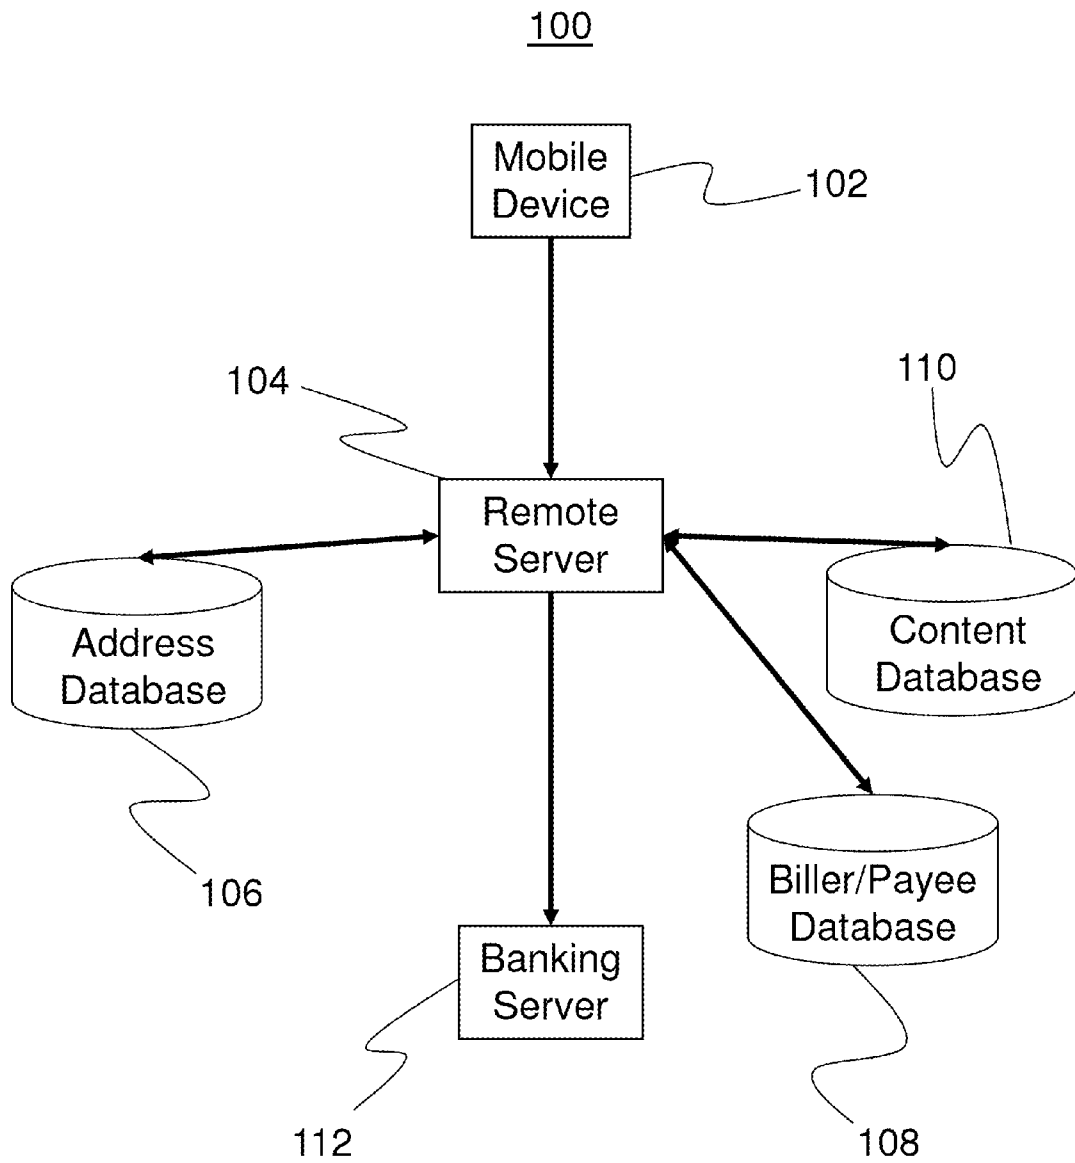


FIG. 1

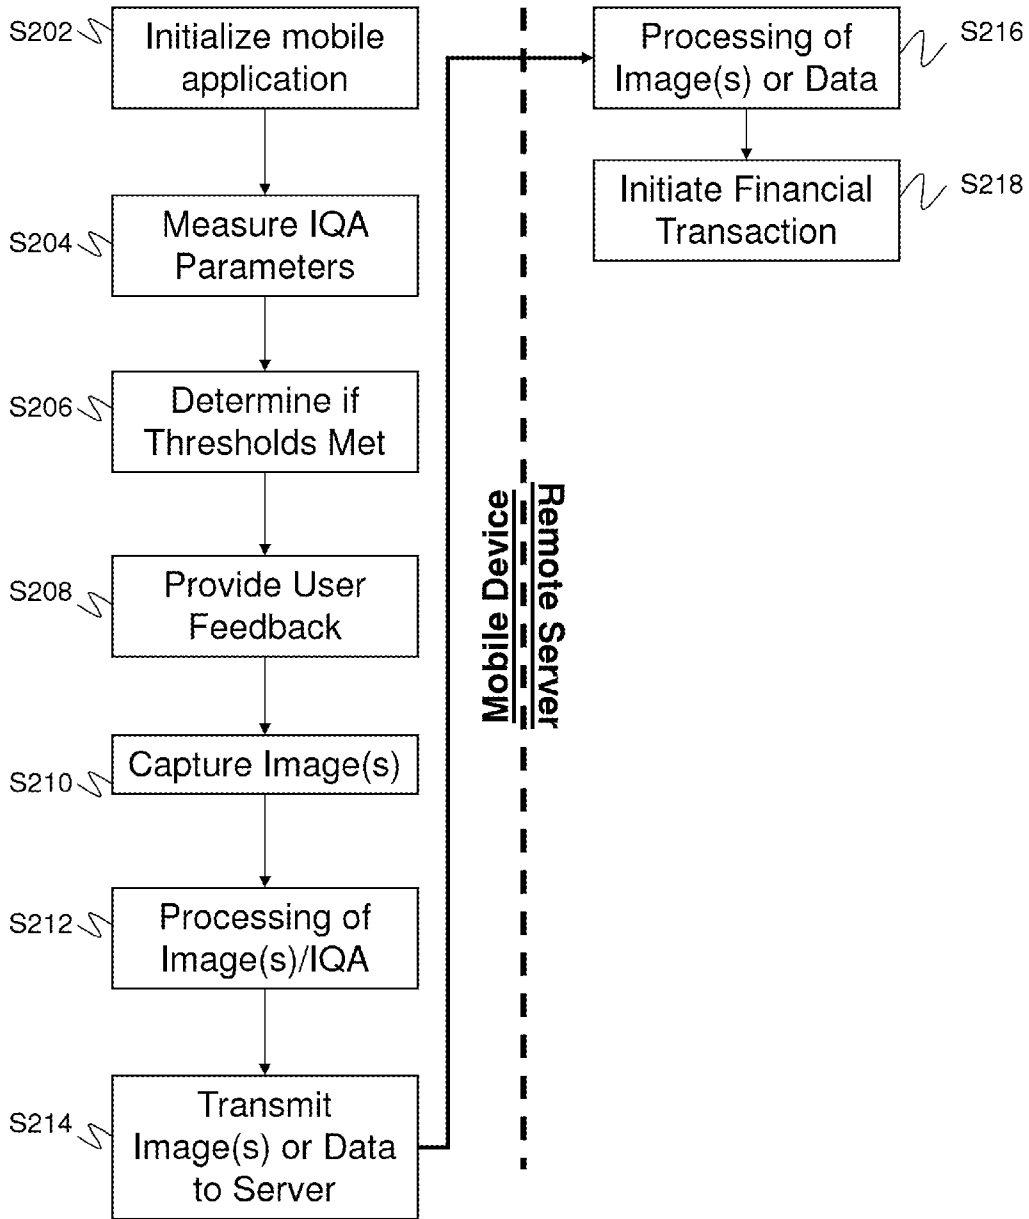


FIG. 2






FIG. 3



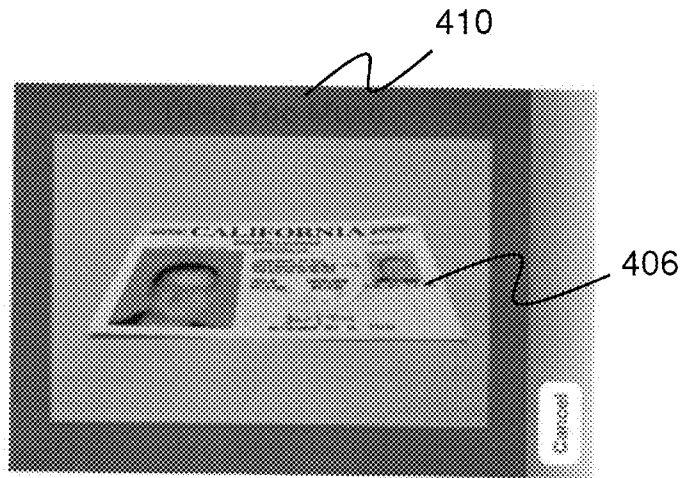





FIG. 4E





FIG. 4F



412

FIG. 4G

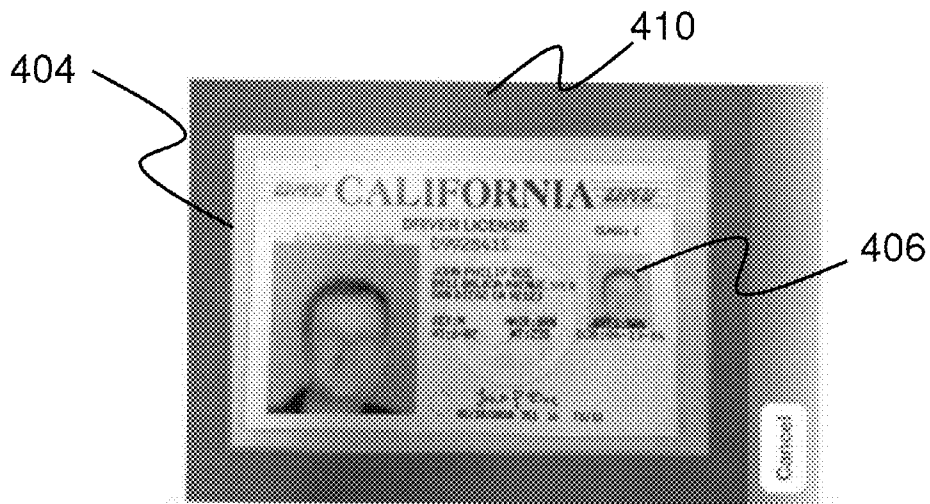


FIG. 4H

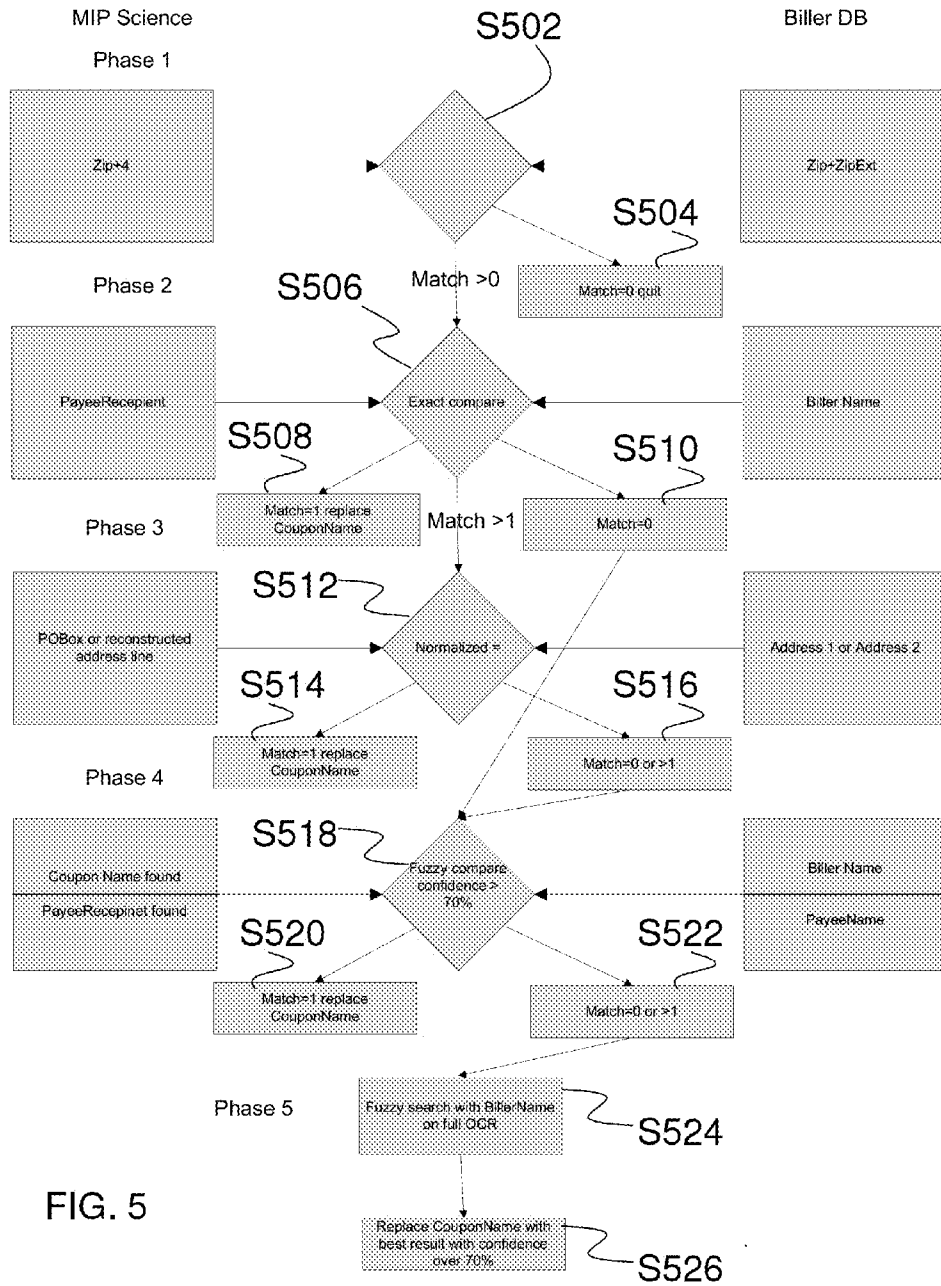


FIG. 5

Second Content Recognition Process

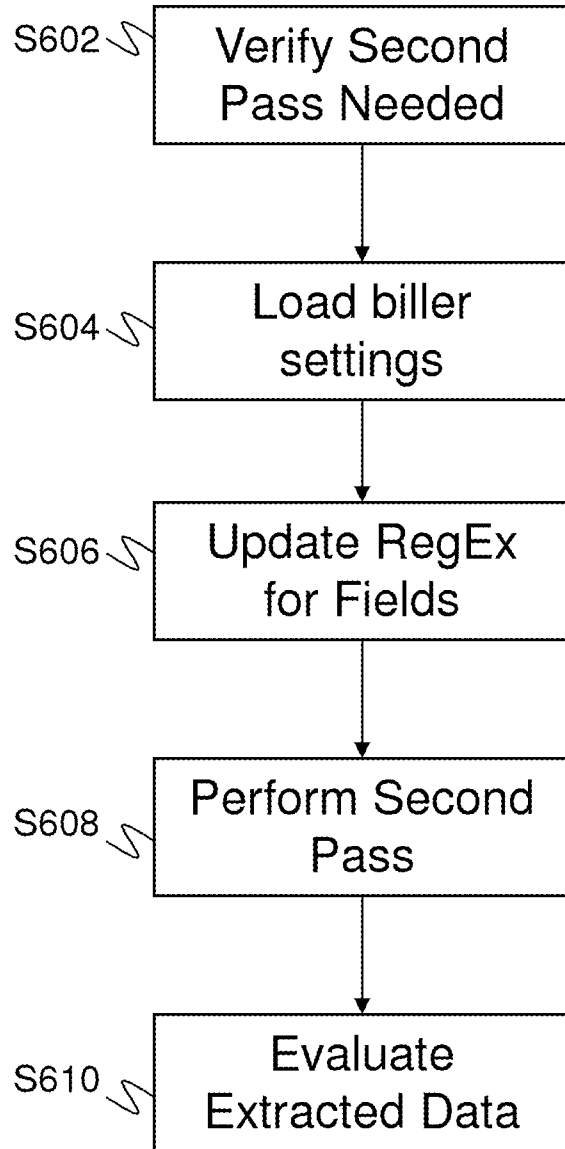




FIG. 6



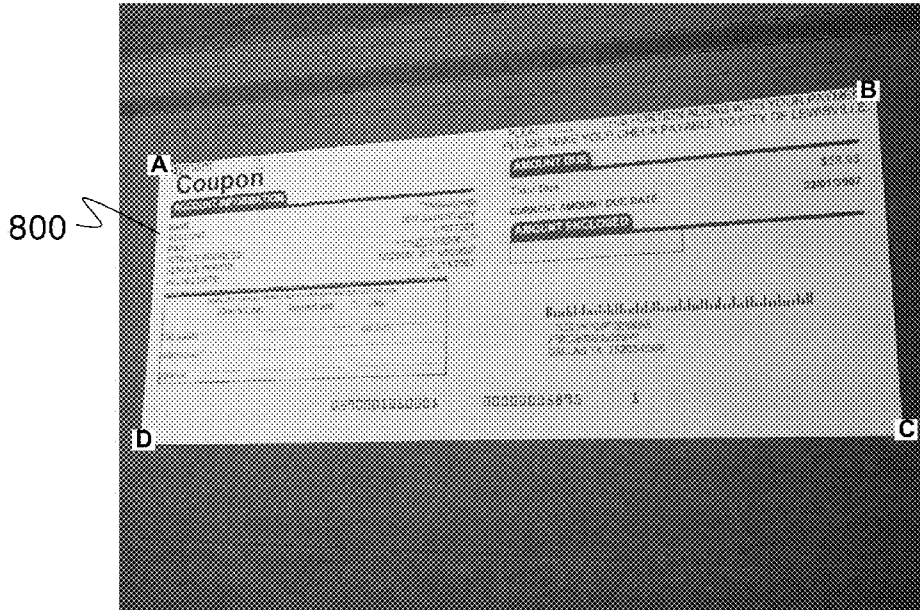


FIG. 8

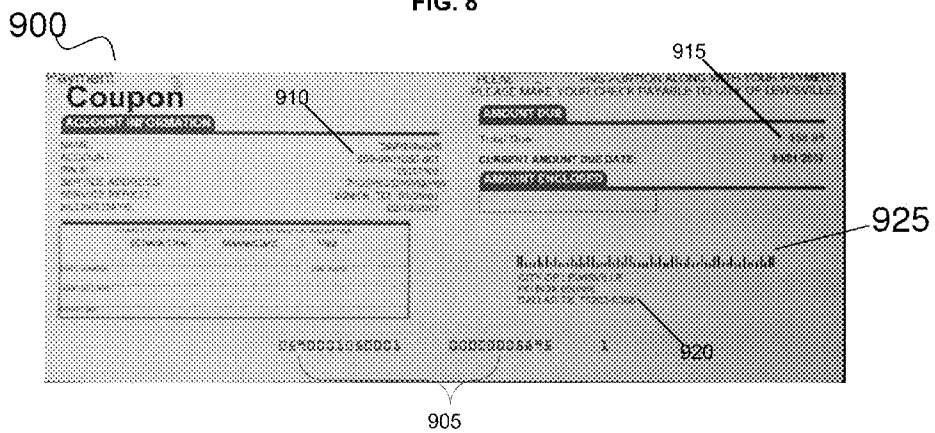


FIG. 9




FIG. 10




FIG. 11

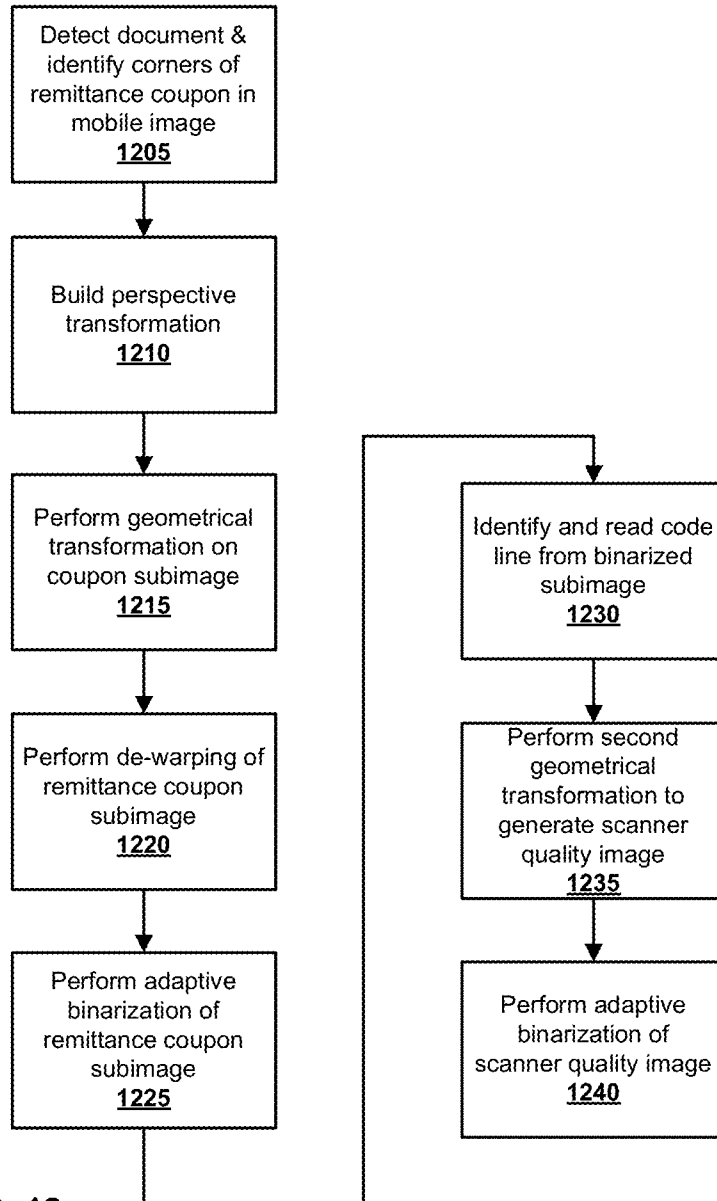


FIG. 12
Image
Correction
Module




FIG. 13

Detect Document in Image

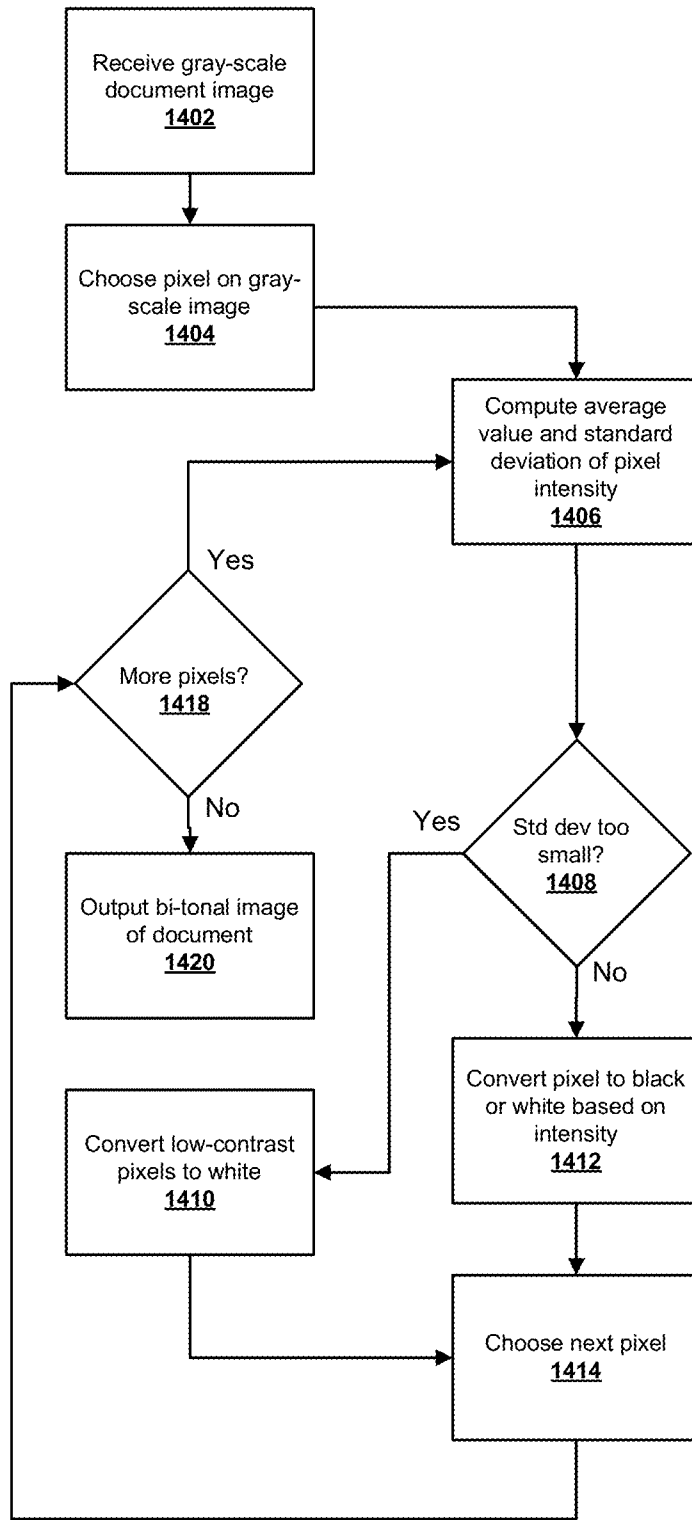


FIG. 14

Payment ³
Coupon

ACCOUNT INFORMATION

NAME: _____
ACCOUNT: 059-0001050-001
PIN #: 03077301
SERVICE ADDRESS: _____
SERVICE PERIOD: 12/29/06 TO: 01/29/07
BILLING DATE: 02/13/2007

PLEASE PRINT THIS PORTION ALONG WITH YOUR PAYMENT
PLEASE MAKE YOUR CHECK PAYABLE TO CITY OF LEWISVILLE.

AMOUNT DUE

Total Due \$58.95
CURRENT AMOUNT DUE DATE: 03/01/2007

AMOUNT ENCLOSED

COMPLETE THE FOLLOWING ONLY IF PAID BY MASTERCARD OR VISA
(Check One) MasterCard Visa

CARD NUMBER _____ EXP. DATE _____
CARD TYPE _____
SIGNATURE _____

0590001050001 00000005895 1

CITY OF LEWISVILLE
PO BOX 850568
DALLAS TX 75265-0568

FIG. 15




FIG. 16
Convert original color image into smaller color image (color "icon" image)

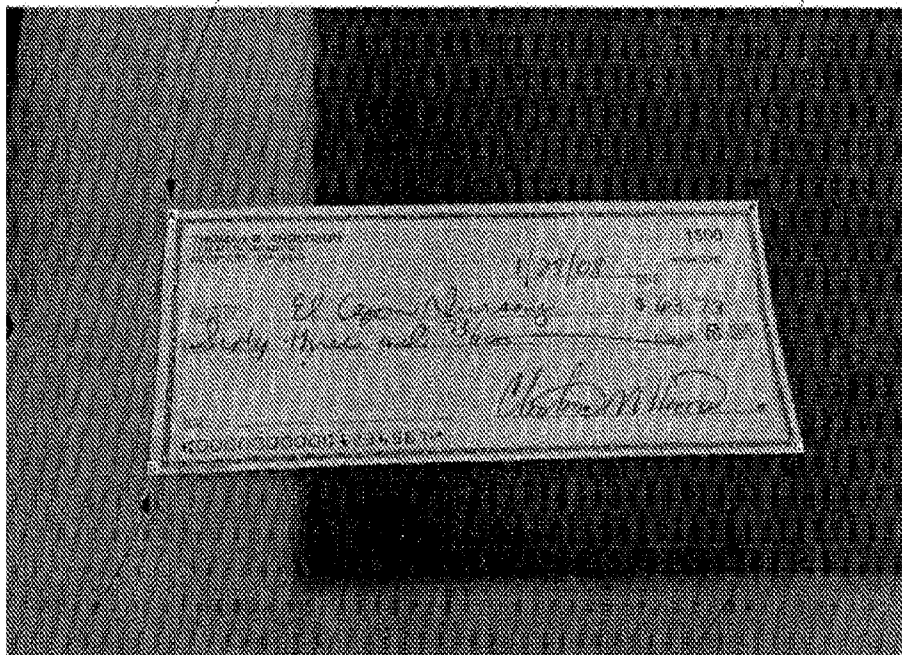


FIG. 17A




FIG. 17B

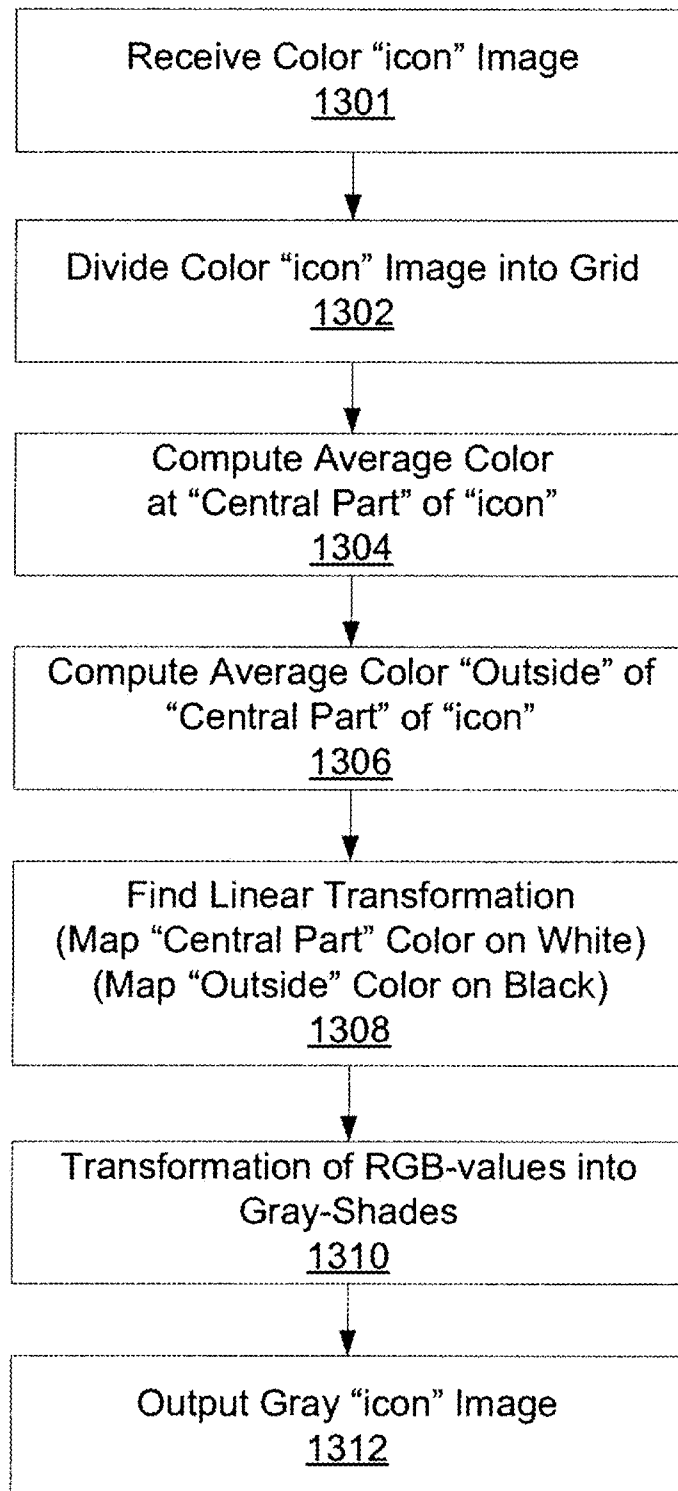


FIG. 18



FIG. 19A




FIG. 19B

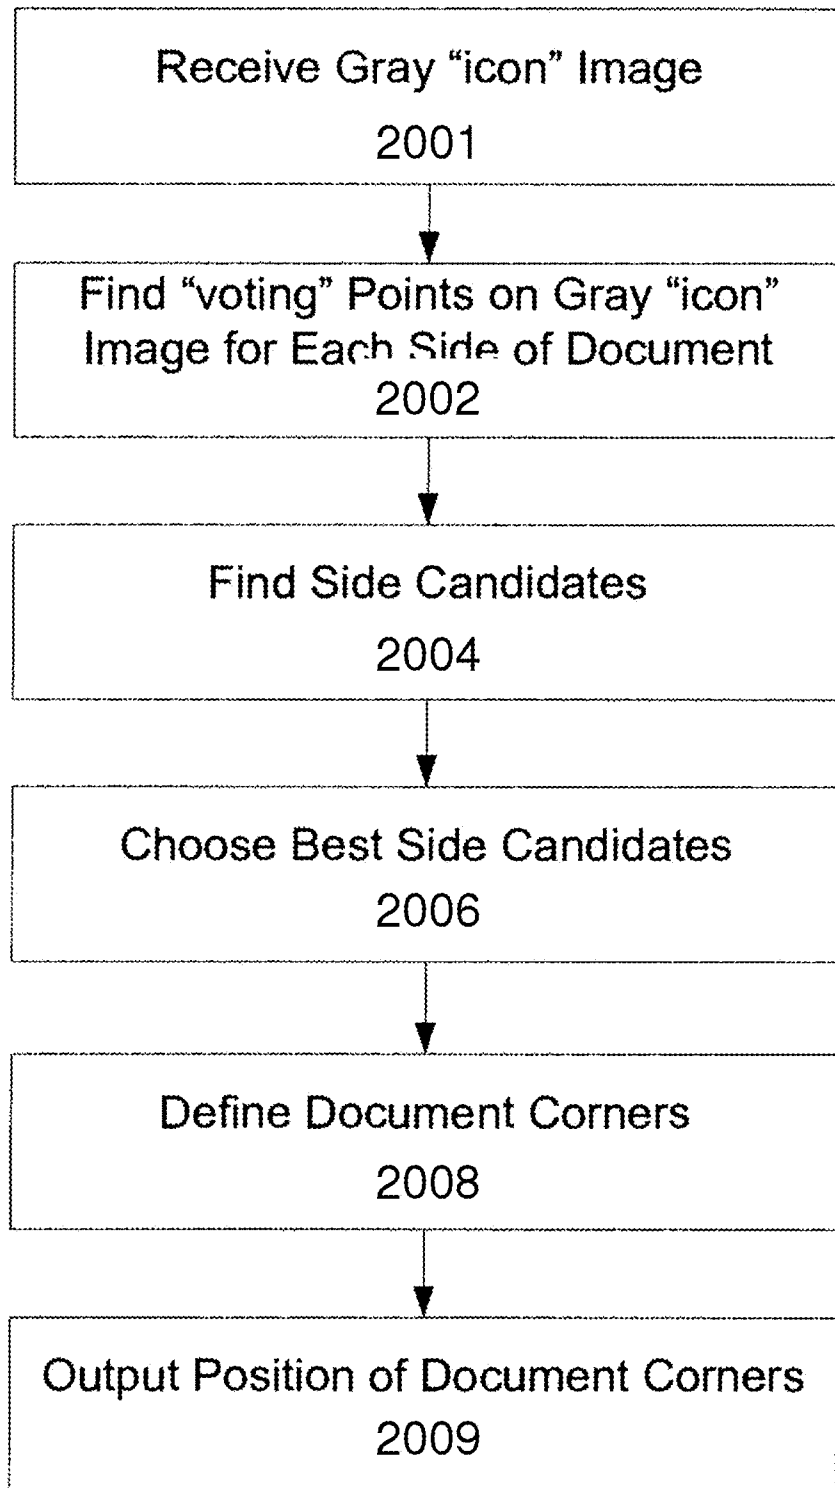


FIG. 20

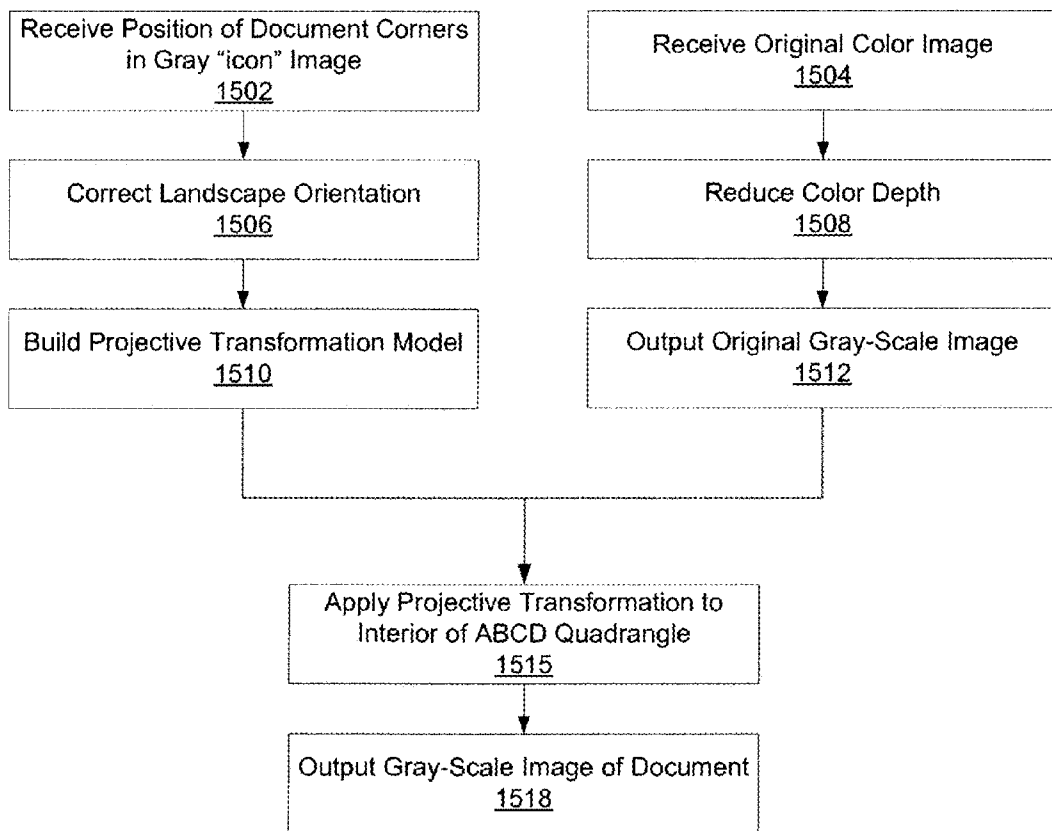


FIG. 21




FIG. 22A
The check in this image is intentionally depicted in landscape orientation.




FIG. 22B

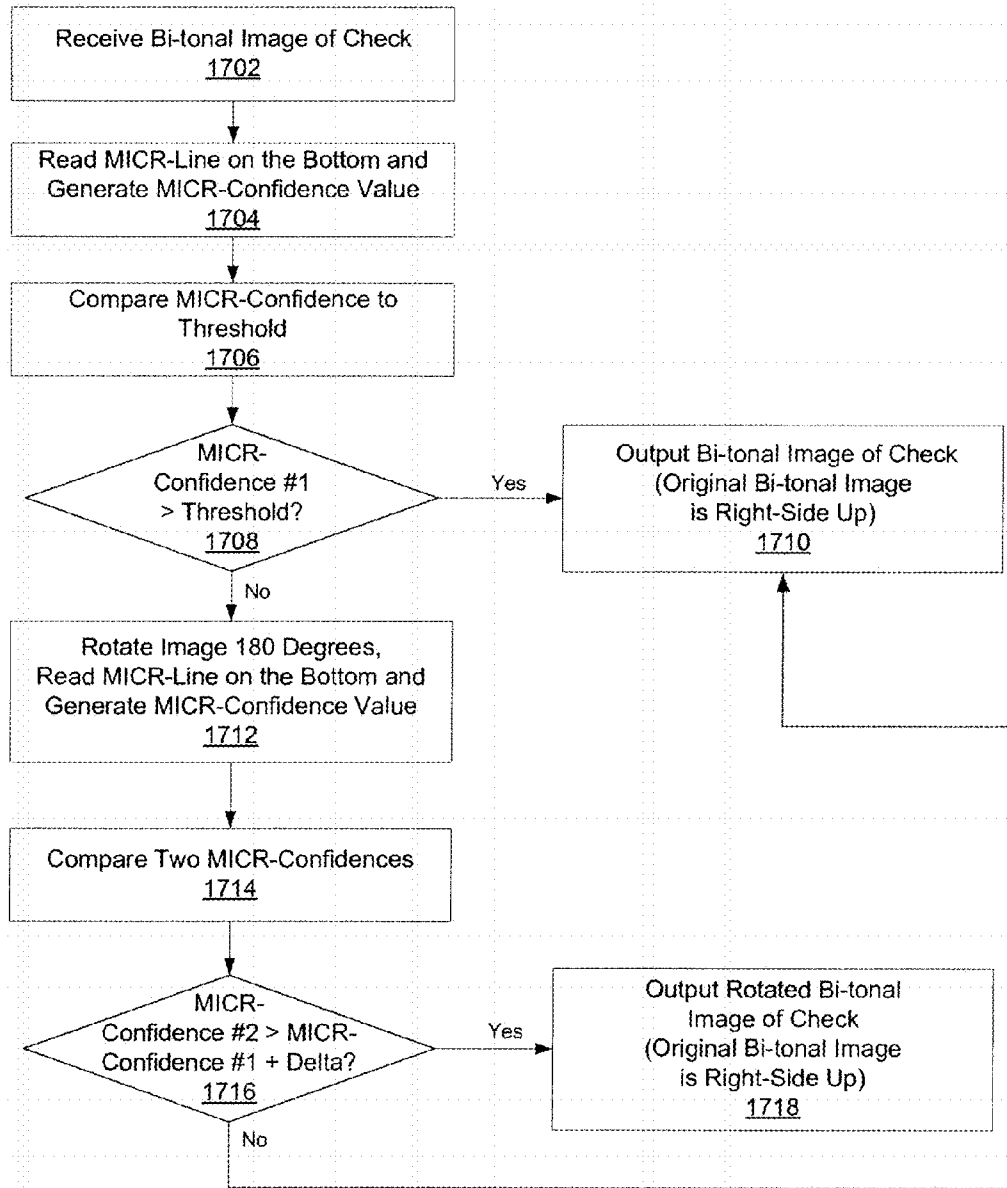


FIG. 23

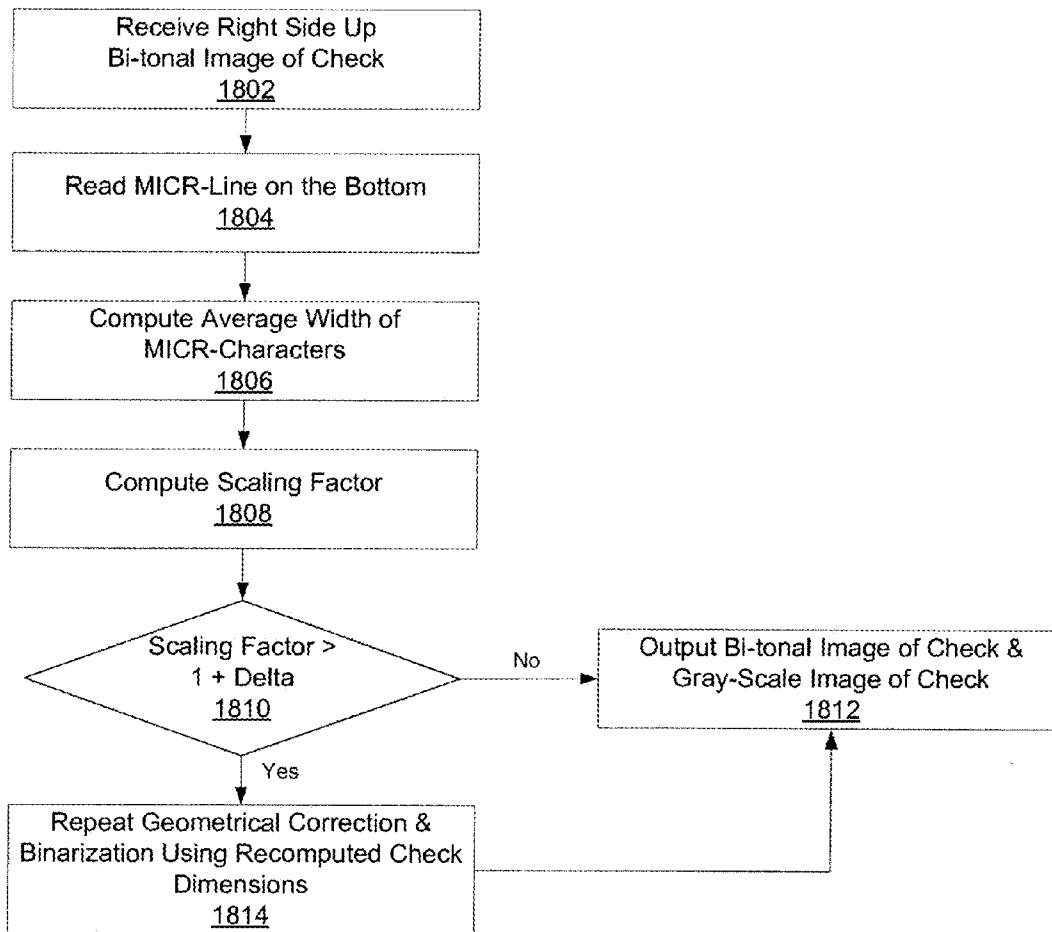


FIG. 24

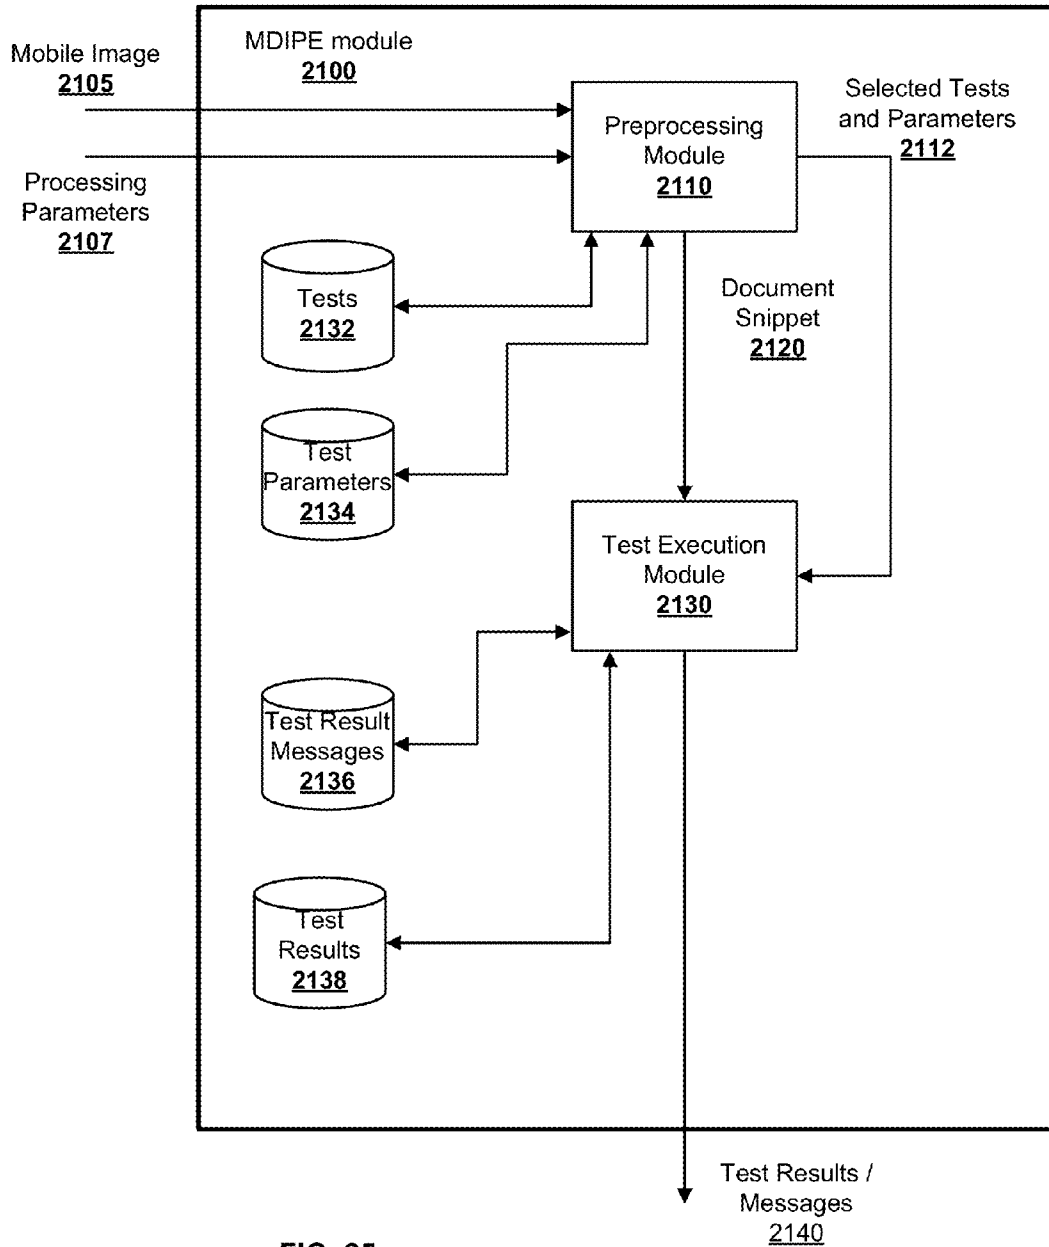
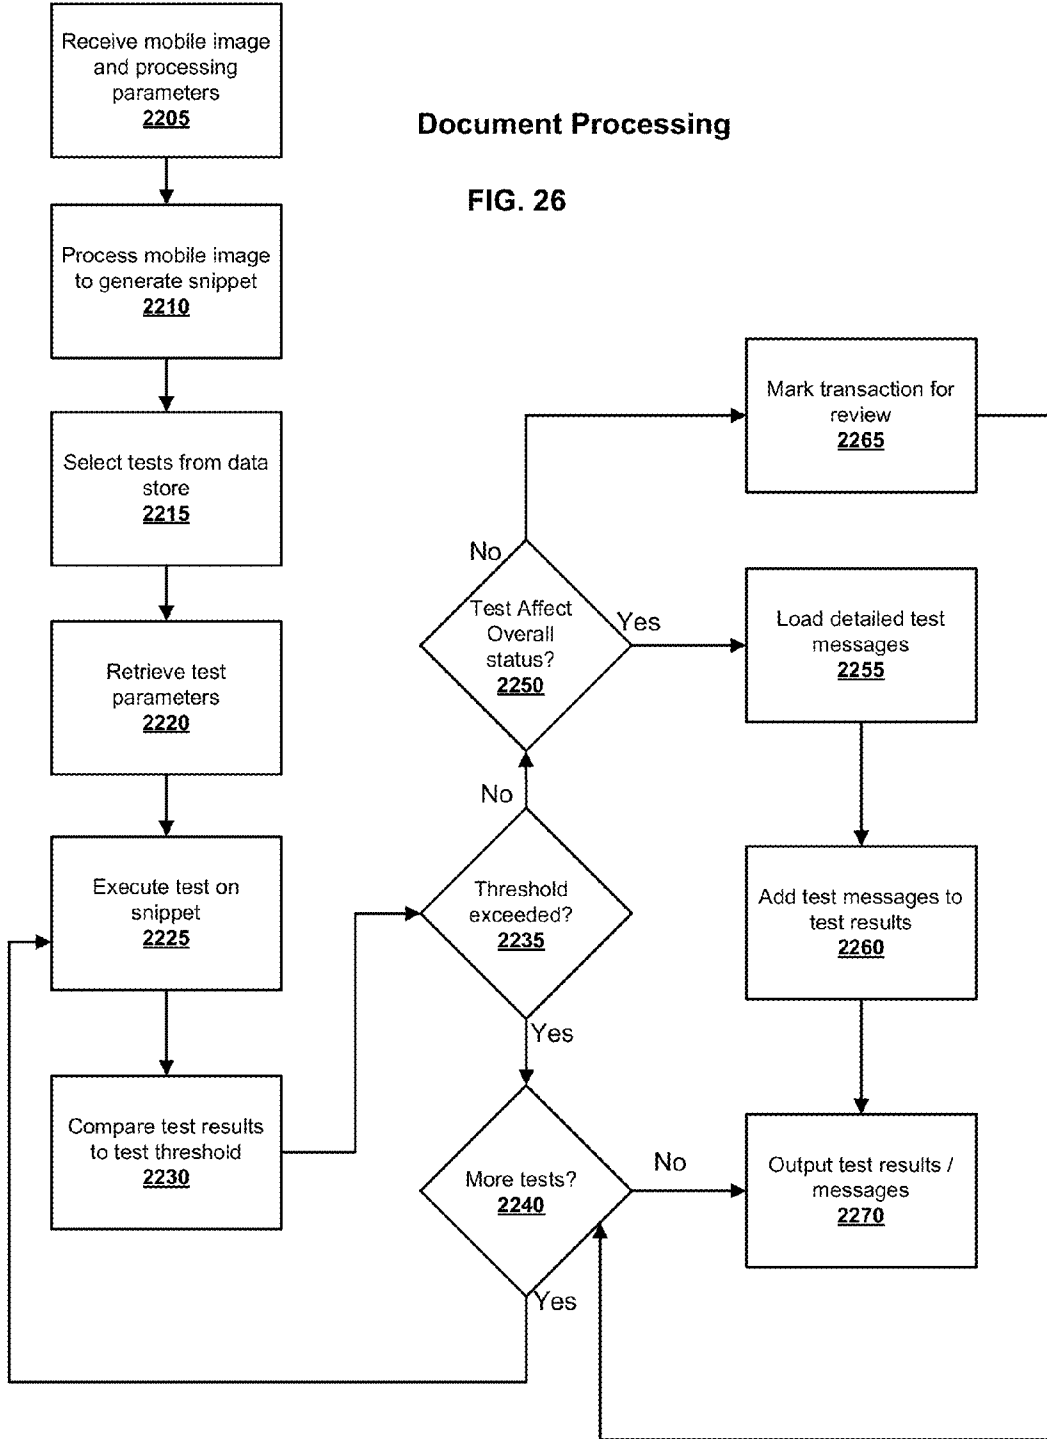
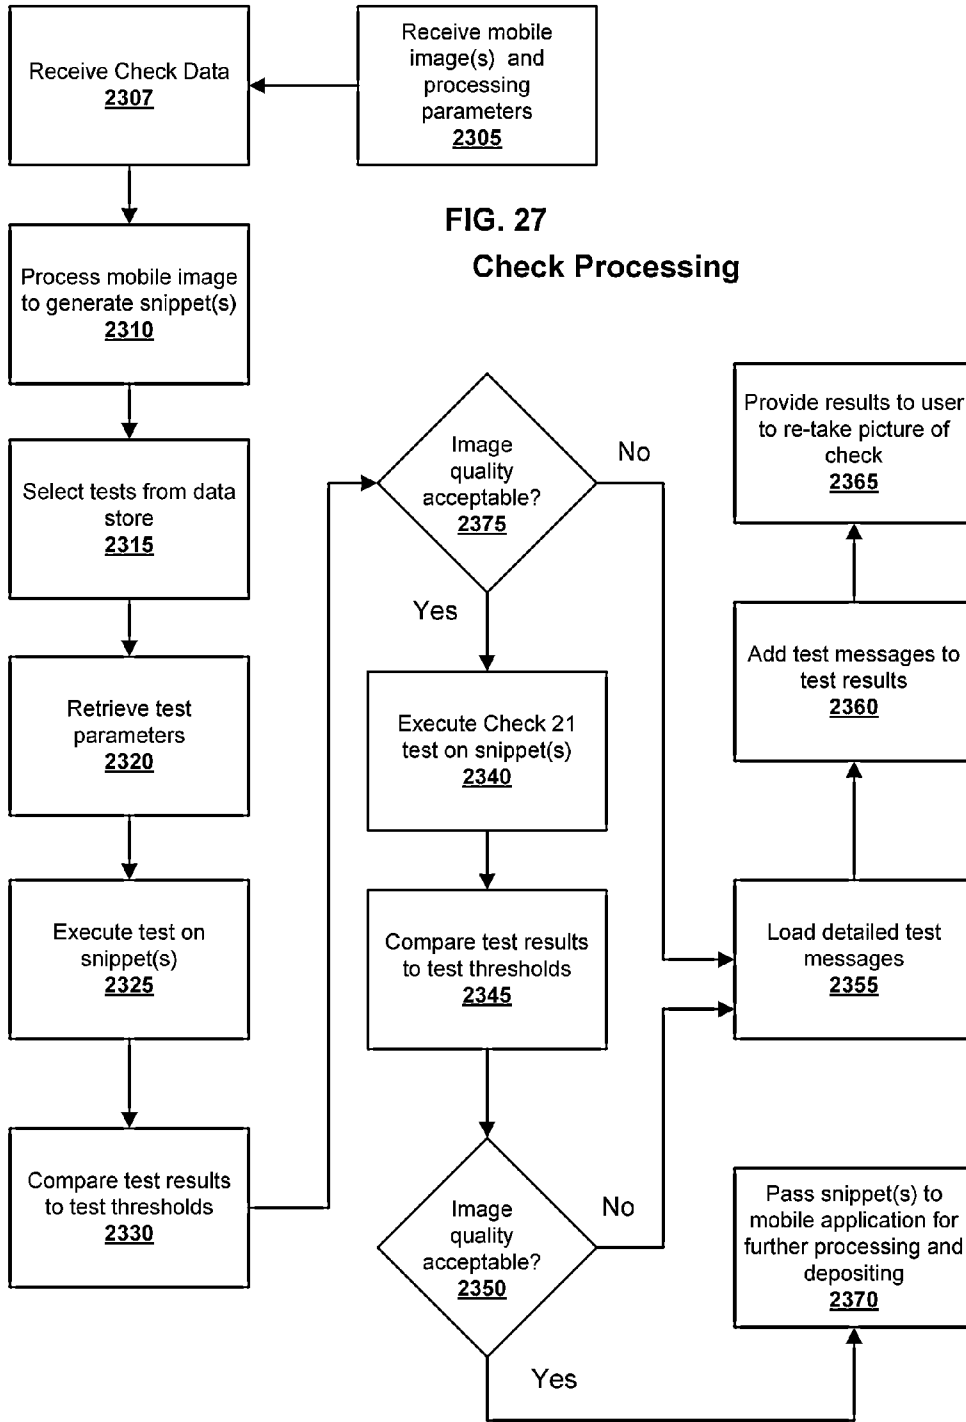


FIG. 25

Document Processing

FIG. 26





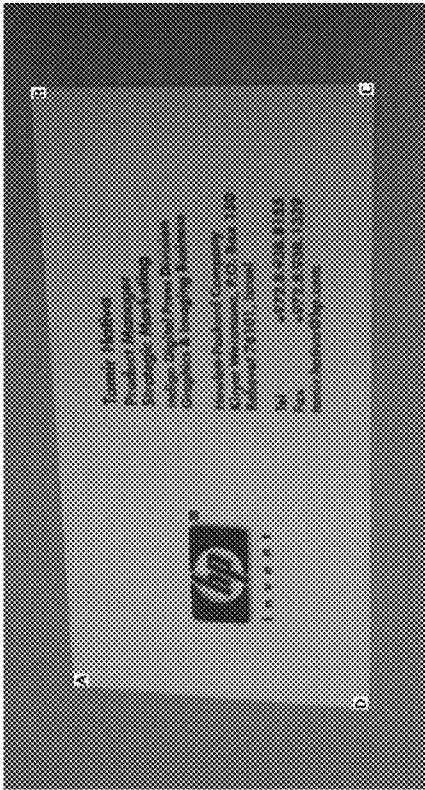


FIG. 28A



FIG. 28B



FIG. 29A




FIG. 29B




FIG. 30




FIG. 31

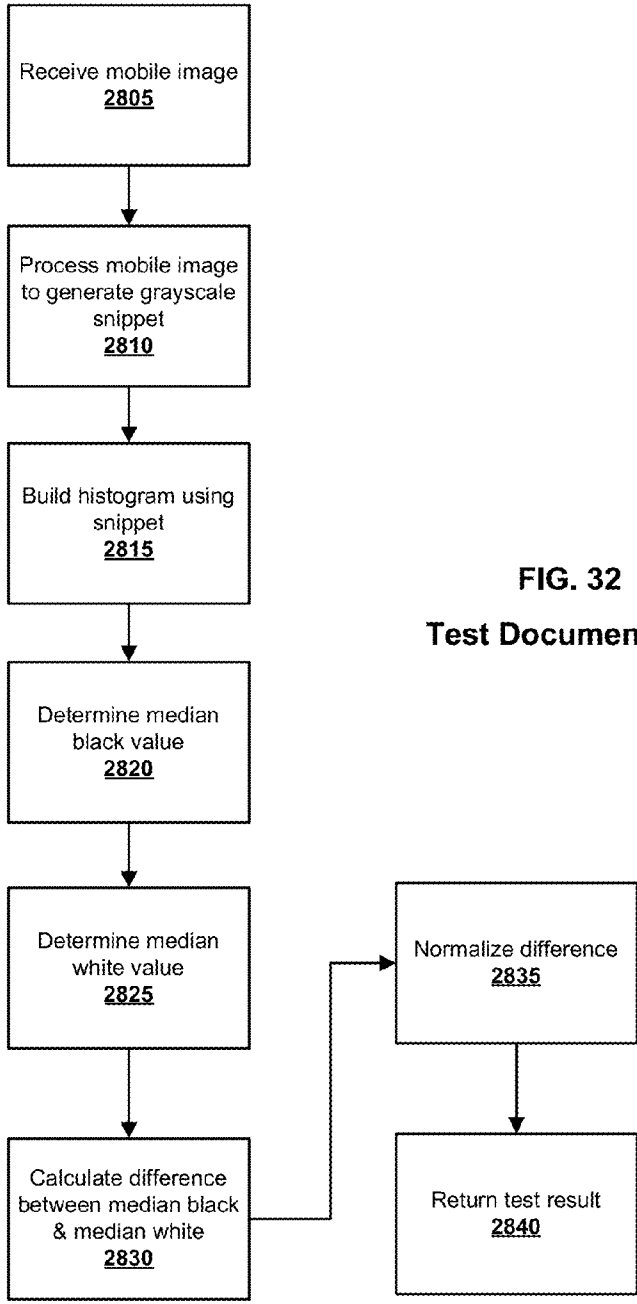


FIG. 32
Test Document Contrast

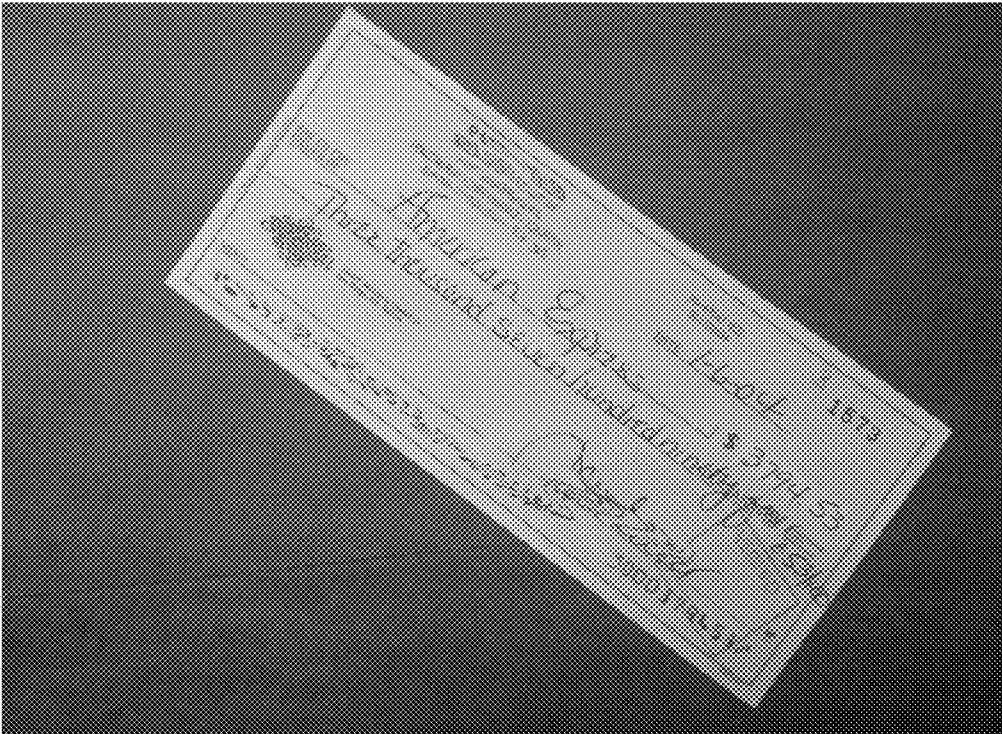


FIG. 33A



FIG. 33B




FIG. 34
Test Document Skew




FIG. 35

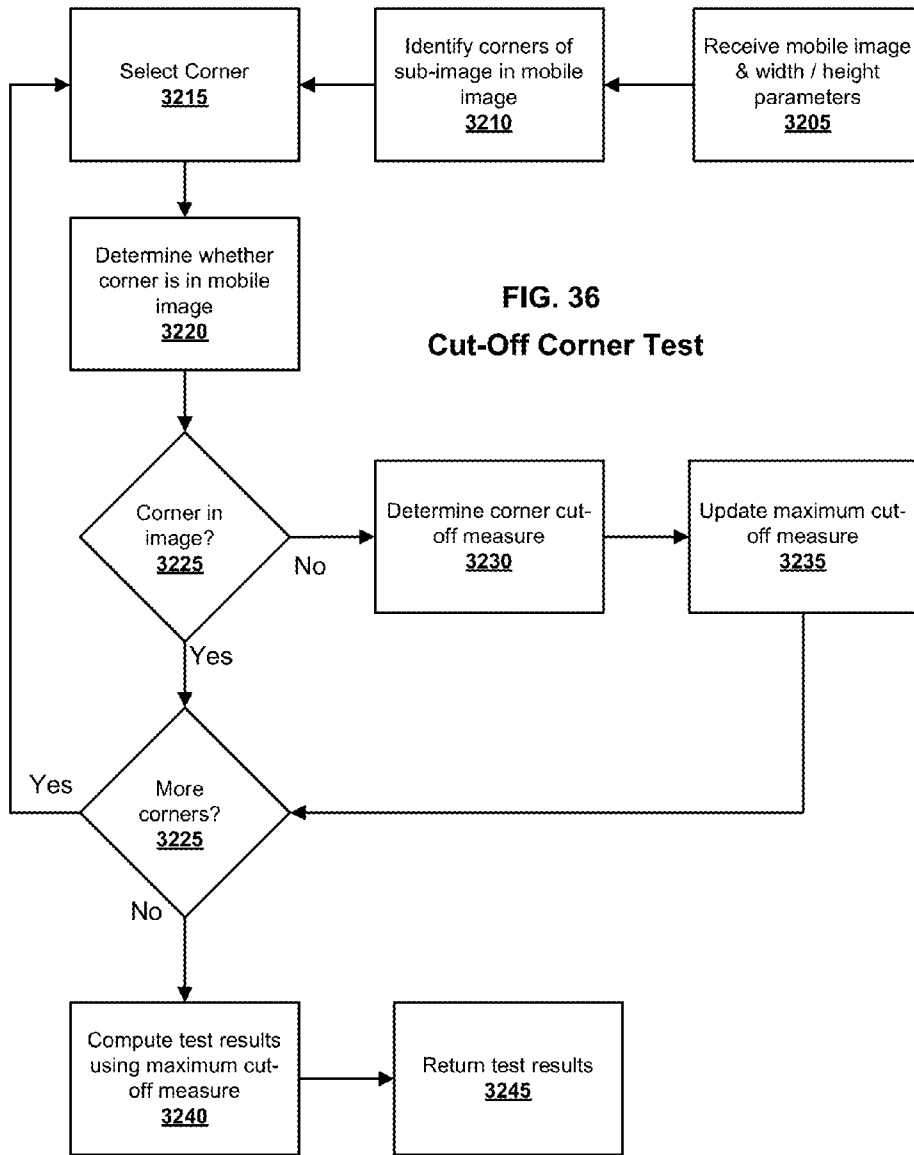


FIG. 36
Cut-Off Corner Test




FIG. 37

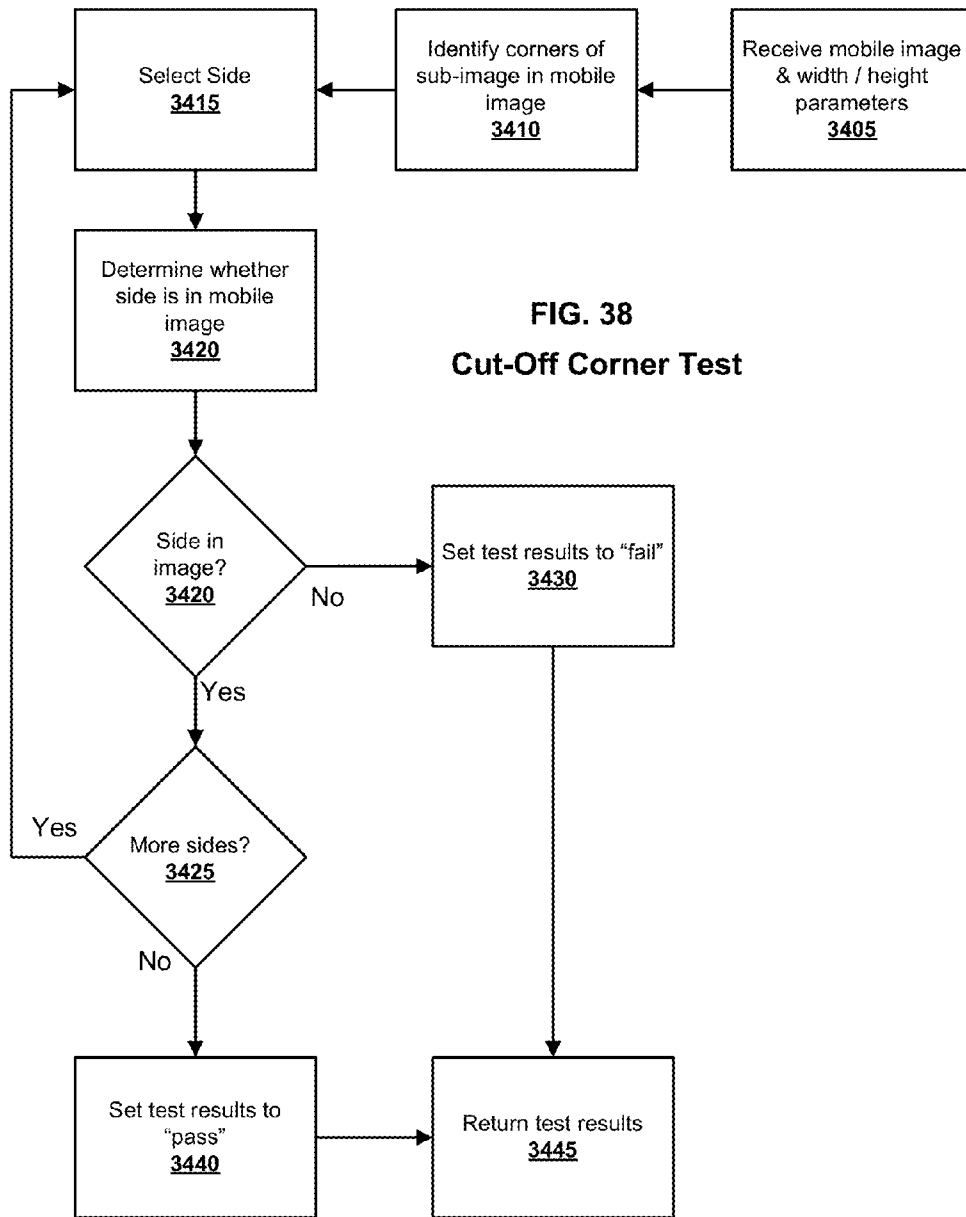


FIG. 38
Cut-Off Corner Test


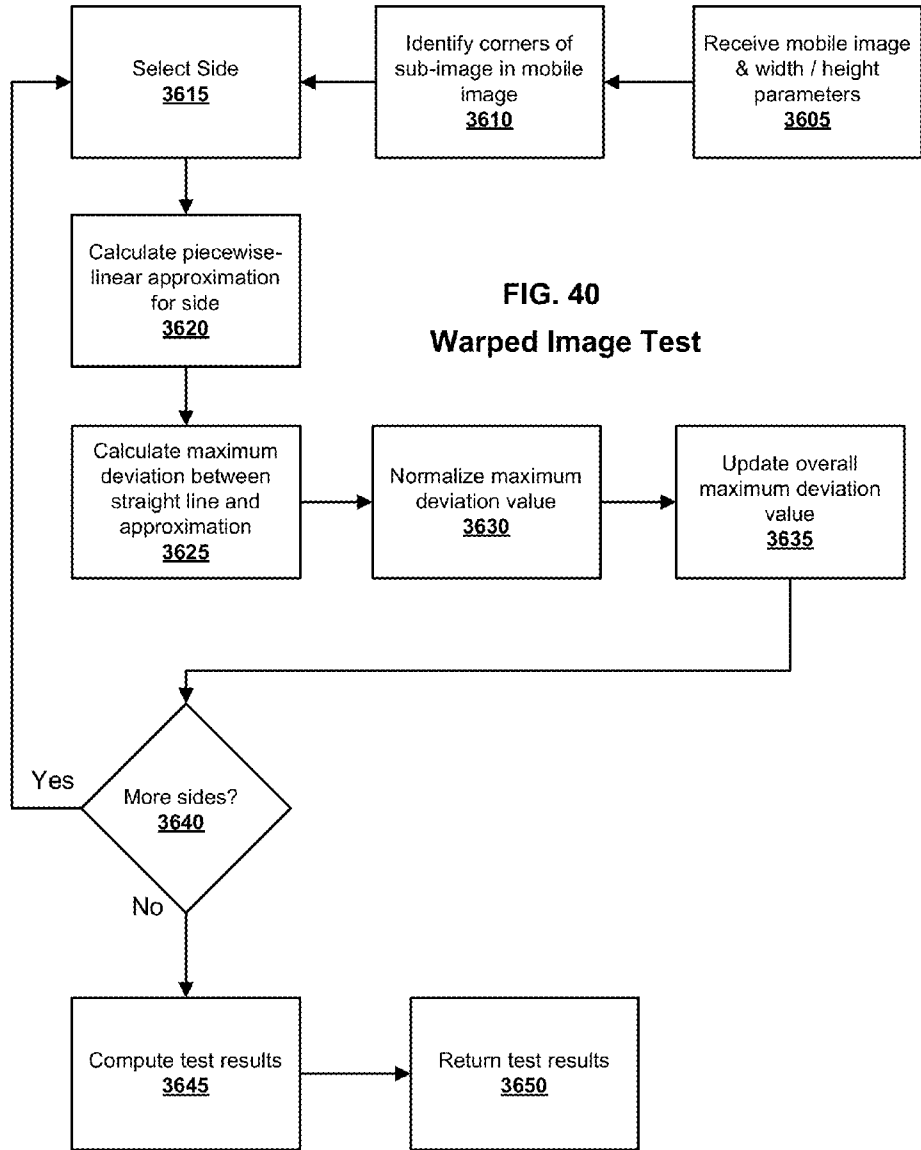


FIG. 39




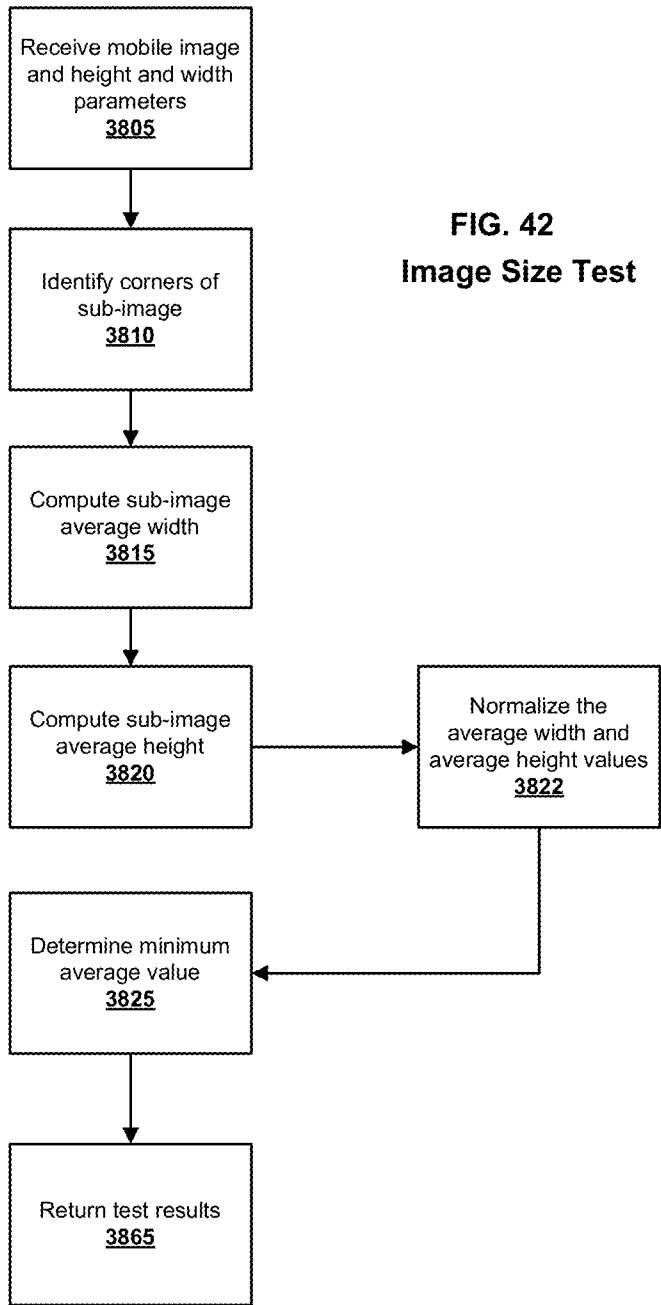


FIG. 41






FIG. 43
Code line Test

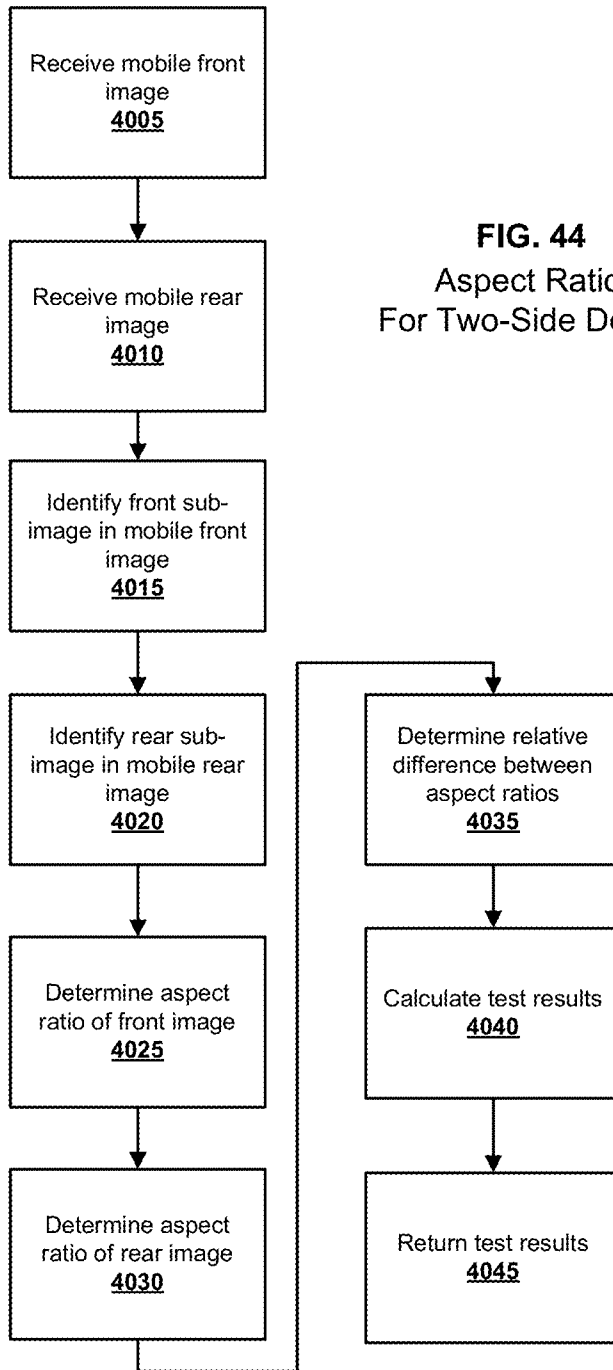


FIG. 44
Aspect Ratio Test
For Two-Side Documents

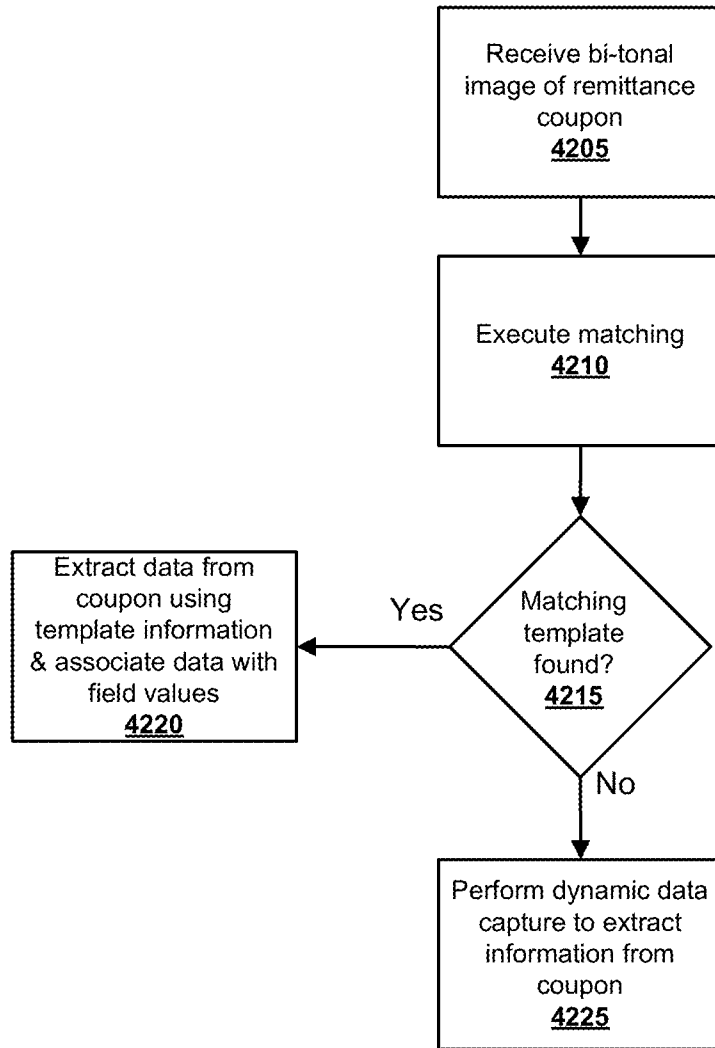


FIG. 45
Form Identification
and Data Extraction




FIG. 46
Dynamic Data
Capture

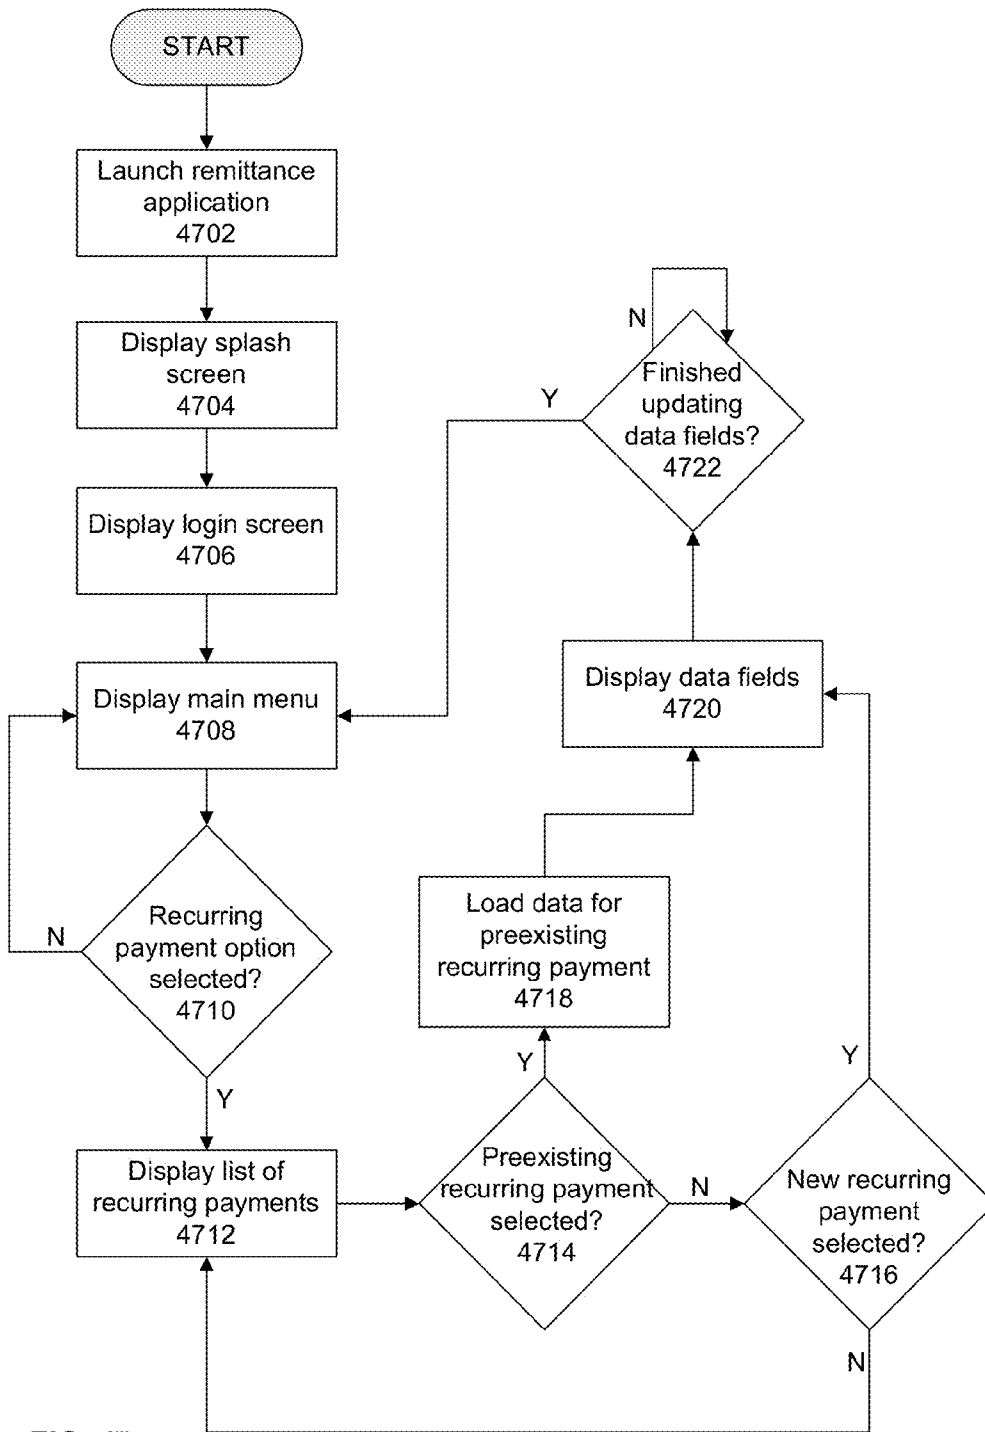


FIG. 47

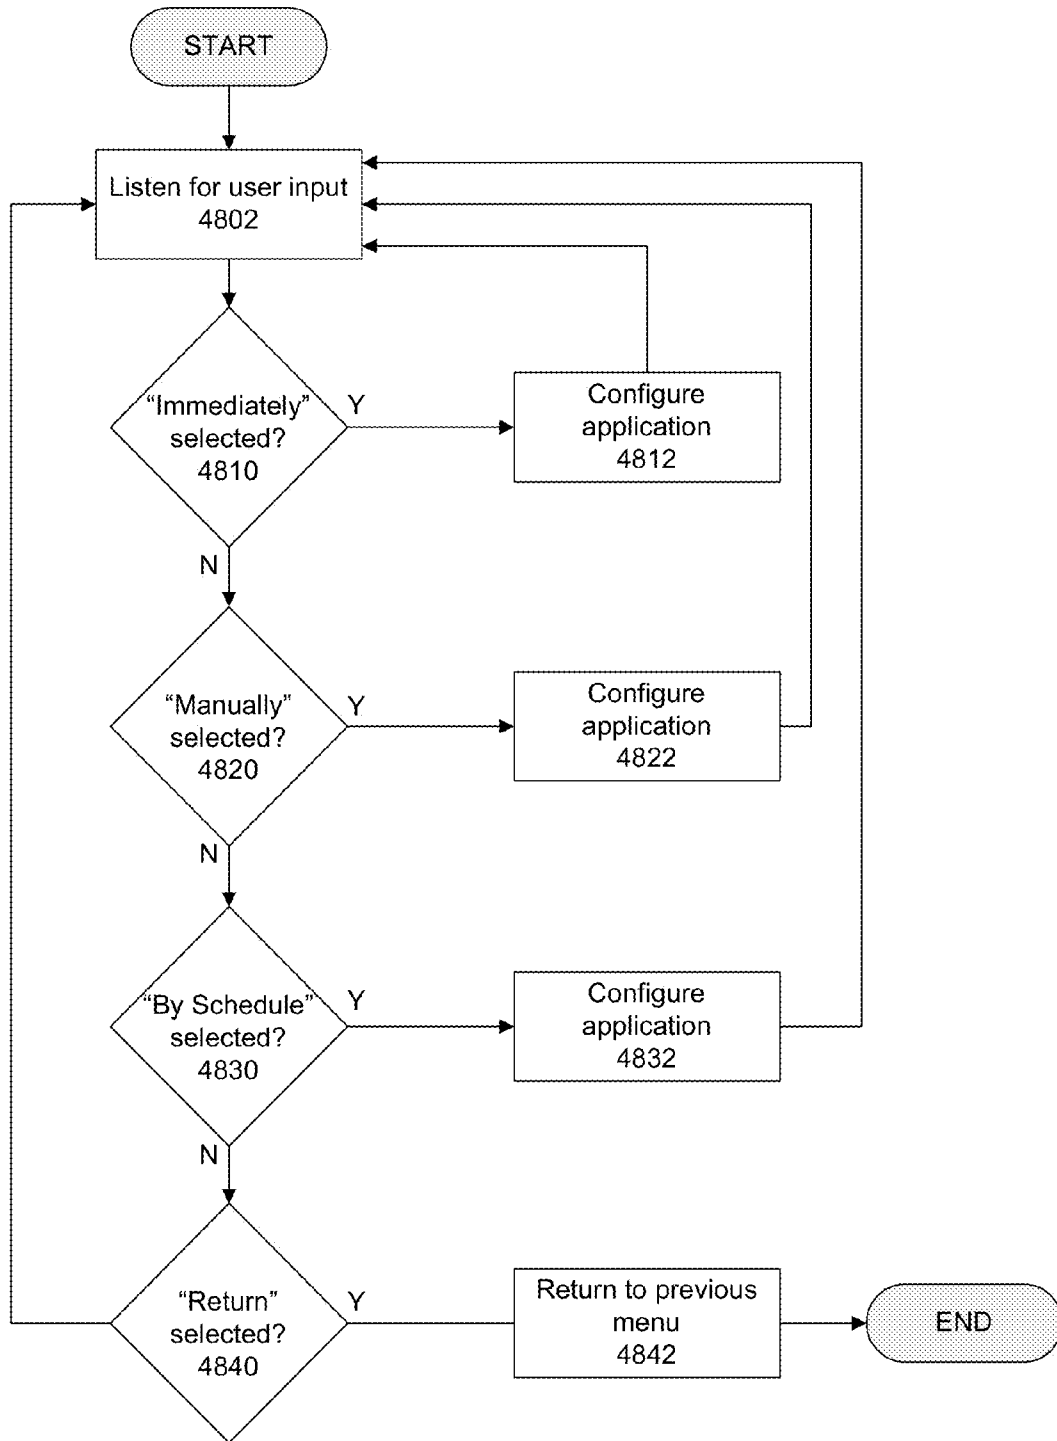


FIG. 48




FIG. 49

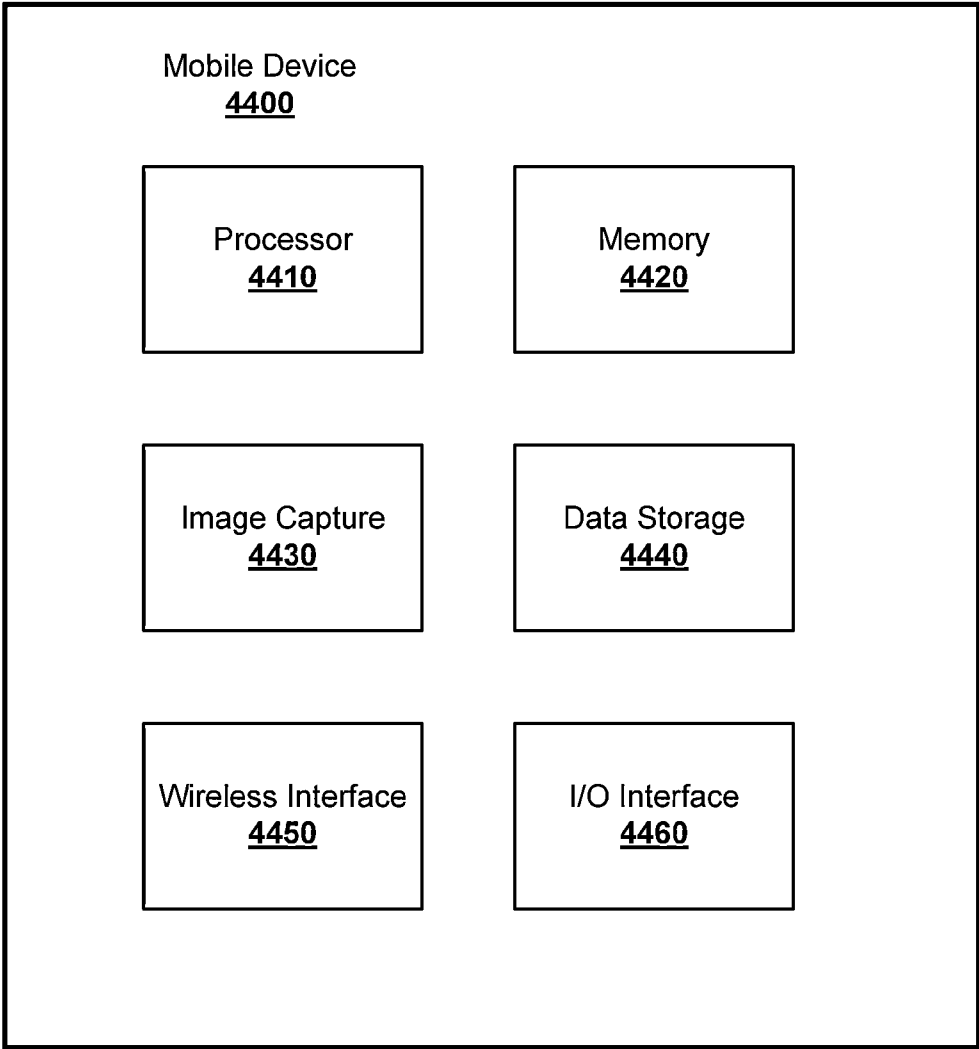


FIG. 50

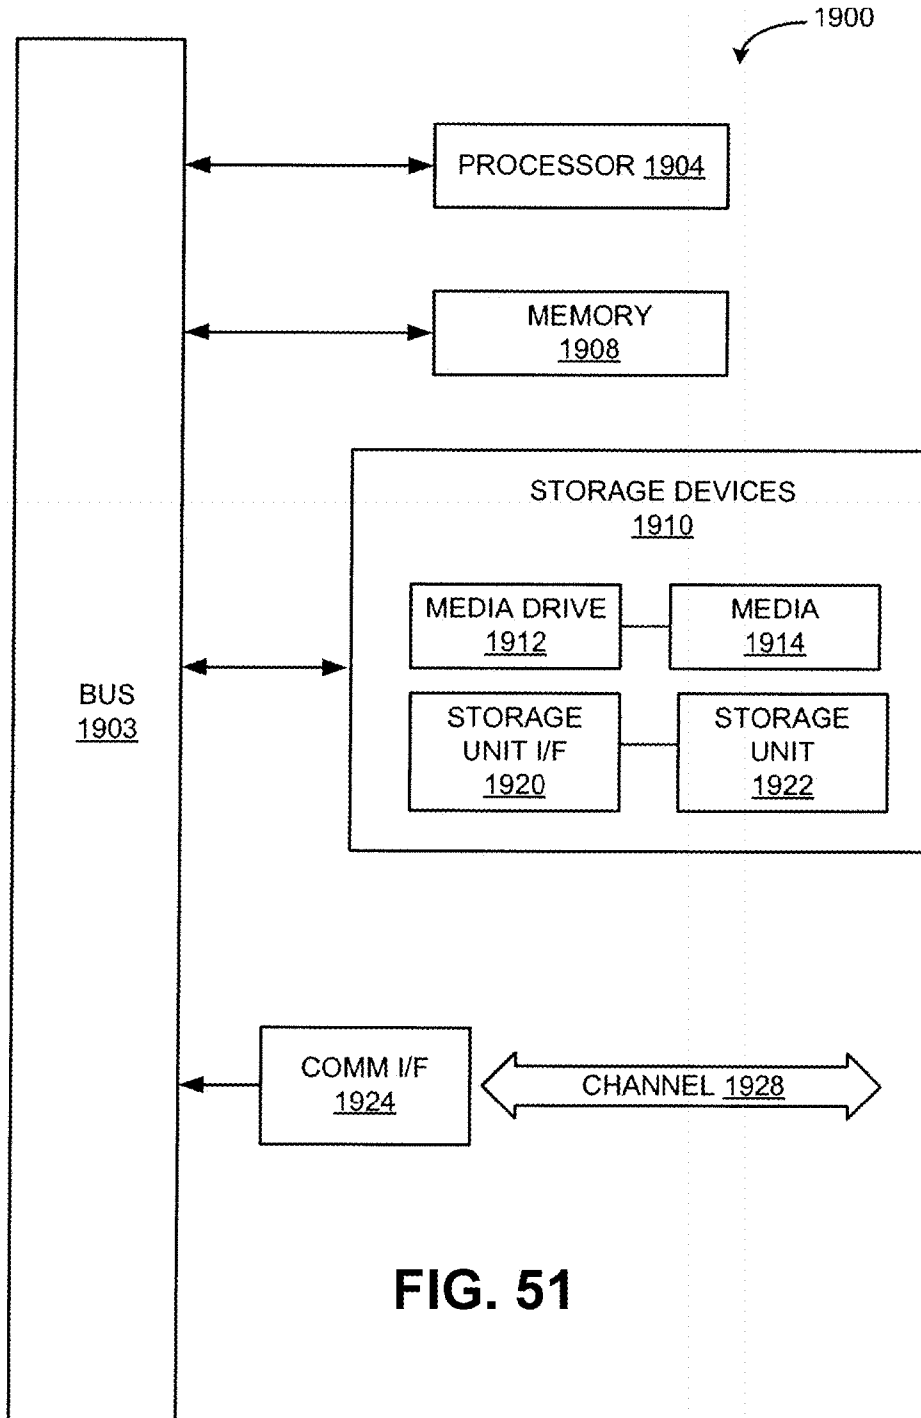


FIG. 51




FIG. 52




FIG. 53A




FIG. 53B




FIG. 53C



FIG. 53D




FIG. 54A

FIG. 54B

FIG. 54C

SYSTEMS AND METHODS FOR AUTOMATIC IMAGE CAPTURE ON A MOBILE DEVICE

RELATED APPLICATIONS INFORMATION

[0001] This application claims priority to U.S. patent application Ser. No. 13/461,726, filed May 1, 2012, which claims priority to U.S. Provisional Patent Application No. 61/561,772, filed Nov. 18, 2011, now pending; and which is a continuation in part of U.S. patent application Ser. No. 12/906,036 filed on Oct. 15, 2010, now pending, which itself is a continuation in part of U.S. patent application Ser. No. 12/778,943 filed on May 12, 2010, now pending, as well as a continuation in part of U.S. patent application Ser. No. 12/346,026 filed Dec. 30, 2008, now U.S. Pat. No. 7,978,900, which in turn claims the benefit of U.S. Provisional Application Ser. No. 61/022,279 filed Jan. 18, 2008, now expired, all of which are incorporated herein by reference in their entirety as if set forth in full. This application is also related to U.S. patent application Ser. No. 12/717,080 filed Mar. 3, 2010, which is now U.S. Pat. No. 7,778,457, which is incorporated herein by reference in its entirety as if set forth in full.

BACKGROUND

[0002] 1. Technical Field

[0003] The embodiments described herein generally relate to automatic capture of an image of a financial or other document on a mobile device, and more particularly to automatically detecting and determining image quality on the mobile device prior to capturing the image of the document.

[0004] 2. Related Art

[0005] Banks and other businesses have become increasingly interested in electronic processing of check and other documents in order to expedite processing of these documents. Users can scan a copy of the document using a scanner or copier to create an electronic copy of the document that can be processed instead of routing a hardcopy of the document from one place to another for processing. For example, some banks can process digital images of checks and extract check information from the image needed to process the check without requiring that the physical check be routed throughout the bank for processing.

[0006] Mobile devices that incorporate cameras have also become ubiquitous. However, the quality of images captured varies greatly, and many factors cause images captured using a mobile device to be of poor quality. Therefore, images captured by mobile devices are often not of sufficiently high quality to be used for electronic processing of documents. For systems which utilize images of documents captured on the mobile device, the process of evaluating a mobile image to determine if it is of sufficient image quality can be time consuming and cumbersome for the user of the mobile device. Therefore, it would be advantageous to streamline and automate the process of capturing images of documents with mobile devices and verifying that the quality of the image is sufficient processing.

SUMMARY

[0007] Systems and methods are provided for real-time evaluation and enhancement of image quality prior to capturing an image of a document on a mobile device. An image capture process is initiated on a mobile device during which a user of the mobile device prepares to capture the image of the document, utilizing hardware and software on the mobile

device to measure and achieve optimal parameters for image capture. Feedback may be provided to a user of the mobile device to instruct the user on how to manually optimize certain parameters relating to image quality, such as the angle, motion and distance of the mobile device from the document. When the optimal parameters for image capture of the document are achieved, at least one image of the document is automatically captured by the mobile device.

[0008] Other features and advantages of the present invention should become apparent from the following description of the preferred embodiments, taken in conjunction with the accompanying drawings, which illustrate, by way of example, the principles of the invention.

BRIEF DESCRIPTION OF THE DRAWINGS

[0009] The various embodiments provided herein are described in detail with reference to the following figures. The drawings are provided for purposes of illustration only and merely depict typical or example embodiments. These drawings are provided to facilitate the reader's understanding of the invention and shall not be considered limiting of the breadth, scope, or applicability of the embodiments. It should be noted that for clarity and ease of illustration these drawings are not necessarily made to scale.

[0010] FIG. 1 is a block diagram which illustrates one embodiment of a system for mobile image capture and remittance processing, according to one embodiment of the invention.

[0011] FIG. 2 illustrates one embodiment of a method of automatic image capture on a mobile device, according to one embodiment of the invention.

[0012] FIG. 3 illustrates one embodiment of a graphical user interface (GUI) which may be presented to the user on a display screen of the mobile device during a mobile image capture process, according to one embodiment of the invention.

[0013] FIGS. 4A-4H illustrate visual feedback displayed on the display screen of the mobile device during the mobile image capture process, according to one embodiment of the invention.

[0014] FIG. 5 is a flowchart illustrating a biller lookup process, according to one embodiment of the invention.

[0015] FIG. 6 is a flow diagram illustrating a process for a second data recognition process on a remittance coupon, according to one embodiment of the invention.

[0016] FIGS. 7A and 7B are images of remittance coupons which illustrate the results of a first data recognition process and a second data recognition process, according to one embodiment of the invention.

[0017] FIG. 8 is an image of a remittance coupon captured by a mobile device.

[0018] FIG. 9 is a geometrically corrected image created using image processing techniques disclosed herein using the mobile image of the remittance coupon illustrated in FIG. 8.

[0019] FIG. 10 and its related description above provide some examples of how a perspective transformation can be constructed for a quadrangle defined by the corners A, B, C, and D according to an embodiment.

[0020] FIG. 11 is a diagram illustrating an example original image, focus rectangle and document quadrangle ABCD in accordance with the example of FIG. 10.

[0021] FIG. 12 is a flow diagram illustrating a method for correcting defects to mobile image according to an embodiment.

[0022] FIG. 13 is a flow chart for a method that can be used to identify the corners of the remittance coupon in a color image according to an embodiment.

[0023] FIG. 14 is a flow diagram of a method for generating a bi-tonal image according to an embodiment.

[0024] FIG. 15 illustrates a binarized image of a remittance coupon generated from the geometrically corrected remittance coupon image illustrated in FIG. 9, according to one embodiment.

[0025] FIG. 16 is a flow diagram of a method for converting a document image into a smaller color icon image according to an embodiment.

[0026] FIG. 17A is a mobile image of a check according to an embodiment.

[0027] FIG. 17B is an example of a color icon image generated using the method of FIG. 12 on the example mobile image of a check illustrated in FIG. 13A according to an embodiment.

[0028] FIG. 18 is a flow diagram of a method for reducing the color depth of an image according to an embodiment.

[0029] FIG. 19A depicts an example of the color "icon" image of FIG. 17B after operation 1302 has divided it into a 3x3 grid in accordance with one embodiment of the invention.

[0030] FIG. 19B depicts an example of the color "icon" image of FIG. 17B converted to a gray "icon" image using the method illustrated in FIG. 18 according to an embodiment.

[0031] FIG. 20 is a flowchart illustrating an example method for finding document corners from a gray "icon" image containing a document according to an embodiment.

[0032] FIG. 21 is a flowchart that illustrates an example method for geometric correction according to an embodiment.

[0033] FIG. 22A is an image illustrating a mobile image of a check that is oriented in landscape orientation according to an embodiment.

[0034] FIG. 22B example gray-scale image of the document depicted in FIG. 17A once a geometrical correction operation has been applied to the image according to an embodiment.

[0035] FIG. 23 is a flow chart illustrating a method for correcting landscape orientation of a document image according to an embodiment.

[0036] FIG. 24 provides a flowchart illustrating an example method for size correction of an image according to an embodiment.

[0037] FIG. 25 illustrates a mobile document image processing engine (MDIPE) module for performing quality assurance testing on mobile document images according to an embodiment.

[0038] FIG. 26 is a flow diagram of a process for performing mobile image quality assurance on an image captured by a mobile device according to an embodiment.

[0039] FIG. 27 is a flow diagram of a process for performing mobile image quality assurance on an image of a check captured by a mobile device according to an embodiment.

[0040] FIG. 28A illustrates a mobile image where the document captured in the mobile document image exhibits view distortion.

[0041] FIG. 28B illustrates an example of a grayscale geometrically corrected subimage generated from the distorted image in FIG. 28A according to an embodiment.

[0042] FIG. 29A illustrates an example of an in-focus mobile document image.

[0043] FIG. 29B illustrates an example of an out of focus document.

[0044] FIG. 30 illustrates an example of a shadowed document.

[0045] FIG. 31 illustrates an example of a grayscale snippet generated from a mobile document image of a check where the contrast of the image is very low according to an embodiment.

[0046] FIG. 32 illustrates a method for executing a Contrast IQA Test according to an embodiment.

[0047] FIG. 33A is an example of a mobile document image that includes a check that exhibits significant planar skew according to an embodiment.

[0048] FIG. 33B illustrates an example of a document subimage that exhibits view skew according to an embodiment.

[0049] FIG. 34 is a flow chart illustrating a method for testing for view skew according to an embodiment.

[0050] FIG. 35 illustrates an example of a mobile document image that features an image of a document where one of the corners of the document has been cut off in the picture.

[0051] FIG. 36 illustrates a Cut-Off Corner Test that can be used for testing whether corners of a document in a document subimage have been cut off when the document was imaged according to an embodiment.

[0052] FIG. 37 illustrates an example of a mobile document image that features a document where one of the ends of the document has been cut off in the image.

[0053] FIG. 38 is a flow diagram of a method for determining whether one or more sides of the document are cut off in the document subimage according to an embodiment.

[0054] FIG. 39 illustrates an example of a mobile document image where the document is warped according to an embodiment.

[0055] FIG. 40 is a flow diagram of a method for identifying a warped image and for scoring the image based on how badly the document subimage is warped according to an embodiment.

[0056] FIG. 41 illustrates an example of a document subimage within a mobile document image that is relatively small in comparison to the overall size of the mobile document image according to an embodiment.

[0057] FIG. 42 is a flow diagram of a process that for performing an Image Size Test on a subimage according to an embodiment.

[0058] FIG. 43 is a flow chart of a method for executing a code line test according to an embodiment.

[0059] FIG. 44 illustrates a method for executing an Aspect Ratio Test according to an embodiment.

[0060] FIG. 45 is a flow chart of a method for processing an image using form identification according to an embodiment.

[0061] FIG. 46 is a flow chart of a method for processing an image using dynamic data capture according to an embodiment.

[0062] FIG. 47 is a flow diagram illustrating an exemplary method of configuring a recurring payment schedule according to an embodiment.

[0063] FIG. 48 is a flow diagram illustrating an exemplary method of selecting a specific scheduling preference according to an embodiment.

[0064] FIG. 49 is a flow diagram illustrating an exemplary method of enabling a user to set one or more reminders associated with a recurring bill payment according to an embodiment.

[0065] FIG. 50 is a block diagram of various functional elements of a mobile device that can be used with the various systems and methods described herein according to an embodiment.

[0066] FIG. 51 is a block diagram of functional elements of a computer system that can be used to implement the mobile device and/or the servers described in the systems and methods disclosed herein.

[0067] FIG. 52 is a flow diagram of a process for edge and feature detection, according to an embodiment.

[0068] FIGS. 53A-53D illustrate images being processed using the edge and feature detection methodologies, according to an embodiment.

[0069] FIGS. 54A-54C illustrate a template-matching method used during feature detection, according to an embodiment.

DETAILED DESCRIPTION

[0070] The embodiments described herein are directed to real-time evaluation and enhancement of image quality prior to capturing an image of a document on a mobile device. An image capture process is initiated on a mobile device during which a user of the mobile device prepares to capture the image of the document, utilizing hardware and software on the mobile device to measure and achieve optimal parameters for image capture. Feedback may be provided to a user of the mobile device to instruct the user on how to manually optimize certain parameters relating to image quality, such as the angle, motion and distance of the mobile device from the document. When the optimal parameters for image capture of the document are achieved, at least one image of the document is automatically captured by the mobile device.

[0071] Automatic capture of an image increases the image quality of the captured image—and the document contained within that image—for many reasons. First, by measuring multiple parameters in real-time, the mobile device will know when the image should be taken in order to maximize the quality of the image. Second, by automatically capturing one or more images when the optimal parameters are achieved, the user of the mobile device is not required to manually press a button on the mobile device that may cause movement of the mobile device and degrade the image quality of the image being captured.

[0072] The capturing and processing of an image of a document with a mobile device may be used for processing a financial transaction using the mobile device, such as depositing a check, paying a bill or transferring money between different bank and credit accounts. The mobile device includes systems and methods for determining a plurality of parameters of the mobile device which affect the quality of an image and automatically capturing one or more images of the document when the parameters fall within acceptable ranges. Other features for ensuring the capture of a high-quality image may be carried out in the form of software running on the mobile device which provides tools for detecting features of the document in real-time prior to image capture and for instructing a user of the mobile device on how to improve the quality of the captured image.

[0073] The embodiments described herein provide for detecting parameters on a mobile device related to image quality in real-time and automatically capturing an image of a document via a camera on the mobile device that is used to extract information for performing an electronic financial transaction.

[0074] The processes described herein utilize the high performance Graphical Processing Units (GPUs) on smartphones and other mobile devices to ensure a high quality capture of an image within a video frame. In addition to the previously mentioned settings that exist at the server, mobile device-based technologies use frame capture and processing techniques that provide a breakthrough in identifying high quality images suitable for OCR/ICR post-processing on the server. This breakthrough adds a significant increase in the accuracy of data recognition of the document image.

[0075] In one embodiment of a method of video-based mobile image capture, as the user prepares to take a photo of a document, an image capture application is switched into a video mode, and video frames are immediately captured (using either the available device APIs or the OpenCV API) and saved in a processing buffer. As the video frames are captured, selected frames are pre-processed on the mobile device to determine the image's suitability for OCR post-processing on the server. This pre-processing is a quick analysis of the image quality to evaluate focus, exposure, contrast, presence of color, reflection, and other criteria as defined by customized and dynamic settings resident on the mobile device or received from the server. Frames that do not meet the criteria are quickly discarded, and another frame is then selected for pre-processing. This pre-processing continues until an acceptable frame is found, at which time the video stream is stopped and the user receives a message that the image capture process is complete.

[0076] Once a suitable video frame is identified, deeper processing continues of the video frames in close proximity to the identified frame, and the best video frame is identified. This best video frame may be combined with other nearby frames to create an even better composite image.

[0077] In real-time, the end-user receives feedback on the quality of the image. This feedback includes:

[0078] 1. Real-time framing around the 'edge' of the document. In one embodiment, a colored line is drawn around the edge of the document to form a frame. For example, the frame color is red if the document is too far away or too close, or if the shape of the document (parallelogram or quadrilateral) indicates the video camera angle is too steep.

[0079] 2. Image crispness feedback on the crispness of the image. This feedback may be presented through messages appearing at the bottom of the screen in a semi-transparent font indicating whether the user is moving the mobile device too much. Crispness is determined via edge detection by looking at how 'crisp' the edges of the document are. The user may be given real-time feedback as they hold the camera steadier, by changing the message or changing the color of the real-time frame.

[0080] 3. Color Indicators. As the user moves the mobile device over the document, a color indicator changes from yellow or red to green when the document is properly placed within the frame and edge detection determines that the document is in focus. Once the green mode is detected for a certain period of time, the current video frame is captured and displayed as a still image to the user, along with a message indicating that a frame (or several frames together) have been captured and uploaded to the server.

[0081] 4. Contrast Indicators. If the contrast of the image is too low, a message is displayed on the mobile device screen which will alert the user to this fact.

[0082] 5. Reflection Indicators. For highly reflective documents such as driver's licenses, an oval or other shape may appear around any reflective areas of the document, and a message will be displayed to the user that they need to change the lighting conditions or perspective of the mobile device with respect to the document to remove the reflective area.

[0083] In addition, the image can optionally be pre-cropped on the mobile device via edge detection so as to reduce the size of the actual image uploaded to the server.

[0084] Using the techniques discussed herein, the mobile device will also be capable of capturing and pre-processing video frames in such a manner as to help uniquely identify the document that is the subject of the image in the video frame. For example, techniques will be employed to identify the document as a driver's license, bank check or credit card bill. Settings, as discussed herein may be received at the mobile device from a remote server in order to guide some of the processing parameters in determining the document type at the mobile device.

[0085] Pre-processing of the video frames will include evaluation of focus, exposure, contrast, etc. as well as identification of document features that will help uniquely identify the category of the document. These document features may be areas of the document that are compared to known document types of known entities in order to find a match. Again, the processing to be performed and the criteria may be defined by the settings received from a remote server. As described above with regard to image capture, once a suitable video frame is identified, deeper processing of the video frames continues in close proximity to the identified frame, and the best video frame is identified. This best video frame may be combined with other nearby frames to create an even better composite image. Additional processing may take place that might include edge detection, image cropping, and compression in order to make the smallest payload possible for submission to the remote server for post-processing.

[0086] FIG. 1 illustrates one embodiment of a system **100** for mobile image capture and remittance processing. The system **100** includes a mobile device **102**, such as a cellular phone, smartphone, tablet, personal digital assistant (PDA) or other portable electronic device that may be connected with a communications network. The mobile device **102** will include an image capture device (not shown), such as a digital camera or a portable scanning device, that uses an image sensor to capture an image of a document. The mobile device **102** is connected with a remote server **104** over a network so that the mobile device **102** can transmit captured images or image data to the remote server **104**. In one embodiment, the remote server **104** may send information to the mobile device **102** (and specifically an application running on the mobile device) regarding the parameters that should be measured and the values of the thresholds required to capture an image. The remote server **104** may perform additional image processing and data extraction, as will be described in further detail below, in order to determine information about the document and identify the appropriate parties and amounts. In one embodiment, the remote server **104** may be connected with an address database **106** which is used to verify address information obtained from the remittance coupon, as will be described in further detail below. The remote server **104** may

also be connected with a biller database **108** which stores information on billers, such as address information and billing formats for the remittance coupons. Once the remote server **104** has extracted and identified all of the relevant data from the image of the remittance coupon, the extracted data and the captured and processed images may be stored in a content database **110** connected with the remote server **104**. The extracted data may then be transmitted to a banking server **112** for processing the payment from a bank account belonging to the user of the mobile device **102**. The extracted data may also be first sent back to the mobile device **102** to display the data to a user for confirmation before a bill is paid.

[0087] The mobile device can comprise a mobile telephone handset, Personal Digital Assistant, or other mobile communication device. The mobile device can include a camera or other imaging device, such as a scanner, or might include functionality that allows it to connect to a camera or other imaging device. The connection to an external camera or other imaging device can comprise a wired or wireless connection. In this way the mobile device can connect to an external camera or other imaging device and receive images from the camera or other imaging device.

[0088] Images of the documents taken using the mobile device or downloaded to the mobile device can be transmitted the remote server via a network. The network can comprise one or more wireless and/or wired network connections. For example, in some cases, the images can be transmitted over a mobile communication device network, such as a code division multiple access ("CDMA") telephone network, or other mobile telephone network. The network can also comprise one or more connections across the Internet. Images taken using, for example, a mobile device's camera, can be 24 bit per pixel (24 bit/pixel) JPG images. It will be understood, however, that many other types of images might also be taken using different cameras, mobile devices, etc.

[0089] The remote server can be configured to perform various image processing techniques on images of remittance coupons, checks, or other documents captured by the mobile device. The remote server can also be configured to perform various image quality assurance tests on images of remittance coupons or documents captured by the mobile device to ensure that the quality of the captured images is sufficient to enable remittance processing to be performed using the images. Examples of various processing techniques and testing techniques that can be implemented on the remote server are described in detail below.

[0090] According to an embodiment, the remote server can be configured to communicate to one or more bank servers via the network. The bank server can be configured to process payments in some embodiments. For example, in some embodiments, mobile device can be used to capture an image of a remittance coupon and an image of a check that can be used to make an electronic payment of the remittance payment. For example, the remote server can be configured to receive an image of a remittance coupon and an image of a check from the mobile device. The bank server can electronically deposit the check into a bank account associated with the entity for which the electronic remittance is being performed (payor). According to some embodiments, the bank server and the remote server can be implemented on the same server or same set of servers.

[0091] In other embodiments, the remote server can handle payment. For example, the remote server can be operated by or on behalf of an entity associated with the coupon of FIG. 8,

such as a utility or business. The user's account can then be linked with a bank, Paypal®, or other account, such that when remote server receives the remittance information, it can charge the appropriate amount to the user's account.

I. Mobile Device Automatic Capture

[0092] In one embodiment, the mobile device is configured to automatically capture an image of the document when certain parameters are met. Real-time analysis of various position settings of the mobile device, image sensor and surrounding environment is performed to ensure that the captured image is as in-focus as possible. Automatic capture allows for the mobile device to be held over the document without the user having to press a button. These position settings may be standardized for all image capture processing or dynamically adjusted based on the type of mobile device and image sensor, the type of document being captured or even the ambient environment of the mobile device.

[0093] The capturing of the image of the document is the first step of an end-to-end solution for processing documents using mobile device cameras, which can be utilized to provide the user with tools and information to improve the quality of the image and decrease the chance of errors from poor image quality. However, by having the mobile device carry out several image processing steps, the overall user experience may be improved, due to the fact that the image of the document which is eventually sent to the remote server will be of substantially higher quality. A higher quality image means it is much less likely that the image will be rejected by the remote server, which would otherwise require the user to capture another image of the document. By running an image capture application on the mobile device, problems with the captured image can be immediately identified and corrected without waiting for transmission of the image to the remote server, analysis at the server, and feedback from the remote server to the user.

[0094] Furthermore, by measuring multiple parameters related to the quality of the image in real-time prior to capturing the image, the quality of the captured image will be significantly increased, the likelihood that the server rejects the image will be reduced, the accuracy of processed information from the image will increase, and the likelihood of needing complex image processing performed will be reduced.

[0095] The parameters being measured on the mobile device are related to the ability of the device, and more particularly the ability of the camera on the mobile device, to capture a high quality image of the document which can be used to accurately extract text and other content from the document that is needed for a financial transaction. Some financial transactions which use digitally-captured images of documents require certain levels of image quality in order to verify the authenticity of the document and ensure that the content of the document—account numbers, addresses, names, monetary values, etc.—is accurately extracted. By utilizing technology within a mobile device to help the user capture a high quality image of the document, the amount of post-capture image processing can be reduced, as is the likelihood of having to reject a captured image and force the user to recapture.

[0096] In one embodiment, the parameters are measured in real time, so that an application running on the mobile device can continuously monitor the image quality settings and instruct the camera to capture one or more images of the

document when the parameters meet certain thresholds of image quality. In one embodiment, the thresholds of image quality may be defined individually for each parameter being measured, where a measurement for each parameter must meet the defined threshold before the image can be captured. In another embodiment, the parameter values may be combined into a group threshold or total overall quality score, such that the combination of all of the measured parameter values must produce an overall quality score which exceeds a defined threshold value before the image can be captured by the camera.

[0097] In a further embodiment, some parameters may be weighed against each other when determining whether to capture the image. A threshold value of one parameter may be adjusted based on the actual value of another parameter—for example if a first parameter with a high value allows for a second parameter to have a low value but still provide a high quality image, the system will adjust the threshold value of the second parameter to a lower value so that an image will still be captured if the first parameter exceeds its threshold value. For example, if the orientation angle falls within 10 degrees of normal while the degree of motion is negligible, the application may determine that an image should be captured if the lack of motion will provide sufficient image quality on its own.

[0098] In a further embodiment, additional parameters may be compared on the captured image after image capture and before being sent to the server. In many cameras, the quality of the real time images, commonly called preview frames or video frames, are of lower quality than the final captured image, where the width or height in pixels is less than the captured image, for various reasons such as frame rate capability. If so, the parameter comparisons of the video frame may require different image quality thresholds based on smaller total pixels or less pixel density in terms of dots per inch (DPI). Hence, the quality thresholds for video frames versus captured image will be different.

[0099] In a further embodiment, some or all quality parameters may not be known when capturing a document. In this case, processing of the preview frames can choose the best preview image from a plurality of preview images over a time period. The processor can store quality measurements and compute quality statistics to determine a quality threshold for a second pass or a sequence of passes.

[0100] In a further embodiment, in the case of high quality preview frames, frames with a sufficient DPI for the type of document being captured, the best quality preview frame passing quality measurements in a sequence of preview frames can be stored temporarily. If, after some variable time, there are no better quality measurements of other preview frames, then the stored preview frame with the best quality is chosen as the image to send for server processing.

[0101] In a further embodiment, automated image quality can be estimated by cooperation with the camera hardware and camera focusing mechanisms. A typical camera will indicate whether camera focusing is in progress, and during that time the preview frame quality is suspect. Once the focusing is not in progress, the camera is likely in sharp focus. If other parameters are passing their thresholds, near the time the camera stops focusing, then the processing might ignore the focus threshold comparison and capture the image. This captured image can then be sent for server processing or compared for final quality before sending.

[0102] In a further embodiment, the mobile processor may interact with the mobile camera features such as focusing, contrast, shutter period, and aperture to iteratively change the real-time preview frames until parameters pass thresholds and image capture commences.

[0103] In a further embodiment, if the preview frame passes its thresholds, but the captured image does not pass its thresholds, the processor can restart processing of the preview frames. This restarting can continue as long as necessary until the captured image passes its quality measurements.

[0104] In a further embodiment, if the image quality does not pass thresholds after a period of time, the thresholds can be adjusted and the preview processing restarted. This can happen due to the wide variety and quality of image capture devices which may not be capable of capturing images which meet the initial thresholds.

[0105] In another embodiment, if the image quality does not pass thresholds after a period of time, the reason may be due to user error, such as an obstruction on the camera lens, whether a finger is over the lens or the lens is dirty. Automatic instructions can be displayed to instruct to clear the lens.

[0106] In a further embodiment, the processor can aggregate quality measurements over different parts of the image to determine if quality thresholds have been achieved overall. For instance, sharpness thresholds may be critical only in some parts of the image, and unimportant in others, such as pictures of a face which may be of no use in server side processing.

[0107] In a further embodiment, the processor can compute quality measurements over different parts of the preview frame or captured image to determine likelihood of image quality prior to server processing. For example, a camera may focus on one side of a document if not all of the document is the same distance from the camera. In such cases, one side of a document may be in focus, and another side out-of-focus. The processor can determine that a significant portion is out of focus and fail a quality measurement.

[0108] In a further embodiment, automated quality measurement may not ensure an image is captured in some time period. In such a case, the processor can stop automated analysis and give camera control to the mobile device user and ask the user to attempt to capture the image manually, either through the mobile device's built in image capture software application or through the processor.

[0109] FIG. 2 illustrates one embodiment of a method of automatic mobile image capture and remittance processing, as will be further described herein. The workflow of the methods described herein start with the mobile device, the steps of which are illustrated on the left portion of FIG. 2. In a first step S202, a mobile application is initialized on the mobile device. The mobile application may be a software application stored on the mobile device which is configured to interface with the image capture device (camera) on the mobile device and communicate with the remote server to send images or data relating to the images and receive information on the parameters and parameter values to use when capturing an image. The mobile application will activate the mobile device camera and begin receiving images from the camera in real-time. The mobile application may also display one or more graphical user interfaces (GUIs) on a display screen of the mobile device that provides instructions and feedback to a user during the mobile image capture process. In one embodiment, the mobile application will display the images being received from the camera on the display screen

in real-time as part of the GUI (see FIG. 3), so that the user can see the image being viewed by the camera in real time and more easily understand the adjustments that may need to be made.

[0110] In a next step S204, at least one parameter related to the image quality of the images being captured by the camera is measured. The parameters may be measured using hardware and/or software on the mobile device, such as an accelerometer, the image sensor in the camera, or image processing software modules running in the mobile application. A parameter value is determined for each measured parameter, such as angle of the mobile device, image size (or document size within the image), lighting, focus, etc. Further details regarding the parameters and the hardware and/or software used to measure those parameters are described further below. In step S206, the measured values of the parameters are compared with their respective threshold values (or single group threshold value) and a determination is made as to whether the parameters meet the thresholds.

[0111] In step S208, if a certain parameter value has not met its threshold, the user of the mobile device may be provided with feedback indicating which parameter is deficient or instead providing a specific instruction to the user as to how the parameter can be adjusted so that it will meet the threshold value. In another embodiment, the user feedback may simply affirm that the threshold values have been met, and that the camera will now capture one or more images of the document. Step S208 may be optional if, for example, all of the necessary parameters have met their thresholds and feedback to the user is unnecessary, or if one of the parameters which has not met the threshold cannot be adjusted by the user. One or more of the parameters may require the mobile application to adjust the parameters automatically without user input or feedback, such as when the parameter relates to an image processing setting within the image processing software running on the mobile device, or such as when image processing quality is insufficient in a variable amount of time.

[0112] In step S210, one or more images of the document are then captured by the camera and stored on the mobile device. In one embodiment, a plurality of images is captured in a short period of time to ensure that at least one of the images is of high quality. The plurality of images may be captured in sequence—at several frames per second—so that the images will look substantially identical but be slightly different if the user is moving the mobile device as the images are being captured. In one embodiment described further below, the plurality of images may be processed to create a single high quality image using the highest quality aspects of each image. In another embodiment, the camera may be set into a video mode so that it captures a high rate of images—for example 24 frames per second—which may be analyzed individually or together in order to generate a high quality image of the document.

[0113] In step S212, the captured images may be processed at the mobile device to perform image quality analysis (IQA) and determine if the captured images are of sufficient image quality to be used for the financial transaction requested by the user. If the images are of sufficient quality, the processing step S212 may be skipped. However, the captured images may undergo post-capture image processing, to correct skew, warping, orientation, contrast, shadow, lighting, focus, etc. (all of which are described in further detail below), in order to ensure that the content of the document can be extracted from the image. The image or images may also be processed in

accordance with previously-described image processing steps to binarize the image, create a grayscale image or otherwise clarify the content of the document within the image. Edge detection and focus algorithms may be used to determine whether all four sides of the document are within the captured image, whether the angle of the image capture device is within an acceptable range relative to the remittance coupon and whether the size of the image is too small. For example, the mobile device can be configured to convert the captured image from a color image to a grayscale image or to bitonal image, identify the corners of the remittance coupon, and to perform geometric corrections and/or warping corrections to correct defects in the mobile image.

[0114] Once an image is captured which meets the required parameters, in step S214, the image may then be re-sized, compressed, encrypted and converted to a base-64 format before being uploaded to the remote server connected with the mobile device over a network using a secure socket layer (SSL) connection. The image may be re-sized due to the fact that some mobile cameras capture images with file sizes of up to 8 megapixels (MP), and so re-sizing and compressing these images allows for faster upload time.

[0115] The content of the document may be extracted during this step as well to create data relating to the image, which may be useful if the image processing on the mobile device is sufficient to extract the contents and avoid having to transmit the captured image or images from the mobile device to the remote server. In step S216, the processed (or un-processed) images or the data extracted from the image may be transmitted to the remote server. Once at the remote server, additional processing of the image may be performed to verify the quality of the document and the reliability of its content. If the content was already extracted at the mobile device, the data may be further processed to verify its accuracy. Once the data from the image has been extracted, the financial transaction desired by the user may be initiated (S218).

[0116] Additional details of the various image pre-processing techniques are described in further detail, below. The pre-processing of the images at the mobile device also allows for a further round of image quality assurance (IQA) beyond the measurement of the parameters prior to capturing the images. If the initial processing of the image identifies problems with the image, the user may be provided with feedback to request that another image be taken, including requesting that settings relating to the image capture device or the positioning of the remittance coupon be altered.

[0117] FIG. 3 illustrates one embodiment of a graphical user interface (GUI) which may be presented to the user on a display screen 302 of a mobile device 300 during the mobile image capture process. The user launches an application on the mobile device 300 which activates the camera (not shown) and displays a real-time image 304 of the field-of-view of the camera. The real-time image 304 is therefore continuously updated with new images from the camera, such that movement of the mobile device 300 would correspond in real-time to movement of the image being detected by the camera. The user can therefore position the mobile device and the camera over a document 308 and immediately see a picture of the document on the GUI. This will aid the user in adjusting the parameters related to image quality while immediately seeing the effect of the adjustments.

[0118] According to one embodiment, the mobile device can also be configured to optionally receive additional information from the user. For example, in some embodiments, the

mobile device can be configured to prompt the user to enter data, such as a payment amount that represents an amount of the payment that the user wishes to make. The payment amount can differ from the account balance or minimum payment amount shown on the remittance coupon. For example, the remittance coupon might show an account balance of \$1000 and a minimum payment amount of \$100, but the user might enter a payment amount of \$400.

[0119] In a further embodiment, the mobile device can prompt the user to point to, or touch, a particular information field on the image being shown on the mobile device screen to indicate where that field is on the document. This may occur if a document information structure is unknown or variable for different printings. The image being shown on the mobile device may be real-time or the captured image.

[0120] In a further embodiment, the mobile device can prompt the user to point to, or touch, a sequence of information fields, one at a time, to indicate where those fields are on the document.

[0121] In a further embodiment, the mobile device can prompt the user to adjust a rectangle or polygon being displayed on the mobile device screen to drag it over an information field or fields to indicate where the field is on the document. The adjustment may be a drag gesture, a resize gesture, a rotation gesture, or other gestures common to touch screens in order to crop the rectangle or polygon over the information field.

II. Parameters

Angle

[0122] In one embodiment, one of the parameters being measured is an angle at which the mobile device is oriented with respect to the document. The angle is measured based on the mobile device camera being positioned parallel to, or directly facing, the document when the document is placed on a flat, level surface. An accelerometer and gyroscope present on the mobile device measure the orientation and movement of the mobile device. In one embodiment, a degree of orientation of approximately 5 degrees is set as a maximum threshold, such that the mobile device would permit automatic capture of the document if the degree of variance from the parallel orientation is approximately equal to or less than 5 degrees. In another embodiment, the threshold of variation in the angle of orientation is approximately 2 degrees. By limiting the variance of the angle of orientation of the camera, the amount of perspective distortion, warping and other image defects will be minimized. Furthermore, by using the gyroscope in the mobile device, the orientation of the mobile device can be automatically determined by the mobile application running on the mobile device. The user can then be provided with feedback as to whether the orientation of the mobile device is adequate or needs to be corrected. Once the orientation falls within the acceptable threshold, the application may instruct the camera to immediately capture an image without requiring the user to manually depress a button or other input function. The automatic capture avoids introducing additional orientation distortion and other disturbances that may occur when the user must depress a button on the device to capture an image.

Motion

[0123] In one embodiment, one of the parameters being measured is a length of time at which the mobile device

remains still or is not in motion. The degree of motion of the mobile device may be measured by an accelerometer or gyroscope in the mobile device. The application may set a threshold period of time for which the mobile device must remain still before triggering automatic capture, thereby decreasing the chance that the captured image will be blurred. The time period may be only a few milliseconds in order to quickly capture the image at the instant that the phone stops moving, thereby requiring the user to hold the phone motionless for as little time as possible. In one embodiment, the time period may be approximately 80 milliseconds (ms), although in another embodiment, the time period may be approximately 400 ms.

[0124] In another embodiment, the shutter speed of the camera may be adjusted depending on the length of time that the mobile device is held still, so that the camera will capture an image in the time that the mobile device is primarily motionless.

[0125] In a further embodiment, the time period may be set to or beyond the length of the shutter speed on the camera to ensure the user has kept the camera steady during the shutter open time.

Viewfinder Bounding Box

[0126] In one embodiment illustrated in FIGS. 4A-4H, a user may be provided with a bounding box 404 in the form of a semi-transparent outline of a quadrilateral shape (such as a rectangle) on the display screen 402 of the mobile device to guide the user to make sure the document is fully contained in the display screen during the mobile image capture process. The bounding box 404 is generally designed to help the user align their mobile device over the document 406 to-be captured so that the document fits within the boundaries 408 of the bounding box. As shown in FIGS. 4A-4H, the bounding box may also be accompanied by a written instruction 410 which indicates one or more reasons why the view of the document by the camera is inadequate to capture the image.

[0127] The bounding box may also be presented in color combinations which represent how close the user is at achieving optimal image quality. For example, in FIG. 4A, the image of the document 406 (in this case a driver's license) is too large as a result of the camera being placed too close to the license. Therefore, the bounding box 404 may appear red, indicating that the document does not fit and that the auto-capture process cannot capture an image. The written instruction 410 above the bounding box simply states "Too Close," indicating to the user that the camera is too close to the driver's license and should be ported. The color of the bounding box 404 may vary dynamically as the user alters the camera or the document so that the document fits within the bounding box. When the document 406 is properly framed within the bounding box 404, the bounding box 404 may turn green to indicate that one or more of the parameters have been satisfied. Other visual or auditory aides may be provided to aid the user in aligning the camera on the mobile device with the document.

[0128] The bounding box may be useful to help the user fix numerous different parameters, such as image size, motion, edge detection and orientation angle, and may also be useful to help identify the type of document being captured based on the aspect ratio of the document dimensions. For example, FIGS. 4A-4H depict an image capture process for a driver's license, which often has a specific, standardized size and aspect ratio. Therefore, if the mobile application knows that

the user will be capturing an image of a driver's license (either through manual selection by the user or through automated process steps which require certain documents), the bounding box can be generated at the specific dimensions of the driver's license that will make it easier to capture the license with the camera and also to verify the authenticity of the license.

[0129] To further illustrate the benefits of the bounding box, FIG. 4B illustrates the display screen 402 of the mobile device when the document 406 is too far away from the camera. Thus, the bounding box 404 is colored in red and the written instruction 410 "too far away" is written up at the top. FIG. 4C illustrates an out-of-focus document 406, which provides the specific written instruction 410 "Out of Focus" to the user so that the user will pay more attention to focusing when taking a picture of the document. FIG. 4D illustrates an image of the driver's license 406 that is taken from an angle of orientation ("tilt") that is significantly greater than desired, thus causing distortion and skew in the images and making the driver's license very difficult to read. The written instruction 410 "Tilted Too Much" is then displayed on the screen. FIG. 4E illustrates the automatic capture process where the image is blurry as a result of too much movement of the camera or mobile device. The user may then see the red rectangle 404 along with the written instruction 410 "Hold Steady" which directly instructs the user on how to fix the problem. FIG. 4F is an illustration of a real-time image being displayed on the display screen 402 when there is insufficient light. As a result, the red rectangle 404 is displayed as the bounding box and the written instruction 410 "Not Enough Light" is displayed. In contrast, FIG. 4G illustrates an image capture process of the driver's license 406 where too much light has caused a large reflective area 412 to form on a left portion of the driver's license 406, making it difficult to read (for human or machine). Once again, written instructions 410 indicating "Too Bright" is displayed at the top, providing the user with the opportunity to adjust the lighting with the hope of fixing the problem. Finally, in FIG. 4H, a correctly proportioned driver's license 406 is presented, and as a result, a green bounding box 404 is displayed along with an affirmative written instruction 410 telling the user that the parameters have been properly aligned.

[0130] In a further embodiment, a bounding box can be replaced with a centered reticule, such as a rifle scope reticule, indicating the center of the display which the user would center over the document with movement.

[0131] In a further embodiment, an icon or image may be displayed in real-time above the detected document to give the user a visual indication of the offset from the mobile device display center. This icon or image appears to move above the detected document as the user moves the camera.

Automatic Aspect Ratio Correction

[0132] The outline may be provided in real-time during the image capture process to aid the user in capturing the entirety of the document at a correct aspect ratio. The rectangle represents the dimensions of the document being framed within the image, and guides the user in centering the document in the image so that the document is completely within a field of view of the image capture device. The user may be instructed to match the sides of the document with the sides of the rectangle, which will encourage the user to capture an image of the document that includes the entire document at an appropriate size and aspect ratio. The rectangle may have a specific width and height based on the type of document being

captured, such as a check, remit coupon, credit card, etc. The size of the document may be stored in a database on the phone or a remote server, and the user may be prompted to select the type of document that is being captured in advance so that the application can produce a rectangle of the appropriate dimension. If the type of document is known to have varying dimensions (such as a remittance coupon), the rectangle outline may be turned off.

[0133] In a further embodiment, the aspect ratio may be calculated from the real-time outline of the document. This calculated aspect ratio may then be applied to the rectangle framing the display, guiding the user to move the camera over the document into the rectangle or center the document underneath a centered reticule image.

Automatic Flash Detection

[0134] The application may also control the use of a flash on the mobile device to fire the flash in specific instances where the type of document or ambient lighting conditions requires the use of a flash. The use of the flash affects the lighting of the document and the shutter speed of the camera. A decision as to whether or not to fire the flash may be provided locally by analysis of the lighting conditions provided by the camera's image sensor or by parameters stored on a remote server or locally on the phone, such as information on the selected type of document to be captured. The remote server may also communicate with the mobile device to determine whether or not to fire the flash based on stored parameters such as the type of document or the specifications of the image sensor on the mobile device. The use of the flash usually requires a faster shutter speed on the camera and ensures more consistent lighting of the document. The faster shutter speed reduces the risk of motion blur as well, improving the quality of the image and the ability to read the content of the image using optical character recognition (OCR) and other image-processing steps described further herein.

[0135] In one embodiment, the flash may be turned off if the type of document is known to have reflectivity which would overexpose the image. A driver's license or credit card may be too reflective to allow for the use of the flash.

[0136] In one embodiment, the application on the phone may request a plurality of image capture settings from the remote server which will aid in controlling the flash, phone-based aspect ratio correction and other automatic capture settings described above.

Edge and Feature Detection

[0137] Edge detection at the mobile device also allows for filtering of images that have a high likelihood of being sub-quality, and allows, in real-time, the ability to indicate to the user various reasons for altering the phone's parameters or capturing the image again. Edge detection may be used to identify the borders of a document within the captured image and to determine the quality of the captured image. Edge detection capability may run on the mobile device, using its graphical and processing CPU's. The capability allows the detection of, and rejection of images with one or more of the above issues, based on if, and where, the edges have been found, and their position and relationships within the image.

[0138] Edge detection may be used to determine whether all four corners and all four sides of a document are within the captured image. In order to identify the borders of a document, one embodiment of edge detection may be performed

using document snippet detection, as described in U.S. Pat. No. 8,000,514, the contents of which are incorporated herein by reference in their entirety. The first step compresses the mobile image in such a way that some or all intra-document edges are suppressed whereas majority of document-to-background edges remain strong. This step makes the edge detection faster and, on documents with no large high-contrast internal areas (such as checks and remittance coupons), helps to avoid false positive edges. Second step finds edge "primitives", which are linear or piecewise linear segments separating high-contrast areas within the compressed mobile image. Such "primitives" are classified into left/top/right and bottom ones. For example, any "primitive" located in the left-most third of the image and having roughly vertical orientation will be classified as a "left" one etc. Third step joins same-category "primitives", making them candidates for left/top/right and bottom sides of the document snippet. For example, collinear left "primitives" are merged into a candidate for the left document snippet side etc. Fourth and last step combines the "candidates" into a complete snippet candidate, assigning each complete candidate a confidence which reflects how well the candidate meets the document-specific assumptions about proportions, orientation, level of geometrical distortions, color contrast etc. Then the highest-confidence candidate is chosen to represent the document snippet's border and its confidence may be used as an indication of how reliably the snippet was found.

[0139] In another embodiment, edge detection helps to determine the focus quality of the captured image. If an edge is blurry or fuzzy, the remainder of the image, including the actual content on the remittance coupon or document, is also likely to be blurry and unreadable. A blurry or fuzzy image will produce high "out-of-focus scores," as described in U.S. Pat. No. 8,000,514, the contents of which are incorporated herein by reference in their entirety. Blurry or fuzzy images will also produce low confidence scores for the borders of the document during the border detection embodiment described immediately above. The drop in confidence occurs due to worsening of contrasts along one or more of the snippet sides. Therefore, the high out-of-focus score and low confidence scores will result in the application determining that the image quality is blurry or out-of-focus, and request that the user capture another image.

[0140] In another embodiment, edge detection may be used to determine an orientation angle of the mobile device with respect to the document and allow the user to correct perspective distortion of the captured image. Ideally, if the camera angle was exactly perpendicular to the document (and the camera did not have any optical distortions), the document snippet would be rectangular. That "ideal" rectangle gets distorted into a quadrilateral, often a trapezoid, when the camera angle deviates from perpendicular. Assuming the document corners have been detected by the document snippet's border detection algorithm (see above), the distortion could be measured using a deviation of the quadrilateral's angles from 90 degrees and/or the size difference between opposite sides of the quadrilateral. The orientation angle of the camera closely correlates with a View Angle Image Quality Assessment (IQA) score, which explains how the latter is computed based on these distortion characteristics. Further descriptions are available in U.S. Pat. No. 8,000,514, the contents of which are incorporated herein by reference in their entirety. Depending on the document type, the minimum value of a View Angle IQA (an IQA threshold) could be

chosen between **900** (camera view close to perpendicular, small distortion of document) and **700** (camera view deviates from perpendicular by about 15% causing more pronounced distortions of the document).

[0141] In a further embodiment, edge detection may also be used to determine whether the remittance coupon within the image is too small, based on the amount of space within the photograph outside of the four detected sides.

[0142] Edge detection may also be able to determine whether the background is busy, based on detection of edges that are either outside or orthogonal to those detected on the images. A busy background is one that interferes with the detection of edges of the desired document. For instance, a plain bill having black text on white paper, when placed on a larger white paper, would interfere with detecting the edge of the document. Another busy background is when a user holds a document in front of themselves, facing a picturesque scene of ocean and sky, and the camera sees the document in front of a dark and light background simultaneously. Other busy backgrounds may be plaid cloth, window shades, and other non-uniform colored or shaded background.

[0143] The result of applying an edge detector to an image may lead to a set of corners and document edges, both bounding the document being sought within the image, as well as any other objects outside it, or within it. This typically indicates the boundaries of objects, the boundaries of surface markings as well as curves that may correspond to discontinuities in surface orientation. By applying an edge detection algorithm to an image, the amount of data to be processed may be significantly reduced. The application may therefore filter out information such as detection of an out-of-focus image or an image that doesn't contain the entire document being captured.

[0144] Edges extracted from non-trivial images are often hampered by fragmentation, meaning that the edges are not connected. Certain issues such as missing edge segments and/or false edges not corresponding to the rectangular document being searched for in the document can complicate the subsequent task of determining the document type through classification, as well as hampering the ability to apply knowledge about the structure layout and context of the document. One example is a user holding a document with a finger or thumb covering a portion of one edge, where the information fields are not covered and otherwise the captured image would be processed accurately if the edge detection were acceptable.

[0145] Edge detection on the mobile device is carried out using the graphical and processing units of the mobile device. The edge detection capability allows the detection of, and rejection of images with one or more of the above list of issues, based on if, and where, the edges are found, their position, and their relationships within the image.

Edge Detection at the Mobile Device

[0146] There are many ways to perform edge detection of both documents, as well as other objects within an image, from a smartphone or mobile device. In the ideal case, the result of applying an edge detector to an image may lead to a set of corners and document edges, both bounding the document being sought within the image, as well as any other objects outside it, or within it. This typically indicates the boundaries of objects, the boundaries of surface markings as well as curves that may correspond to discontinuities in surface orientation. By applying an edge detection algorithm to

an image of a document, we can significantly reduce the amount of data to be processed and may therefore filter out information such as detection of an out-of-focus image or an image that doesn't contain the entire document being captured. Edges extracted from non-trivial images are often hampered by fragmentation, meaning that the edges are not connected. Certain issues such as missing edge segments and/or false edges not corresponding to the rectangular document being searched for in the document can complicate the subsequent task of determining the document type through classification, as well as hamper the ability to apply knowledge about the structure layout and context of the document.

[0147] Therefore, edge detection at the mobile device allows for initial filtering of images that have a high likelihood of being sub-quality, and allows, in real-time, the ability to indicate to the user various reasons for taking the picture again.

[0148] Edge detection helps directly and indirectly determine the following: the focus quality of the image; whether all four corners and four sides of the document are within the photographic image; what the camera angle is with respect to the document (based on the perspective distortion of the quadrilateral within the image compared to the expected rectangle's dimensions); whether the document within the image is too far away (based on the amount of space within the photograph outside of the four sides detected); and whether the background is busy, based on detection of edges that are either outside or orthogonal to those detected on the image.

[0149] The edge detection capabilities run on the mobile device, using their graphical and processing CPU's. The capabilities allow the detection—and rejection—of images with one or more of the above issues relating to angle, distance, business, etc., based on if and where the edges are found, as well as their position and relationship within the image.

[0150] The first stage of edge detection is image cropping. The quality of edge detection depends on how existing information about a document is used. In most cases, a document has a rectangular shape and is placed on some distinct background, as show in FIG. 53A. Cropping based on edge detection is the proper method for these types of images with clearly-defined documents and contrasting backgrounds. Cropping based on edge detection may include several steps: 1) detecting the edges of the document (as shown in FIG. 53B); 2) extracting the lines (which can be done by edge tracing or using Hough transformation) and 3) detecting four corners of the document through refinement and then cropping the image using the found corners.

[0151] But in some cases, for example if the document is lying on a non-distinct background **5312** (as illustrated in FIG. 53C) or on a cluttered background **5314** (as illustrated in FIG. 53D), the edge detection steps described above fail. In these cases, prior knowledge of a document logo, text or picture can be used for cropping enhancement. For example, for a driver's license, prior knowledge of key words and the word locations can be used as a template image **5400**, as illustrated in FIG. 54A. The template image **5400** may be matched with the document **5402** in the captured image in FIG. 54B based on these key words. However, direct template matching may require significant computational resources; and therefore feature matching is more preferable.

[0152] In one embodiment, the first stage of proposed cropping is feature detection, which can be implemented using a multi-scale Hessian operator. Next, feature points are

detected in the local maxima of the Hessian operator output. In the next stage, the description of feature points is built. The distributions of local gradients are calculated in the area of feature points. Because the document can have some distinct colors (as shown by the red 5308 and blue 5310 keywords in FIG. 53C), color distribution may be added to feature point descriptors for enhancing feature matching performance, which is the next step of the algorithm. A matching example is shown in FIG. 54B). The last stage is cropping using a transformation matrix, which is calculated based on matched feature points.

[0153] Sometimes, a document can have several different templates, for example the template image 5400 in FIG. 54A and a second template image 5406 as illustrated in FIG. 54C. With two or more templates, the feature matching steps should be completed for both templates, and the best cropping parameters should be selected based on an analysis of output transformation matrixes and prior information about the document (such as size and aspect ratio). In one embodiment illustrated in FIG. 52, two channels of image cropping are proposed, where an input image 5202 is cropped in a first channel using edge-based cropping 5204, and also in a second channel using feature-based cropping 5206. The merging block 5208 combines results of the two cropping methodologies based on prior information about document (size, aspect ratio) and results of recognition (if OCR was applied on the output of each cropper).

Real-time Feedback

[0154] In one embodiment, one or more image quality assessment (IQA) tests are performed on the captured image to ensure that the image is of sufficient quality for further processing. If the image does not pass one or more of these IQAs, the user may be provided with a feedback message. Feedback messages from the system to the user help the user understand and eliminate obstacles to successful processing. These messages can originate from the mobile device, from the remote server, or from the financial institution's or the billers' own system. For alerts to be useful, they should be specific, which is often difficult. Feedback alerts typically fall into the following three categories:

- [0155] a. Image quality issues, which the user can correct by moving the camera appropriately or re-capturing the image.
 - [0156] b. Technical issues which prevent the user from completing an action related to the captured document, such as a wrong account password or failed network connection.
 - [0157] c. Business issues, such as unknown account number or unknown document type, that prevents the user from completing the action related to the captured document.
- [0158] The user may be able to correct the image quality issues, while the other two issues prevent the user from completing the desired action related to the document, such as paying a bill or depositing a check. However, it is still important to let the user know exactly why the transaction could not be completed and what they can do about it (e.g. capture the image again, wait for an Internet connection, or contact a customer support service).
- [0159] Of the three types of alerts, the first—image quality issues—is the most difficult to offer actionable feedback. The system must offer as much specific assistance as possible to allow the user to take better pictures, especially considering

the large variety of potential image issues (insufficient lighting, cut-off corners, blurry image, etc.). The mobile image processing steps described above, such as orientation angle, amount of motion, edge detection, de-warping and shadow detection, provide the specific image analyses needed to generate effective feedback to the user in order to correct image defects. In accordance with the embodiments above, the user can be provided with feedback to adjust the angle of the mobile device, hold it steady to prevent motion, line-up the quadrilateral outline with the edges of the document or eliminate a shadow on the document. Numerous additional feedback messages may be generated based on additional processing steps which take place on the mobile device or even at the remote server.

[0160] In one embodiment, the feedback may be displayed on the display of the mobile device, such as a text displayed to the user. The feedback may also be non-text visual feedback, such as a check-symbol, an “x” symbol or a color-coded status bar indicating the quality of the image (green for high quality, yellow for medium quality, red for poor quality). The feedback may also be audio—either a spoken voice telling the user what to correct, or a non-vocal sound (ring, beep chime, etc) that indicates if the image capture is acceptable or needs fixing. The feedback may also be tactile as well, with one or more vibrations produced by the mobile device to indicate whether an image capture quality is acceptable or not. Other types of feedback are possible as well, and the aforementioned list should be considered non-limiting.

Document Identification

[0161] There are various technologies that can be used to identify the document in the captured image and identify the document type on the mobile device. The benefits of document identification at the mobile device include the ability to detect the document and the document type in real-time without the user needing to manually select it or determine it during server-side processing. Furthermore, the document can then be reconstituted in its proper dimensions and cropped on the mobile device, so that a smaller image can be sent to the server instead of the considerably larger entire image. The document type can also be provided to the server to avoid the need for significant document type identification processing on the server-side. Various features described above may be used for document identification, including edge detection and pre-cropping. The dimensions of the cropped image can then be utilized as one of several clues as to the document type. In addition, detection of the presence of photos, icons, logos, colors and color locations and reflectivity may also be used to determine the document type.

[0162] Specific examples of how the detection of photos, icons, logos, colors and reflectivity lead to document identification include:

- [0163] 1. Knowledge of the various document types can be hosted in a biller database and utilized by the mobile device-side technology via phone applications, and updated dynamically with meta-data sent from the server when the mobile application initially connects.
- [0164] 2. Presence of photos, including the photo positions, significantly narrows down the choice of possible document types (e.g. 1-2 photos are typical on Driver's Licenses but not on remittances).
- [0165] 3. Presence of rounded corners significantly narrows down the choice of possible document types (e.g.

rounded corners are typical on Driver's Licenses and Credit Cards but not on remittances).

- [0166] 4. Detection of characteristic "key points" using a scale-invariant and rotation-invariant feature transform algorithm can identify type and position of the document within the captured image.
- [0167] 5. Detection of certain image elements, including geometrical lines, boxes and text blocks (normalized to achieve scale-invariance and rotation-invariance) can uniquely identify some known templates which are rich in such image elements.
- [0168] 6. Detection of a readable MICR line can uniquely identify that the document is a check.
- [0169] 7. Detection of barcodes, PDF417 codes, QR codes or other verifiable codes significantly narrows down the selection of possible document types.
- [0170] 8. Color-map description (normalized to achieve scale-invariance and rotation-invariance) can identify known templates with unique color distribution.
- [0171] 9. Detection of reflectivity, including that of holographic elements, significantly narrows down the choice of possible document types (e.g. reflections/glare are typical on plastic documents such as Driver's Licenses and Credit Cards but not on paper-based documents such as remittances).
- [0172] Other methods of database-assisted and dynamic data capture-based form identification are described herein in the sections entitled "Form Identification" and "Dynamic Data Capture," the methods and features of which may be implemented on the mobile device or the server.

Compression and Mobile Cropping

[0173] In one embodiment, the capture process may include a final "packaging" step, during which the image is compressed and optimized for transmission. This results in a smaller document size, and faster transmission speed. In a further embodiment, the process may include a mobile cropping algorithm, cropping the image before compression, eliminating background presence on the image, which is ultimately of no value to the intended transaction. This further reduces the size of the image before transmission.

II. Post-Capture Image Processing

[0174] The process that occurs once the image or images have been captured includes one or more additional image processing and content extracting steps to capture the content of the document. In one embodiment, processing at the server also includes use of at least one database to compare known information about a biller with the extracted content and confirm the accuracy of the extracted content. An overview of one embodiment of the workflow of the processing steps which occur at the mobile device or remote server is provided in FIG. 2.

Image Correction

[0175] In one embodiment, the captured image undergoes one or more image processing steps to further correct various aspects of the image, improving the overall quality and readability of the remittance coupon before the content is extracted. The captured image may first undergo conversion from three-dimensions (3D) to two-dimensions (2D) to correct perspective distortion, and may then be cropped and reconstituted into a rectangular shape that resembles the

dimensions of the original document. Rotation and skew correction may also be completed (described in further detail below). In one embodiment, a pixel-level update is performed (not shown) to ensure that the characters, fields, logos and other data found on the image are converted back to the 2D version of the document. Numerous additional image correction steps may be executed on the server, as will be described in detail in Section III, below.

Codeline Read

[0176] In one embodiment, a next step is to read a code line on the remittance coupon at a code-reading unit which contains important information about the biller and the bill. FIG. 9 illustrates one embodiment of a codeline 905. The code-reading unit may also be configured to read a barcode on the remittance coupon, which will be described further below. FIG. 9 illustrates one embodiment of a barcode 925. The type of document in the captured image may then be identified by a document classifier unit, if possible, through a variety of methods which may compare the remittance coupon with a database of documents, look to specific codes, fields or content on the remittance coupon.

First Content Recognition Process (First Pass)

[0177] A first content recognition pass of the image may be made using optical character recognition (OCR) or Intelligent Character Recognition (ICR) to capture all of the data and fields on the remittance coupon. A processing engine unit may be provided to coordinate the OCR/ICR of the captured image with an OCR engine. An output of the OCR/ICR may include both character-level and field-level strings, the coordinates where the content is found, the confidence level of the recognition of the content, as well as the cleaned up and cropped images.

[0178] In one embodiment, dynamic field extraction, described further herein, may be performed to find fields on an unstructured document where there are no standards with regard to the location or context of the information and fields. The methods of identifying the type of document are described further herein.

[0179] Detection of field coordinates and confidences is part of the dynamic data capture process described separately herein. The process starts with accepting an image (which maybe bitonal, grey-scale or full color) and rules for capturing fields of interest. Since bills are usually printed in black-and-white, it's sufficient for data capture to use bitonal (1 bit/pixel) images. The fields of interest may include Account Number, Amount Due, Payee Address and Payee Name, Amount and Date Due etc. Each such field is defined by a set of rules which help the data capture process to distinguish this field from others. The rules usually (but not always) include restrictions on field location (e.g. in the left-top quadrant of the document), format (e.g. contains from 3 to 10 digits and up to 3 alphas), textual clues/keywords (e.g. adjacent to "Account No"), relation to the keywords and/or other fields (e.g. located to the right of Amount Due, which by itself is a field of interest) etc.

[0180] Whatever the color depth of the first pass' input image is, it gets always detected, cropped and geometrically corrected by the snippet's border detection algorithm described above.

[0181] The dynamic data capture system usually starts with full-page OCR of the image (to speed it up, only part of the

image defined by the rules may be used). OCR results, in addition to ASCII code, contain location, confidence and some other information for each character. Then, depending on the rules, the data capture system applies various techniques to locate each field within the OCR result. For example, if a field is defined by its format, a fuzzy-search method is used to find a subset of OCR-result which meets the format. If the field is defined by its limited search area, only part of the OCR result will be used. If the field is defined via its relation to certain keywords and/or other fields, the latter are found prior to finding the field etc.

[0182] Whatever the rule is, it always produces a confidence value—a numeric measure of how consistent the rule and found field location (also called field alternative) are. For example, the “Located in left-top quadrant” rule will produce confidence of 1000 (maximum) if given field alternative is located entirely in the quadrant and only 500 if one half of the alternative is outside of the quadrant. The “Field is entirely numeric” format rule will produce confidence of 1000 if all characters in the alternative are numeric and reduce the score for each alpha character (the penalty may vary). Furthermore, the rules may produce character-level confidences: e.g. the alpha character in the previous example will have a format confidence of “0,” whereas other numeric characters will have the format confidence of 1000.

[0183] Once all the rules are executed and the field is found, its overall (field-level) confidence is computed as a function of individual rules’ confidences. Individual character confidences are computed using their OCR-confidence and character-level rule confidences.

[0184] In one embodiment, post-process rejections may be made, where an image is rejected even after successfully cropping and reading the document based on a combination of low scores across multiple fields. This is typically when there is either a bad portion of the image, or the image looked good in the first pass, but showed low confidences when key field level values were analyzed. Therefore, in the aggregate, the image is rejected based on the confidence levels of the extracted content. If the image is rejected at this stage, a message may be provided to the user at the mobile device, indicating that another image must be taken and possibly providing specific advice on how to improve the image capture.

[0185] According to an embodiment, the remote server can be configured to report the results of the image quality assurance testing to the mobile device. This can be useful for informing a user of the mobile device that an image that the user captured of a remittance coupon passed quality assurance testing, and thus, should be of sufficient quality that the mobile image can be processed by the remote server. According to an embodiment, the remote server can be configured to provide detailed feedback messages to the mobile device **102** if a mobile image fails quality assurance testing. Mobile device **102** can be configured to display this feedback information to a user of the device to inform the user what problems were found with the mobile image of the remittance coupon and to provide the user with the opportunity to retake the image in an attempt to correct the problems identified.

[0186] If the mobile image passes the image quality assurance testing, the remote server can submit the mobile image plus any processing parameters received from the mobile device to the remote server for processing.

Barcode Detection

[0187] In one embodiment, a pre-processing step may include barcode detection and recognition. If a barcode is detected on the document, the barcode is read and saved alongside the coordinates. The barcodes on bills may include address information of the biller in the form of the zip code plus four digit identifier (“zip+4”) value, positioned right below an address block on the remittance coupon. A comparison can then be made with the optically-read zip+4 value and the barcode-provided value, and a vote is taken on the two for the best guessed value. The barcodes are typically address-type information, such as the zip code plus four digit identifier (i.e. 92101-6789), but may correspond to a payor and payee of a bill. The location of the barcodes on the document is useful in determining the type of address (payor or payee) which the barcode contains.

Address Search

[0188] Use of dictionaries both at the language level and keyword level, as well as vector location information around particular fields types helps find fields within a larger semantic or syntactic meaning. In one embodiment, various dictionaries, or biller databases, may be used to find the biller based on information captured from the document. For example, in step, a fuzzy search of address database will allow for further qualification and normalization of the address information obtained from the first pass in, which improves the overall accuracy of the system. The address search may be carried out by an address search unit. In one embodiment, this fuzzy search includes search of a database of nationwide biller information that contains the biller name, full address, zip+4 and various aliases. The fuzzy search means that an exact match is not necessary in the event that the OCR/ICR of the image was not exact and certain address or biller name fields are not perfectly accurate. The fuzzy search looks for a best match based on standard algorithms around string comparisons and scoring, and provides a list of billers where the spelling is close. The address database can be searched for address information, such as the zip+4, which corresponds to the zip+4 found during the first pass or through codeline or barcode detection. The address database may be a United States Postal Service (USPS) database of valid addresses that can be used to validate the information read off of a bill with regard to the Payor and Payee.

Biller Lookup

[0189] Once the payee address is known with a certain degree of confidence, a biller lookup process may be initiated by a biller lookup unit to identify the biller (payee) on the remittance coupon. The biller lookup process attempts to identify the entity responsible for creating a bill so that a payment made by a user will be transferred to the correct entity. In one embodiment, the biller lookup process may perform a “fuzzy” search against the customized biller database with the fields identified during the first pass used as input for the search. The biller database may contain biller profile information on numerous billers (payees). The biller profile information may include their addresses, various aliases they might be known as, remittance coupon formats, fields used, as well as any account number formats, address formats, codeline formats and other biller-specific fields (determined with masks/regex). For example, a particular zip+4 zip code, “92101-1234,” may be found on the remittance

coupon by the OCR content capture process. The server-side application may then look up that zip+4 and determine that both billers “City G&E” and “Municipal Water District” process bills at this payee address. In order to determine which biller the remittance coupon is from, the remittance coupon may then be re-read a second time on a second data recognition pass, armed with the two possible biller names, such that in the second read, the application looks for either “City G&E” text or “Municipal Water District” text. This second pass puts the Biller Name in greater context and provides further verification of the biller. Overall, the data in the biller database allows for the system to “read” the remittance coupon multiple times, if needed, with increasing levels of knowledge about the classification of the bill, the biller keywords expected, and account masks (via use of RegEx).

[0190] The biller lookup process may be broken up into five different phases, as illustrated in FIG. 5. Data obtained from the first pass of the data extraction processing engines (sometimes called “MIP Science Engines”) is shown in the boxes on the left column under “MIP Science,” while data stored in the biller database is shown in the boxes on the right side under “Biller DB.” In phase one, a query of the billing database of billers for a Zip Code and/or Zip Code plus four which matches the Zip Code found during the first pass. The output of the phase 1 search is either a direct hit on a biller or a list of billers and their potential aliases. If no matches are found, then the search is retried using the Payor Zip Code and/or Zip Code plus four. If still no match is found, the system exits the biller lookup process.

[0191] If a biller is found, then phase two proceeds, where an “exact match” comparison is done for the subset of billers identified during phase one. The exact match comparison compares each biller’s Coupon-Biller-Name from the biller database with the Payee Recipient (payee) found during the first pass. If a single match is found, then the Coupon Name, Payee Recipient Name, and Account Number Format found during the first pass are replaced with the biller profile information for that biller name in the biller database, and the biller lookup process is terminated. If no matches are found, then the process jumps to phase four. If the process of phase two results in more than one biller, then the lookup process proceeds to phase three.

[0192] In phase three, an “exact match” comparison using the subset of billers found in phase two is performed to compare the biller’s Address Line 1 and/or Address Line 2 from the biller database with the PO Box and/or reconstructed Address Line 1 found during the first pass. If a single match is found, then the Coupon Name, Payee Recipient Name, and Account Number RegEx Format are replaced and the biller lookup process is terminated. If no matches are found or more than one match is still found, then the biller lookup process proceeds to phase four.

[0193] In phase four, the application will build a list of billers from the biller database to compare against certain fields from the raw OCR data returned from the OCR engine obtained during the first pass. If a Coupon Name exists, then the application first searches the Biller DB for any matching Biller-Coupon-Name. Any matches are added to a sub-list of Billers. If the MIP Payee Recipient exists, then the system searches the Biller DB for any matching Payee Recipient. Any matches are added to a sub-list of Billers. If a single match is found, then the Coupon Name, Payee Recipient Name, and Account Number Format are replaced and the

biller lookup process is terminated. If no matches are found, then the biller lookup process is terminated.

[0194] If more than one biller is identified during phase four, in phase five, each biller found in phase four is scored by doing a fuzzy comparison of the biller with the raw OCR data. The highest ranked biller is then obtained, and if the score is above a certain threshold level, such as 70%, the Payee Recipient Name, the Coupon Name and Account Number Format are replaced (**S526**) and the biller lookup process is terminated.

[0195] The biller lookup process is configured to identify the billing entity with great confidence. Once the biller is identified, additional biller profile information can be obtained from the biller database. The biller database contains both nationwide biller address information, as well as specific formats and masks for various fields found on bills. The “mask” may be a format or regex (regular expression) that provides details on the format, layout and characters and potential checksums used for formatting things like account numbers, and is sometimes very specific to a particular bill format. This includes account number formats, address formatting, code line formatting and other biller-specific fields found.

[0196] In one embodiment, mask information may provide basic template information on the remittance coupon for that biller (for instance, an account mask may indicate that a particular credit card issuer always has account numbers which start with the number “3” and have 15 digits). With this account mask information, the account number field identified during the first pass may be re-read during a second pass to obtain the account number off the remittance coupon, this time with greater accuracy.

[0197] Additional “dictionaries” may be provided which are specific to payment of bills and focus on phrases that are common, for instance “please pay this amount”, “the amount due is,” and “the check should be made out to,” etc., which may be used to identify the amount due.

Second Content Recognition Process (Second Pass)

[0198] In one embodiment, a second content recognition process (second pass) of the captured image may be performed by an engine processing unit with further hints to the OCR engine based on the information obtained from the address search and biller lookup processes. For example, the hints may include more information on formats masks via regex expressions, as well as information on the biller, document format, location information and so forth.

[0199] In one embodiment, the second pass is used to re-recognize an account number by using a narrowed RegEx (regular expression) provided by the biller lookup process. One embodiment of the second pass process is illustrated in FIG. 6, and includes the following steps, which may vary. In a first step, the system will first verify that a second pass of the document is needed (**S602**), such as when the biller lookup process has located a biller and obtained a new account number format from the biller database. The biller lookup process can, and will, prevent the second pass process from running if the account number already matches the biller’s regular expression for known account number format(s). In a second step **S604**, a plurality of runtime configurable settings are loaded based on the biller profile information obtained from the biller database **108**. The second pass has various logic settings configured around how the field level data and outcomes are examined. This logic is different for different types

of documents, such as a driver's license or a bill. So the logic of this second pass dynamically depends on what type of document is being captured. In a third step (S606), the fields which are to be processed again in the second pass are updated with new RegEx data obtained from the biller lookup process. For example, an account number RegEx may be updated with a new RegEx provided by the biller lookup process.

[0200] With the updated settings loaded, the second pass is now performed (S608) with the updated second pass runtime configuration by executing one or more OCR/ICR engines or other low level processing engines. The engines are this time provided with 'hints' via masks which indicate probable locations of fields and the format of certain fields. For instance, given an account mask where a biller is known to have account numbers which are 15 digits and always start with a "3," the account number field is re-read in the second pass. Further details of the account mask may also be known, such as the use of a space between digits 7 and 8.

[0201] FIGS. 7A and 7B illustrate the effect of the second pass process on a remittance coupon 700 with an account number field 702. The first pass could return results illustrated in FIG. 7A, which shows the identification of incomplete portions 704 of the account number. However, with the execution of the biller lookup process and the second pass process, the complete account number 706 is correctly identified based on the appropriate account number field 702, as illustrated in FIG. 7B. The incomplete portion 704 of the account number has been ignored since it did not fit the known account number mask.

[0202] Once the second pass is complete, the address, account number and other extracted data may be parsed and cleaned up. The newly-extracted data is evaluated (S610) to provide new confidence levels reflective of the additional biller profile information. If the extracted data meets required confidence thresholds, it will be deemed the final value and stored in the content database 110 for output to the user and the appropriate financial institutions for processing the bill payment, as described below.

Billing Information Output

[0203] Once final values are obtained for the biller, payor and other content of the document, these final values are stored in the content database 110 along with other information, including the original JPG image from the mobile device, a cropped grayscale image and one or more bitonal images. The grayscale image and bitonal images may have been created at the mobile device or the remote server for the image correction steps, as described above. More details regarding the use of grayscale and bitonal images is provided below. The extracted data from the remittance coupon may output from the recognition engines in an XML file and stored in the content database 110. The content database 110 will also store all data, locations and confidence values around character field and document characteristics. In one embodiment, datatime, geo-locations, and user session information will also be stored, and may be used for user verification and other security information. Finally, the version of the system in place at the time, on both phone and server, may be stored as well.

[0204] In one embodiment, a final output 418 which may include the final values and images may be presented to the user on a graphical user interface (GUI) on a display of the mobile device so that the user can verify the accuracy of the

extracted data and then approve the payment of the bill. The final values will then be submitted to the banking server 112 which will handle the actual processing of the payment from a bank account of the user to the payee. In another embodiment, the final values may be submitted directly to the banking sever 112 for processing of the payment.

III. Image Processing of Mobile-Captured Images

[0205] The systems and methods provided herein advantageously allow a user to capture an image of a remittance coupon, and in some embodiments, a form of payment, such as a check, for automated processing. Typically, a remittance processing service will scan remittance coupons and checks using standard scanners that provide a clear image of the remittance coupon and accompanying check. Often these scanners produce either gray-scale and bi-tonal images that are then used to electronically process the payment. The systems and methods disclosed herein allow an image of remittance coupons, and in some embodiments, checks to be captured using a camera or other imaging device included in or coupled to a mobile device, such as a mobile phone. The systems and methods disclosed herein can test the quality of a mobile image of a document captured using a mobile device, correct some defects in the image, and convert the image to a format that can be processed by remittance processing service.

[0206] The term "standard scanners" as used herein, but is not limited to, transport scanners, flat-bed scanners, and specialized check-scanners. Some manufacturers of transport scanners include UNISYS®, BancTec®, IBM®, and Canon®. With respect to specialized check-scanners, some models include the TellerScan® TS200 and the Panini® My Vision X. Generally, standard scanners have the ability to scan and produce high quality images, support resolutions from 200 dots per inch to 300 dots per inch (DPI), produce gray-scale and bi-tonal images, and crop an image of a check from a larger full-page size image. Standard scanners for other types of documents may have similar capabilities with even higher resolutions and higher color-depth.

[0207] The term "color images" as used herein, pertains to, but is not limited to, images having a color depth of 24 bits per a pixel (24 bit/pixel), thereby providing each pixel with one of 16 million possible colors. Each color image is represented by pixels and the dimensions W (width in pixels) and H (height in pixels). An intensity function I maps each pixel in the [W×H] area to its RGB-value. The RGB-value is a triple (R,G,B) that determines the color the pixel represents. Within the triple, each of the R(Red), G(Green) and B(Blue) values are integers between 0 and 255 that determine each respective color's intensity for the pixel.

[0208] The term "gray-scale images" as used herein may be considered, but is not limited to, images having a color depth of 8 bits per a pixel (8 bit/pixel), thereby providing each pixel with one of 256 shades of gray. As a person of ordinary skill in the art would appreciate, gray-scale images also include images with color depths of other various bit levels (e.g. 4 bit/pixel or 2 bit/pixel). Each gray-scale image is represented by pixels and the dimensions W (width in pixels) and H (height in pixels). An intensity function I maps each pixel in the [W×H] area onto a range of gray shades. More specifically, each pixel has a value between 0 and 255 which determines that pixel's shade of gray.

[0209] Bi-tonal images are similar to gray-scale images in that they are represented by pixels and the dimensions W

(width in pixels) and H (height in pixels). However, each pixel within a bi-tonal image has one of two colors: black or white. Accordingly, a bi-tonal image has a color depth of 1 bit per a pixel (1 bit/pixel). The similarity transformation, as utilized by some embodiments of the invention, is based off the assumption that there are two images of $[W \times H]$ and $[W' \times H']$ dimensions, respectively, and that the dimensions are proportional (i.e. $W/W' = H/H'$). The term “similarity transformation” may refer to a transformation ST from $[W \times H]$ area onto $[W' \times H']$ area such that ST maps pixel $p = p(x, y)$ on pixel $p' = p'(x', y')$ with $x' = x * W'/W$ and $y' = y * H'/H$.

[0210] FIG. 8 is an image illustrating an example remittance coupon **800** that can be imaged with the systems and methods described herein. The mobile image capture and processing systems and methods described herein can be used with a variety of documents, including documents such as personal checks, business checks, cashier’s checks, certified checks, and warrants. By using an image of the remittance coupon **800**, the remittance process can be automated and performed more efficiently. As would be appreciated by those of skill in the art, remittance coupons are not the only types of documents that might be processed using the system and methods described herein. For example, in some embodiments, a user can capture an image of a remittance coupon and an image of a check associated with a checking account from which the remittance payment will be funded.

[0211] FIG. 9 is a geometrically corrected image **900** created using image processing techniques disclosed herein and using the mobile image of the remittance coupon **800** illustrated in FIG. 8. A remittance coupon may include various fields, and some fields in the documents might be considered “primary” fields. For example, some remittance coupons also include computer-readable bar codes or code lines **905** that include text or other computer-readable symbols that can be used to encode account-related information. The account-related information can be used to reconcile a payment received with the account for which the payment is being made. Code line **905** can be detected and decoded by a computer system to extract the information encoded therein. The remittance coupon can also include an account number field **910** and an amount due field **915**. Remittance coupons can also include other fields, such as the billing company name and address **920**, a total outstanding balance, a minimum payment amount, a billing date, and payment due date. The examples are merely illustrative of the types of information that may be included on a remittance coupon and it will be understood that other types of information can be included on other types of remittance coupons.

[0212] Once the image is captured and corrected, and the data is extracted and adjusted, then the image, data, and any required credential information, such as username, password, and phone or device identifier, can be transmitted to the remote server for further processing. This further processing is described in detail with respect to the remaining figures in the description below.

Image Processing

[0213] Mobile device and remote server can be configured to perform various processing on a mobile image to correct various defects in the image quality that could prevent the remote server or the banking server from being able to process the remittance due to poor image quality.

[0214] For example, an out of focus image of a remittance coupon or check, in embodiments where the mobile device

can also be used to capture check images for payment processing, can be impossible to read and process electronically. For example, optical character recognition of the contents of the imaged document based on a blurry mobile image could result in incorrect payment information being extracted from the document. As a result, the wrong account could be credited for the payment or an incorrect payment amount could be credited. This may be especially true if a check and a payment coupon are both difficult to read or the scan quality is poor.

[0215] Many different factors may affect the quality of an image and the ability of a mobile device based image capture and processing system. Optical defects, such as out-of-focus images (as discussed above), unequal contrast or brightness, or other optical defects, can make it difficult to process an image of a document, e.g., a check, payment coupon, deposit slip, etc. The quality of an image can also be affected by the document position on a surface when photographed or the angle at which the document was photographed. This affects the image quality by causing the document to appear, for example, right side up, upside down, skewed, etc. Further, if a document is imaged while upside-down it might be impossible or nearly impossible to for the system to determine the information contained on the document.

[0216] In some cases, the type of surface might affect the final image. For example, if a document is sitting on a rough surface when an image is taken, that rough surface might show through. In some cases the surface of the document might be rough because of the surface below it. Additionally, the rough surface may cause shadows or other problems that might be picked up by the camera. These problems might make it difficult or impossible to read the information contained on the document.

[0217] Lighting may also affect the quality of an image, for example, the location of a light source and light source distortions. Using a light source above a document can light the document in a way that improves the image quality, while a light source to the side of the document might produce an image that is more difficult to process. Lighting from the side can, for example, cause shadows or other lighting distortions. The type of light might also be a factor, for example, sun, electric bulb, florescent lighting, etc. If the lighting is too bright, the document can be washed out in the image. On the other hand, if the lighting is too dark, it might be difficult to read the image.

[0218] The quality of the image can also be affected by document features, such as, the type of document, the fonts used, the colors selected, etc. For example, an image of a white document with black lettering may be easier to process than a dark colored document with black letters. Image quality may also be affected by the mobile device used. Some mobile camera phones, for example, might have cameras that save an image using a greater number of mega pixels. Other mobile cameras phones might have an auto-focus feature, automatic flash, etc. Generally, these features may improve an image when compared to mobile devices that do not include such features.

[0219] A document image taken using a mobile device might have one or more of the defects discussed above. These defects or others may cause low accuracy when processing the image, for example, when processing one or more of the fields on a document. Accordingly, in some embodiments, systems and methods using a mobile device to create images of documents can include the ability to identify poor quality

images. If the quality of an image is determined to be poor, a user may be prompted to take another image.

Detecting an Out of Focus Image

[0220] Mobile device and remote server can be configured to detect an out of focus image. A variety of metrics might be used to detect an out-of-focus image. For example, a focus measure can be employed. The focus measure can be the ratio of the maximum video gradient between adjacent pixels measured over the entire image and normalized with respect to an image's gray level dynamic range and "pixel pitch". The pixel pitch may be the distance between dots on the image. In some embodiments a focus score might be used to determine if an image is adequately focused. If an image is not adequately focused, a user might be prompted to take another image.

[0221] According to an embodiment, the mobile device can be configured to detect whether an image is out of focus using the techniques disclosed herein. In an embodiment, the remote server can be configured to detect out of focus images. In some embodiments, the remote server can be configured to detect out of focus images and reject these images before performing mobile image quality assurance testing on the image. In other embodiments, detecting and out of focus image can be part of the mobile image quality assurance testing.

[0222] According to an embodiment, an image focus score can be calculated as a function of maximum video gradient, gray level dynamic range and pixel pitch. For example, in one embodiment:

$$\text{Image Focus Score} = (\text{Maximum Video Gradient}) * (\text{Gray Level Dynamic Range}) * (\text{Pixel Pitch}) \quad (\text{eq. 1})$$

[0223] The video gradient may be the absolute value of the gray level for a first pixel "i" minus the gray level for a second pixel "i+1". For example:

$$\text{Video Gradient} = \text{ABS} [(\text{Grey level for pixel "i"}) - (\text{Gray level for pixel "i+1"})] \quad (\text{eq. 2})$$

[0224] The gray level dynamic range may be the average of the "n" lightest pixels minus the average of the "n" darkest pixels. For example:

$$\text{Gray Level Dynamic Range} = [\text{AVE}(\text{"N" lightest pixels}) - \text{AVE}(\text{"N" darkest pixels})] \quad (\text{eq. 3})$$

[0225] In equation 3 above, N can be defined as the number of pixels used to determine the average darkest and lightest pixel gray levels in the image. In some embodiments, N can be chosen to be 64. Accordingly, in some embodiments, the 64 darkest pixels are averaged together and the 64 lightest pixels are averaged together to compute the gray level dynamic range value.

[0226] The pixel pitch can be the reciprocal of the image resolution, for example, in dots per inch.

[0227] In other words, as defined above, the pixel pitch is the distance between dots on the image because the Image Resolution is the reciprocal of the distance between dots on an image.

$$\text{Pixel Pitch} = [1/\text{Image Resolution}] \quad (\text{eq. 4})$$

[0228] In other words, as defined above, the pixel pitch is the distance between dots on the image because the Image Resolution is the reciprocal of the distance between dots on an image.

Detecting and Correcting Perspective Distortion

[0229] FIG. 10 is a diagram illustrating an example of perspective distortion in an image of a rectangular shaped document. An image can contain perspective transformation distortions 2500 such that a rectangle can become a quadrangle ABCD 2502, as illustrated in the figure. The perspective distortion can occur because an image is taken using a camera that is placed at an angle to a document rather than directly above the document. When directly above a rectangular document it will generally appear to be rectangular. As the imaging device moves from directly above the surface, the document distorts until it can no longer be seen and only the edge of the page can be seen.

[0230] The dotted frame 2504 comprises the image frame obtained by the camera. The image frame is sized h x w, as illustrated in the figure. Generally, it can be preferable to contain an entire document within the h x w frame of a single image. It will be understood, however, that some documents are too large or include too many pages for this to be preferable or even feasible.

[0231] In some embodiments, an image can be processed, or preprocessed, to automatically find and "lift" the quadrangle 2502. In other words, the document that forms quadrangle 502 can be separated from the rest of the image so that the document alone can be processed. By separating quadrangle 2502 from any background in an image, it can then be further processed.

[0232] The quadrangle 2502 can be mapped onto a rectangular bitmap in order to remove or decrease the perspective distortion. Additionally, image sharpening can be used to improve the out-of-focus score of the image. The resolution of the image can then be increased and the image converted to a black-and-white image. In some cases, a black-and-white image can have a higher recognition rate when processed using an automated document processing system in accordance with the systems and methods described herein.

[0233] An image that is bi-tonal, e.g., black-and-white, can be used in some systems. Such systems can require an image that is at least 200 dots per inch resolution. Accordingly, a color image taken using a mobile device can need to be high enough quality so that the image can successfully be converted from, for example, a 24 bit per pixel (24 bit/pixel) RGB image to a bi-tonal image. The image can be sized as if the document, e.g., check, payment coupon, etc., was scanned at 200 dots per inch.

[0234] FIG. 11 is a diagram illustrating an example original image, focus rectangle and document quadrangle ABCD in accordance with the example of FIG. 10. In some embodiments it can be necessary to place a document for processing at or near the center of an input image close to the camera. All points A, B, C and D are located in the image, and the focus rectangle 2602 is located inside quadrangle ABCD 2502. The document can also have a low out-of-focus score and the background surrounding the document can be selected to be darker than the document. In this way, the lighter document will stand out from the darker background.

Image Correction

[0235] FIG. 12 is a flow diagram illustrating a method for correcting defects to mobile image according to an embodiment. According to an embodiment, the method illustrated in FIG. 12 can be performed by the image correction unit 404 implemented on the remote server. The method illustrated in

FIG. 12 can be implemented as part of step S212 of the method illustrated in FIG. 2. The image correction unit can also receive a mobile image and processing parameters from the mobile device. According to some embodiments, some or all of the image correction functionality of the image correction unit can be implemented on the mobile device, and the mobile device can be configured to send a corrected mobile image to the remote server for further processing.

[0236] According to an embodiment, the image correction unit can also be configured to detect an out of focus image using the technique described above and to reject the mobile image if the image focus score for the image falls below a predetermined threshold without attempting to perform other image correction techniques on the image. According to an embodiment, the image correction unit can send a message to the mobile device 340 indicating that the mobile image was too out of focus to be used and requesting that the user retake the image.

[0237] The image correction unit can be configured to first identify the corners of a coupon or other document within a mobile image (step 1205). One technique that can be used to identify the corners of the remittance coupon in a color image is illustrated in FIG. 12 and is described in detail below. The corners of the document can be defined by a set of points A, B, C, and D that represent the corners of the document and define a quadrangle.

[0238] The image correction unit can be configured to then build a perspective transformation for the remittance coupon (step 1210). As can be seen in FIG. 8, the angle at which an image of a document is taken can cause the rectangular shape of the remittance coupon to appear distorted. FIG. 10 and its related description above provide some examples of how a perspective transformation can be constructed for a quadrangle defined by the corners A, B, C, and D according to an embodiment. For example, the quadrangle identified in step 1210 can be mapped onto a same-sized rectangle in order to build a perspective transformation that can be applied to the document subimage, i.e. the portion of the mobile image that corresponds to the remittance coupon, in order to correct perspective distortion present in the image.

[0239] A geometrical transformation of the document subimage can be performed using the perspective transformation built in step 1210 (step 1215). The geometrical transformation corrects the perspective distortion present in the document subimage. An example of results of geometrical transformation can be seen in FIG. 9 where a document subimage of the remittance coupon pictured in FIG. 8 has been geometrically corrected to remove perspective distortion.

[0240] A “dewarping” operation can also be performed on the document subimage (step 1220). An example of a warping of a document in a mobile image is provided in FIG. 38. Warping can occur when a document to be imaged is not perfectly flat or is placed on a surface that is not perfectly flat, causing distortions in the document subimage. A technique for identifying warping in a document subimage is illustrated in FIG. 39.

[0241] According to an embodiment, the document subimage can also be binarized (step 1225). A binarization operation can generate a bi-tonal image with color depth of 1 bit per a pixel (1 bit/pixel). Some automated processing systems, such as some Remote Deposit systems require bi-tonal images as inputs. A technique for generating a bi-tonal image is described below with respect to FIG. 13. FIG. 15 illustrates a binarized version of the geometrically corrected mobile

document image of the remittance coupon illustrated in FIG. 9. As illustrated, in the bi-tonal image of FIG. 15, the necessary information, such as payees, amounts, account number, etc., has been preserved, while extra information has been removed. For example, background patterns that might be printed on the coupon are not present in the bi-tonal image of the remittance coupon. Binarization of the subimage also can be used to remove shadows and other defects caused by unequal brightness of the subimage.

[0242] Once the image has been binarized, the code line of the remittance coupon can be identified and read (step 1230). As described above, many remittance coupons include a code line that comprises computer-readable text that can be used to encode account-related information that can be used to reconcile a payment received with the account for which the payment is being made. Code line 905 of FIG. 9 illustrates an example of code line on a remittance coupon.

[0243] Often, a standard optical character recognition font, the OCR-A font, is used for printing the characters comprising the code line. The OCR-A font is a fixed-width font where the characters are typically spaced 0.10 inches apart. Because the OCR-A font is a standardized fixed-width font, the image correction unit can use this information to determine a scaling factor for the image of the remittance coupon. The scaling factor to be used can vary from image to image, because the scaling is dependent upon the position of the camera or other image capture device relative to the document being imaged and can also be dependent upon optical characteristics of the device used to capture the image of the document. FIG. 23 illustrates a scaling method that can be used to determine a scaling factor to be applied according to an embodiment. The method illustrated in FIG. 23 is related to scaling performed on a MICR-line of a check, but can be used to determine a scaling factor for an image of a remittance coupon based on the size of the text in the code line of the image of the remittance coupon.

[0244] Once the scaling factor for the image has been determined, a final geometrical transformation of the document image can be performed using the scaling factor (step 1235). This step is similar to that in step 1215, except the scaling factor is used to create a geometrically altered subimage that represents the actual size of the coupon at a given resolution. According to an embodiment, the dimensions of the geometrically corrected image produced by set 635 are identical to the dimensions of an image produced by a flat bed scanner at the same resolution.

[0245] During step 1235, other geometrical corrections can also be made, such as correcting orientation of the coupon subimage. The orientation of the coupon subimage can be determined based on the orientation of the text of the code line.

[0246] Once the final geometrical transformation has been applied, a final adaptive binarization can be performed on the grayscale image generated in step 1235 (step 1240). The bi-tonal image output by this step will have the correct dimensions for the remittance coupon because the bi-tonal image is generated using the geometrically corrected image generated in step 1235.

[0247] According to an embodiment, the image correction unit can be configured to use several different binarization parameters to generate two or more bi-tonal images of the remittance coupon. The use of multiple images can improve

data capture results. The use of multiple bi-tonal images to improve data captures results is described in greater detail below.

Detecting Document within Color Mobile Image

[0248] Referring now to FIG. 13, a flowchart is provided illustrating an example method for automatic document detection within a color image from a mobile device. According to an embodiment, the method illustrated in FIG. 13 can be used to implement step 1205 of the method illustrated in FIG. 12. Typically, the operations described within method of FIG. 13 are performed within an automatic document detection unit of the remote server; however, embodiments exist where the operations reside in multiple units. In addition, generally the automatic document detection unit takes a variety of factors into consideration when detecting the document in the mobile image. The automatic document detection unit can take into consideration arbitrary location of the document within the mobile image, the 3-D distortions within the mobile image, the unknown size of the document, the unknown color of the document, the unknown color(s) of the background, and various other characteristics of the mobile engine, e.g. resolution, dimensions, etc.

[0249] The method of FIG. 13 begins at step 1502 by receiving the original color image from the mobile device. Upon receipt, this original color image is converted into a smaller color image, also referred to as a color "icon" image, at operation 1504. This color "icon" image preserves the color contrasts between the document and the background, while suppressing contrasts inside the document. A detailed description of an example conversion process is provided with respect to FIG. 16.

[0250] A color reduction operation is then applied to the color "icon" image at step 1506. During the operation, the overall color of the image can be reduced, while the contrast between the document and its background can be preserved within the image. Specifically, the color "icon" image of operation 1504 can be converted into a gray "icon" image (also known as a gray-scale "icon" image) having the same size. An example, color depth reduction process is described with further detail with respect to FIG. 18.

[0251] The corners of the document are then identified within the gray "icon" image (step 1310). As previously noted above with respect to FIG. 10, these corners A, B, C, and D make up the quadrangle ABCD (e.g. quadrangle ABCD 2502). Quadrangle ABCD, in turn, makes up the perimeter of the document. Upon detection of the corners, the location of the corners is outputted (step 1310).

Binarization

[0252] FIG. 14 illustrates a binarization method that can be used to generate a bi-tonal image from a document image according to an embodiment. The method illustrated in FIG. 10 can be used to implement the binarization step 1225 of the method illustrated in FIG. 12. In an embodiment, the steps of the method illustrated in FIG. 14 can be performed within unit of the remote server.

[0253] A binarization operation generates a bi-tonal image with color depth of 1 bit per a pixel (1 bit/pixel). In the case of documents, such as checks and deposit coupons, a bi-tonal image is required for processing by automated systems, such as Remote Deposit systems. In addition, many image processing engines require such an image as input. The method of FIG. 14 illustrates binarization of a gray-scale image of a document as produced by geometrical operation 1004. This

particular embodiment uses a novel variation of well-known Niblack's method of binarization. As such, there is an assumption that the gray-scale image received has the dimensions W pixel \times H pixels and an intensity function $I(x,y)$ gives the intensity of a pixel at location (x,y) in terms one of 256 possible gray-shade values (8 bit/pixel). The binarization operation will convert the 256 gray-shade value to a 2 shade value (1 bit/pixel), using an intensity function $B(x,y)$. In addition, to apply the method, a sliding window with dimensions w pixels \times h pixels is defined and a threshold T for local (in-window) standard deviation of gray image intensity $I(x,y)$ is defined. The values of w , h , and T are all experimentally determined.

[0254] A gray-scale image of the document is received at step 1402, the method 1400 chooses a pixel $p(x,y)$ within the image at step 1404. In FIG. 14, the average (mean) value ave and standard deviation σ of the chosen pixel's intensity $I(x,y)$ within the $w \times h$ current window location (neighborhood) of pixel $p(x,y)$ are computed (step 1406). If the standard deviation σ is determined to be too small at operation 1408 (i.e. $\sigma < T$), pixel $p(x,y)$ is considered to low-contrast and, thus, part of the background. Accordingly, at step 1410, low-contrast pixels are converted to white, i.e. set $B(x,y)$ set to 1, which is white; however, if the deviation σ is determined to be larger or equal to the threshold T , i.e. $\sigma \geq T$, the pixel $p(x,y)$ is considered to be part of the foreground. In step 1412, if $I(p) < \text{ave} - k * \sigma$, pixel p is considered to be a foreground pixel and therefore $B(x,y)$ is set to 0 (black). Otherwise, the pixel is treated as background and therefore $B(x,y)$ is set to 1. In the formula above, k is an experimentally established coefficient.

[0255] Subsequent to the conversion of the pixel at either step 1410 or operation 1412, the next pixel is chosen at step 1414, and operation 1406 is repeated until all the gray-scale pixels (8 bit/pixel) are converted to a bi-tonal pixel (1 bit/pixel). However, if no more pixels remain to be converted 1418, the bi-tonal image of the document is then outputted at step 1420.

Conversion of Color Image to Icon Image

[0256] Referring now to FIG. 16, a flowchart is provided describing an example method for conversion of a color image to a smaller "icon" image according to an embodiment. This method can be used to implement step 1304 of the method illustrated FIG. 13. The smaller "icon" image preserves the color contrasts between the document depicted therein and its background, while suppressing contrasts inside the document. Upon receipt of the original color image from the mobile device (step 1601), over-sharpening is eliminated within the image (step 1602). Accordingly, assuming the color input image I has the dimensions of $W \times H$ pixels, operation 1602 averages the intensity of image I and down-scales image I to image I' , such that image I' has dimensions that are half that of image I (i.e. $W' = W/2$ and $H' = H/2$). Under certain embodiments, the color transformation formula can be described as the following:

$$C(p') = \text{ave}\{C(q) : q \text{ in } S \times S\text{-window of } p\}, \text{ where} \quad (\text{eq. 5})$$

[0257] C is any of red, green or blue components of color intensity;

[0258] p' is any arbitrary pixel on image I' with coordinates (x',y') ;

[0259] p is a corresponding pixel on image I : $p = p(x,y)$, where $x = 2 * x'$ and $y = 2 * y'$;

[0260] q is any pixel included into $S \times S$ -window centered in p ;

[0261] S is established experimentally; and

[0262] ave is averaging over all q in the $S \times S$ -window.

[0263] Small “dark” objects within the image can then be eliminated (step 1604). Examples of such small “dark” objects include, but are not limited to, machine-printed characters and hand-printed characters inside the document. Hence, assuming operation 1604 receives image I' from step 1602, step 1604 creates a new color image I'' referred to as an “icon” with width W' set to a fixed small value and height H' set to $W' \cdot (H/W)$, thereby preserving the original aspect ratio of image I . In some embodiments, the transformation formula can be described as the following:

$$C(p'') = \max\{C(q') : q' \text{ in } S' \times S' \text{-window of } p'\}, \text{ where} \quad (\text{eq. 6})$$

[0264] C is any of red, green or blue components of color intensity;

[0265] p'' is an arbitrary pixel on image I'' ;

[0266] p' is a pixel on image I' which corresponds to p'' under similarity transformation, as previously defined;

[0267] q' is any pixel on image I' included into $S' \times S'$ -window centered in p' ;

[0268] \max is maximum over all q' in the $S' \times S'$ -window;

[0269] W' is established experimentally;

[0270] S' is established experimentally for computing the intensity I'' ; and

[0271] $I''(p'')$ is the intensity value defined by maximizing the intensity function I' (p') within the window of corresponding pixel p' on image I' , separately for each color plane.

The reason for using the “maximum” rather than “average” is to make the “icon” whiter (white pixels have a RGB-value of (255,255,255)).

[0272] In the next operation 1606, the high local contrast of “small” objects, such as lines, text, and handwriting on a document, is suppressed, while the other object edges within the “icon” are preserved. Often, these other object edges are bold. In various embodiments of the invention, multiple dilation and erosion operations, also known as morphological image transformations, are utilized in the suppression of the high local contrast of “small” objects. Such morphological image transformations are commonly known and used by those of ordinary skill in the art. The sequence and amount of dilation and erosion operations used is determined experimentally. Subsequent to the suppression operation 1606, a color “icon” image is outputted at operation 1608. FIG. 17B depicts an example of the mobile image of a check illustrated in FIG. 17A after being converted into a color “icon” image according to an embodiment.

Color Depth Reduction

[0273] Referring now to FIG. 18, a flowchart is provided illustrating an example method that provides further details with respect to the color depth reduction operation 1306 as illustrated in FIG. 13. At step 1301, a color “icon” image for color reduction is received. The color “icon” image is divided into a grid (or matrix) of fixed length and width with equal size grid elements at operation 1302. In some embodiments, the preferred grid size is such that there is a center grid element. For example, a grid size of 3×3 may be employed. FIG. 19A depicts an example of the color “icon” image of FIG. 19B after operation 1302 has divided it into a 3×3 grid in accordance with one embodiment of the invention.

[0274] Then, at step 1304, the “central part” of the icon, which is usually the center most grid element, has its color averaged. Next, the average color of the remaining parts of the icon is computed at step 1306. More specifically, the grid elements “outside” the “central part” of the “icon” have their colors averaged. Usually, in instances where there is a central grid element, e.g. 3×3 grid, the “outside” of the “central part” comprises all the grid elements other than the central grid element.

[0275] Subsequently, a linear transformation for the RGB-space is determined at step 1308. The linear transformation is defined such that it maps the average color of the “central part” computed during operation 1304 to white, i.e. 255, while the average color of the “outside” computed during operation 1306 maps to black, i.e. 0. All remaining colors are linearly mapped to a shade of gray. This linear transformation, once determined, is used at operation 1310 to transform all RGB-values from the color “icon” to a gray-scale “icon” image, which is then outputted at operation 1312. Within particular embodiments, the resulting gray “icon” image, also referred to as a gray-scale “icon” image, maximizes the contrast between the document background, assuming that the document is located close to the center of the image and the background. FIG. 15B depicts an example of the color “icon” image of FIG. 13B once it has been converted to a gray “icon” image in accordance with one embodiment.

[0276] Referring now to FIG. 20, a flowchart is provided illustrating an example method for finding document corners from a gray “icon” image containing a document. The method illustrated in FIG. 20 can be used to implement step 1308 of the method illustrated in FIG. 13. Upon receiving a gray “icon” image at operation 2001, the “voting” points on the gray “icon” image are found in step 2002 for each side of the document depicted in the image. Consequently, all positions on the gray “icon” image that could be approximated with straight line segments to represent left, top, right, and bottom sides of the document are found.

[0277] In accordance with one embodiment, this goal is achieved by first looking for the “voting” points in the half of the “icon” that corresponds with the current side of interest. For instance, if the current side of interest is the document’s top side, the upper part of the “icon” ($Y < H/2$) is examined while the bottom part of the “icon” ($Y \geq H/2$) is ignored.

[0278] Within the selected half of the “icon,” the intensity gradient (contrast) in the correct direction of each pixel is computed. This is accomplished in some embodiments by considering a small window centered in the pixel and, then, breaking the window into an expected “background” half where the gray intensity is smaller, i.e. where it is supposed to be darker, and into an expected “doc” half where the gray intensity is higher, i.e. where it is supposed to be whiter. There is a break line between the two halves, either horizontal or vertical depending on side of the document sought to be found. Next the average gray intensity in each half-window is computed, resulting in an average image intensity for the “background” and an average image intensity of the “doc.” The intensity gradient of the pixel is calculated by subtracting the average image intensity for the “background” from the average image intensity for the “doc.”

[0279] Eventually, those pixels with sufficient gray intensity gradient in the correct direction are marked as “voting” points for the selected side. The sufficiency of the actual gray intensity gradient threshold for determining is established experimentally.

[0280] Continuing with method **2000**, candidate sides, i.e. line segments that potentially represent the sides of the document, i.e. left, top, right, and bottom sides, are found. In order to do so, some embodiments find all subsets within the “voting” points determined in step **2002** that could be approximated by a straight line segment (linear approximation). In many embodiments, the threshold for linear approximation is established experimentally. This subset of lines is defined as the side “candidates.” As an assurance that the set of side candidates is never empty, the gray “icon” image’s corresponding top, bottom, left, and right sides are also added to the set.

[0281] Next, in step **2006** chooses the best candidate for each side of the document from the set of candidates selected in operation **2004**, thereby defining the position of the document within the gray “icon” image. In accordance with some embodiments, the following process is used in choosing the best candidate for each side of the document:

[0282] The process starts with selecting a quadruple of line segments $\{L, T, R, B\}$, where L is one of the candidates for the left side of the document, T is one of the candidates for the top side of the document, R is one of the candidates for the right side of the document, and B is one of the candidates for the bottom side of the document. The process then measures the following characteristics for the quadruple currently selected.

[0283] The amount of “voting” points is approximated and measured for all line segments for all four sides. This amount value is based on the assumption that the document’s sides are linear and there is a significant color contrast along them. The larger values of this characteristic increase the overall quadruple rank.

[0284] The sum of all intensity gradients over all voting points of all line segments is measured. This sum value is also based on the assumption that the document’s sides are linear and there is a significant color contrast along them. Again, the larger values of this characteristic increase the overall quadruple rank.

[0285] The total length of the segments is measured. This length value is based on the assumption that the document occupies a large portion of the image. Again, the larger values of this characteristic increase the overall quadruple rank.

[0286] The maximum of gaps in each corner is measured. For example, the gap in the left/top corner is defined by the distance between the uppermost point in the L-segment and the leftmost point in the T-segment. This maximum value is based on how well the side-candidates suit the assumption that the document’s shape is quadrangle. The smaller values of this characteristic increase the overall quadruple rank.

[0287] The maximum of two angles between opposite segments, i.e. between L and R, and between T and R, is measured. This maximum value is based on how well the side-candidates suit the assumption that the document’s shape is close to parallelogram. The smaller values of this characteristic increase the overall quadruple rank.

[0288] The deviation of the quadruple’s aspect ratio from the “ideal” document aspect ratio is measured. This characteristic is applicable to documents with a known aspect ratio, e.g. checks. If the aspect ratio is unknown, this characteristic should be excluded from computing the quadruple’s rank. The quadruple’s aspect ratio is computed as follows:

[0289] a) Find the quadrangle by intersecting the quadruple’s elements;

[0290] b) Find middle-point of each of the four quadrangle’s sides;

[0291] c) Compute distances between middle-points of opposite sides, say D1 and D2;

[0292] d) Find the larger of the two ratios: $R = \max(D1/D2, D2/D1)$;

[0293] e) Assuming that the “ideal” document’s aspect ratio is known and Min/MaxAspectRatio represent minimum and maximum of the aspect ratio respectively, define the deviation in question as:

[0294] 0, if $\text{MinAspectRatio} \leq R \leq \text{MaxAspectRatio}$

[0295] $\text{MinAspectRatio} - R$, if $R < \text{MinAspectRatio}$

[0296] $R - \text{MaxAspectRatio}$, if $R > \text{MaxAspectRatio}$.

[0297] f) For checks, MinAspectRatio can be set to 2.0 and MaxAspectRatio can be set to 3.0.

This aspect ratio value is based on the assumption that the document’s shape is somewhat preserved during the perspective transformation. The smaller values of this characteristic increase the overall quadruple rank.

[0298] Following the measurement of the characteristics of the quadruple noted above, the quadruple characteristics are combined into a single value, called the quadruple rank, using weighted linear combination. Positive weights are assigned for the amount of “voting” points, the sum all of intensity gradients, and the total length of the segments. Negative weights are assigned for maximum gaps in each corner, maximum two angles between opposite segments, and the deviation of the quadruple’s aspect ratio. The exact values of each of the weights are established experimentally.

[0299] The operations set forth above are repeated for all possible combinations of side candidates, eventually leading to the “best” quadruple, which is the quadruple with the highest rank. The document’s corners are defined as intersections of the “best” quadruple’s sides, i.e. the best side candidates.

[0300] In, step **2008** the corners of the document are defined using the intersections of the best side candidates. A person of ordinary skill in the art would appreciate that these corners can then be located on the original mobile image by transforming the corner locations found on the “icon” using the similarity transformation previously mentioned. Method **2000** concludes at step **2010** where the locations of the corners defined in step **2008** are output.

Geometric Correction

[0301] FIG. 21 provides a flowchart that illustrates an example method for geometric correction in accordance with the invention according to an embodiment. According to an embodiment, the method illustrated in FIG. 21 can be used to implement steps **1210**, **1215**, and **1235** of the method illustrated in FIG. 12. As previously mentioned, geometric correction is needed to correct any possibly perspective distortions that exist in the original mobile image. Additionally, geometric correction can correct the orientation of the documentation within the original mobile image, e.g. document is orientated at 90, 180, or 270 degrees where the right-side-up orientation is 0 degrees. It should be noted that in some embodiments, the orientation of the document depends on the type of document depicted in the mobile image, as well as the fields of relevance on the document.

[0302] In instances where the document is in landscape orientation (90 or 270 degrees), as illustrated by the check in FIG. 22A, geometric correction is suitable for correcting the orientation of the document. Where the document is at 180 degree orientation, detection of the 180 degree orientation and its subsequent correction are suitable when attempting to

locate an object of relevance on the document. A codeline for a remittance coupon can be located in various locations on the remittance coupon, and might not be located along the bottom of the coupon. The ability to detect a codeline in an image of the remittance coupon changes significantly after the document has been rotated 180-degrees. In contrast, the MICR-line of check is generally known to be at a specific location along the bottom of the document, and the MICR-line can be used to determine the current orientation of the check within the mobile image. In some embodiments, the object of relevance on a document depends on the document's type. For example, where the document is a contract, the object of relevance may be a notary seal, signature, or watermark positioned at a known position on the contract. Greater detail regarding correction of a document (specifically, a check) having upside-down orientation (180 degree orientation) is provided with respect to FIG. 23.

[0303] According to some embodiments, a mathematical model of projective transformations is built and converts the distorted image into a rectangle-shaped image of predefined size. According to an embodiment, this step corresponds to step 1210 of FIG. 12. In an example, where the document depicted in mobile image is a check, the predefined size is established as 1200x560 pixels, which is roughly equivalent to the dimensions of a personal check scanned at 200 DPI. In other embodiments, where the document depicted is a remittance coupon, the size of the remittance coupons may not be standardized. However, the size and spacing of the characters comprising the code line can be used to determine a scaling factor to be applied to the image to correct the size of the image of the remittance coupon relative to a specific resolution.

[0304] Continuing with reference to the method of FIG. 21, there are two separate paths of operations that are either performed sequentially or concurrently, the outputs of which are eventually utilized in the final output. One path of operations begins at step 1504 where the original mobile image in color is received. In step 1508, the color depth of the original mobile image is reduced from a color image with 24 bit per pixel (24 bit/pixel) to a gray-scale image with 8 bit per pixel (8 bit/pixel). This image is subsequently outputted to step 1516 as a result of step 1512.

[0305] The other path of operations begins at step 1502, where the positions of the document's corners within the gray "icon" image are received. Based off the location of the corners, the orientation of the document is determined and the orientation is corrected (step 1506). In some embodiments, this operation uses the corner locations to measure the aspect ratio of the document within the original image. Subsequently, a middle-point between each set of corners can be found, wherein each set of corners corresponds to one of the four sides of the depicted document, resulting in the left (L), top (T), right (R), and bottom (B) middle-points (step 1506). The distance between the L to R middle-points and the T to B middle points are then compared to determine which of the two pairs has the larger distance. This provides step 1506 with the orientation of the document.

[0306] In some instances, the correct orientation of the document depends on the type of document that is detected. For example, as illustrated in FIG. 22A, where the document of interest is a check, the document is determined to be in landscape orientation when the distance between the top middle-point and bottom middle-point is larger than the dis-

tance between the left middle-point and the right middle-point. The opposite might be true for other types of documents.

[0307] If it is determined in step 1506 that an orientation correction is necessary, then the corners of the document are shifted in a loop, clock-wise in some embodiments and counter-clockwise in other embodiments.

[0308] At step 1510, the projective transformation is built to map the image of the document to a predefined target image size of width of W pixels and height of H pixels. In some embodiments, the projective transformation maps the corners A, B, C, and D of the document as follows: corner A to (0,0), corner B to (W,0), corner C to (W,H), and corner D to (0,H). Algorithms for building projective transformation are commonly known and used amongst those of ordinary skill in the art.

[0309] At step 1516, the projective transformation created during step 1514 is applied to the mobile image in gray-scale as outputted as a result of step 1512. The projective transformation as applied to the gray-scale image of step 1512 results in all the pixels within the quadrangle ABCD depicted in the gray-scale image mapping to a geometrically corrected, gray-scale image of the document alone. FIG. 22B is an example gray-scale image of the document depicted in FIG. 17A once a geometrical correction operation in accordance with the invention is applied thereto. The process concludes at operation 1518 where the gray-scale image of the document is outputted to the next operation.

Correcting Landscape Orientation

[0310] FIG. 23 is a flow chart illustrating a method for correcting landscape orientation of a document image according to an embodiment. As previously noted, the geometric correction operation as described in FIG. 21 is one method in accordance with the invention for correcting a document having landscape orientation within the mobile image. However, even after the landscape orientation correction, the document still may remain in upside-down orientation. In order to the correct upside-down orientation for certain documents, some embodiments of the invention require the image containing the document be binarized beforehand. Hence, the orientation correction operation included in step 1235 usually follows the binarization operation of 1225. While the embodiment described herein uses the MICR-line of a check or determine the orientation of an image, the code line of a remittance coupon can be used to determine the orientation of a remittance coupon using the technique described herein.

[0311] Upon receiving the bi-tonal image of the check at operation 1702, the MICR-line at the bottom of the bi-tonal check image is read at operation 1704 and an MICR-confidence value is generated. This MICR-confidence value (MC1) is compared to a threshold value T at operation 1706 to determine whether the check is right-side-up. If $MC1 > T$ at operation 1708, then the bi-tonal image of the check is right side up and is outputted at operation 1710.

[0312] However, if $MC1 \leq T$ at operation 1708, then the image is rotated 180 degrees at operation 1712, the MICR-line at the bottom read again, and a new MICR-confidence value generated (MC2). The rotation of the image by 180 degree is done by methods commonly-known in the art. The MICR-confidence value after rotation (MC2) is compared to the previous MICR-confidence value (MC1) plus a Delta at operation 1714 to determine if the check is now right-side-up.

If $MC2 > MC2 + \Delta$ at operation **1716**, the rotated bi-tonal image has the check right-side-up and, thus, the rotated image is outputted at operation **1718**. Otherwise, if $MC2 \leq MC2 + \Delta$ at operation **1716**, the original bi-tonal image of the check is right-side-up and outputted at operation **1710**. Δ is a positive value selected experimentally that reflects a higher a priori probability of the document initially being right-side-up than upside-down.

Size Correction

[0313] FIG. **24** provides a flowchart illustrating an example method for size correction of an image according to an embodiment. The method of FIG. **24** can be used to implement the size correction step described in relation to step **1230** of FIG. **12**. Specifically, FIG. **24** illustrates an example method, in accordance with one embodiment, for correcting the size of a remittance coupon within a bi-tonal image, where the remittance coupon is oriented right-side-up. A person of ordinary skill in the art would understand and appreciate that this method can operate differently for other types of documents, e.g. deposit coupons, remittance coupons.

[0314] Since many image processing engines are sensitive to image size, it is crucial that the size of the document image be corrected before it can be properly processed. For example, a form identification engine may rely on the document size as an important characteristic for identifying the type of document that is being processed. Generally, for documents such as remittance coupons, the image size should be equivalent to the image size produced by a standard scanner running at 200 DPI.

[0315] In addition, where the document is a remittance coupon, the size of the remittance coupons vary widely across different biller. Hence, in order to restore the size of remittance coupons that have been geometrically corrected in accordance with the invention at a predefined image size of 1200x560 pixels, the size correction operation is performed.

[0316] Referring now to FIG. **24**, after receiving a bi-tonal image containing a remittance coupon that is orientated right-side-up at operation **1802**, the codeline at the bottom of the remittance coupon is read at operation **1804**. This allows the average width of the codeline characters to be computed at operation **1806**. In doing so, the computer average width gets compared to the average size of a codeline character at 200 DPI at operation **1808**, and a scaling factor is computed accordingly. In some embodiments of the invention, the scaling factor SF is computer as follows:

$$SF = AW_{200} / AW, \text{ where} \quad (\text{eq. 7})$$

[0317] AW is the average width of the MICR-character found; and

[0318] $A W_{200}$ is the corresponding "theoretical" value based on the ANSI x9.37 standard (Specifications for Electronic Exchange of Check and Image Data) at 200 DPI.

[0319] The scaling factor is used at operation **1810** to determine whether the bi-tonal image of the remittance coupon requires size correction. If the scaling SF is determined to be less than or equal to $1.0 + \Delta$, then the most recent versions of the remittance coupon's bi-tonal image and the remittance coupon's the gray-scale image are output at operation **1812**. Δ defines the system's tolerance to wrong image size.

[0320] If, however, the scaling factor SF is determined to be higher than $1.0 + \Delta$, then at operation **1814** the new dimensions of the remittance coupon are computed as follows:

$$AR = H_S / W_S \quad (\text{eq. 8})$$

$$W' = W * SF \quad (\text{eq. 9})$$

$$H' = AR * W', \text{ where} \quad (\text{eq. 10})$$

[0321] H_S and W_S are the height and width of the remittance coupon snippet found on the original image;

[0322] AR is the remittance coupon aspect ratio which we want to maintain while changing the size;

[0323] W is the width of geometrically corrected image before it's size is adjusted;

[0324] W' is the adjusted remittance coupon's width in pixels; and

[0325] H' is the adjusted remittance coupon's height in pixels.

Subsequent to re-computing the new dimensions, operation **1814** repeats geometrical correction and binarization using the newly dimensioned remittance coupon image. Following the repeated operations, operation **1812** outputs the resulting bi-tonal image of the remittance coupon and gray-scale image of the remittance coupon.

Image Quality Assurance

[0326] Once the remote server has processed a mobile image (see step **S216** of the method illustrated in FIG. **2**), the remote server can be configured to perform image quality assurance processing on the mobile image to determine whether the quality of the image is sufficient to submit to banking server **112**.

[0327] FIG. **25** illustrates a mobile document image processing engine (MDIPE) unit **2100** for performing quality assurance testing on mobile document images according to an embodiment. The MDIPE unit **2100** can receive a mobile document image captured by a mobile device, or multiple mobile images for some tests; perform preprocessing on the mobile document image; select tests to be performed on the mobile document image; and execute the selected tests to determine whether the quality of the image of a high enough quality for a particular mobile application. The MDIPE unit **2100** includes a preprocessing unit **2110** and test execution unit **2130**. The preprocessing unit **2110** can be configured to receive a mobile image **2105** captured using a camera of a mobile device as well as processing parameters **2107**. According to an embodiment, the mobile image **2105** and the processing parameters **2107** can be passed to MDIPE **2100** by a mobile application on the mobile device where the mobile application provides the mobile image **2105** to the MDIPE **2100** to have the quality of the mobile image **2105** assessed.

[0328] The processing parameters **2107** can include various information that the MDIPE **2100** can use to determine which tests to run on the mobile image **2105**. For example, the processing parameters **2107** can identify the type of device used to capture the mobile image **2105**, the type of mobile application that will be used to process the mobile image if the mobile image passes the IQA testing, or both. The MDIPE **2100** can use this information to determine which tests to select from test data store **2132** and which test parameters to select from test parameter data store **2134**. For example, if a mobile image is being tested for a mobile deposit application that expects an image of a check, a specific set of tests related to assessing the image quality for a mobile image of a check can be selected, such as an MICR-line test, or a test for whether an image is blurry, etc. The MDIPE **2100** can also select test parameters from test parameters data store **2134**

that are appropriate for the type of image to be processed, or for the type of mobile device that was used to capture the image, or both. In an embodiment, different parameters can be selected for different mobile phones that are appropriate for the type of phone used to capture the mobile image. For example, some mobile phones might not include an autofocus feature.

[0329] The preprocessing unit **2110** can process the mobile document image to extract a document snippet that includes the portion of the mobile document that actually contains the document to be processed. This portion of the mobile document image is also referred to herein as the document subimage. The preprocessing unit **2110** can also perform other processing on the document snippet, such as converting the image to a grayscale or bi-tonal document snippet, geometric correction of the document subimage to remove view distortion, etc. Different tests can require different types of preprocessing to be performed, and the preprocessing unit **2110** can produce mobile document snippets from a mobile document image depending on the types of mobile IQA tests to be executed on the mobile document image.

[0330] The test execution unit **2130** receives the selected tests and test parameters **2112** and the preprocessed document snippet (or snippets) **120** from the preprocessing mobile **110**. The test execution unit **2130** executes the selected tests on the document snippet generated by the preprocessing unit **2110**. The test execution unit **2130** also uses the test parameters provided by the preprocessing unit **2110** when executing the test on the document snippet. The selected tests can be a series of one or more tests to be executed on the document snippets to determine whether the mobile document image exhibits geometrical or other defects.

[0331] The test execution unit **2130** executes each selected test to obtain a test result value for that test. The test execution unit **2130** then compares that test result value to a threshold value associated with the test. If the test result value is equal to or exceeds the threshold, then the mobile image has passed the test. Otherwise, if the test result value is less than the threshold, the mobile document image has failed the test. According to some embodiments, the test execution unit **2130** can store the test result values for the tests performed in test results data store **2138**.

[0332] According to an embodiment, the test threshold for a test can be stored in the test parameters data store **2134** and can be fetched by the preprocessing unit **2110** and included with the test parameters **2112** provided to the test execution unit **2130**. According to an embodiment, different thresholds can be associated with a test based on the processing parameters **2107** received by the preprocessing unit **2110**. For example, a lower threshold might be used for an image focus IQA test for image capture by camera phones that do not include an autofocus feature, while a higher threshold might be used for the image focus IQA test for image capture by camera phones that do include an autofocus feature.

[0333] According to an embodiment, a test can be flagged as “affects overall status.” These tests are also referred to here as “critical” tests. If a mobile image fails a critical test, the MDIPE **2100** rejects the image and can provide detailed information to the mobile device user explaining why the image was not of a high enough quality for the mobile application and that provides guidance for retaking the image to correct the defects that caused the mobile document image to fail the test, in the event that the defect can be corrected by retaking the image.

[0334] According to an embodiment, the test result messages provided by the MDIPE **2100** can be provided to the mobile application that requested the MDIPE **2100** perform the quality assurance testing on the mobile document image, and the mobile application can display the test results to the user of the mobile device. In certain embodiments, the mobile application can display this information on the mobile device shortly after the user takes the mobile document image to allow the user to retake the image if the image is found to have defects that affect the overall status of the image. In some embodiments, where the MDIPE **2100** is implemented at least in part on the mobile device, the MDIPE **2100** can include a user interface unit that is configured to display the test results message on a screen of the mobile device.

[0335] FIG. **25** merely provides a description of the logical components of the MDIPE **2100**. In some embodiments, the MDIPE **2100** can be implemented on the mobile device **340**, in software, hardware, or a combination thereof. In other embodiments, the MDIPE **2100** can be implemented on the remote server, and the mobile device can send the mobile image **2105** and the processing parameters **2107**, e.g., via a wireless interface, to the remote server for processing, and the remote server can send the test results and test messages **2140** to the mobile device to indicate whether the mobile image passed testing. In some embodiments, part of the functionality of the MDIPE **2100** can be implemented on the mobile device while other parts of the MDIPE **2100** are implemented on the remote server. The MDIPE **2100** can be implemented in software, hardware, or a combination thereof. In still other embodiments, the MDIPE **2100** can be implemented entirely on the remote server, and can be implemented using appropriate software, hardware, or a combination there.

[0336] FIG. **26** is a flow diagram of a process for performing mobile image quality assurance on an image captured by a mobile device according to an embodiment. The process illustrated in FIG. **26** can be performed using the MDIPE **2100** illustrated in FIG. **25**.

[0337] The mobile image **2105** captured by a mobile device is received (step **2205**). The mobile image **2105** can also be accompanied by one or more processing parameters **2107**.

[0338] As described above, the MDIPE **2100** can be implemented on the mobile device, and the mobile image can be provided by a camera that is part of or coupled to the mobile device. In some embodiments, the MDIPE **2100** can also be implemented at least in part on a remote server, and the mobile image **2105** and the processing parameters **2107** can be transmitted to the remote server, e.g., via a wireless interface included in the mobile device.

[0339] Once the mobile image **2105** and the processing parameters **2107** have been received, the mobile image is processed to generate a document snippet or snippets (step **2210**). For example, preprocessing unit **2110** of MDIPE **2100** can be used to perform various preprocessing on the mobile image. One part of this preprocessing includes identifying a document subimage in the mobile image. The subimage is the portion of the mobile document image that includes the document. The preprocessing unit **2110** can also perform various preprocessing on the document subimage to produce what is referred to herein as a “snippet.” For example, some tests can require that a grayscale image of the subimage be created. The preprocessing unit **2110** can create a grayscale snippet that represents a grayscale version of the document subimage. In another example, some tests can require that a bitonal image of the subimage be created. The preprocessing unit **2110** can

create a bitonal snippet that represents a bitonal version of the document subimage. In some embodiments, the MDIPE 2100 can generate multiple different snippets based on the types of tests to be performed on the mobile document image.

[0340] After processing the mobile document image to generate a snippet, the MDIPE 2100 then selects one or more tests to be performed on the snippet or snippets (step 2215). In an embodiment, the tests to be performed can be selected from test data store 2132. In an embodiment, the MDIPE 2100 selects the one or more tests based on the processing parameters 2107 that were received with the mobile image 2105.

[0341] After selecting the tests from the test data store 2132, test parameters for each of the tests can be selected from the test parameters data store 2134 (step 2220). According to an embodiment, the test parameters can be used to configure or customize the tests to be performed. For example, different test parameters can be used to configure the tests to be more or less sensitive to certain attributes of the mobile image. In an embodiment, the test parameters can be selected based on the processing parameters 2107 received with the mobile image 2105. As described above, these processing parameters can include information, such as the type of mobile device used to capture the mobile image as well as the type of mobile application that is going to be used to process the mobile image if the mobile image passes scrutiny of the mobile image IQA system.

[0342] Once the tests and the test parameters have been retrieved and provided to the test execution unit 2130, a test is selected from tests to be executed, and the test is executed on the document snippet to produce a test result value (step 2225). In some embodiments, more than one document snippet may be used by a test. For example, a test can be performed that tests whether images of a front and back of a check are actually images of the same document can be performed. The test engine can receive both an image of the front of the check and an image of the back of the check from the preprocessing unit 2110 and use both of these images when executing the test.

[0343] The test result value obtained by executing the test on the snippet or snippets of the mobile document is then compared to test threshold to determine whether the mobile image passes or fails the test (step 2230) and a determination is made whether the test results exceed the threshold (step 2235). According to an embodiment, the test threshold can be configured or customized based on the processing parameters 2107 received with the mobile image. For example, the test for image blurriness can be configured to use a higher threshold for passing if the image is to be used to for a mobile deposit application where the MICR-line information needs to be recognized and read from the document image. In contrast, the test for blurriness can be configured use a lower threshold for passing the mobile image for some mobile applications. For example, the threshold for image quality may be lowered for if a business card is being imaged rather than a check. The test parameters can be adjusted to minimize the number of false rejects and false accept rate, the number of images marked for reviewing, or both.

[0344] The “affects overall status” flag of a test can also be configured based on the processing parameters 2107. For example, a test can be marked as not affecting the overall status for some types of mobile applications or for documents being processed, or both. Alternatively, a test can also be marked as affecting overall status for other types of mobile

applications or documents being processed, or both. For example, a test that identifies the MICR-line of a check can be marked as “affecting overall status” so that if the MICR-line on the check cannot be identified in the image, the image will fail the test and the image will be rejected. In another example, if the mobile application is merely configured to receive different types of mobile document image, the mobile application can perform a MICR-line test on the mobile document image in an attempt to determine whether the document that was imaged was a check. In this example, the MICR-line may not be present, because a document other than a check may have been imaged. Therefore, the MICR-line test may be marked as not “affecting overall status,” and if a document fails the test, the transaction might be flagged for review but not marked as failed.

[0345] Since different camera phones can have cameras with very different optical characteristics, image quality may vary significantly between them. As a result, some image quality defects may be avoidable on some camera phones and unavoidable on the others and therefore require different configurations. To mitigate the configuration problem, Mobile IQA test can be automatically configured for different camera phones to use different tests, or different thresholds for the tests, or both. For example, as described above, a lower threshold can be used for an image focus IQA test on mobile document images that are captured using a camera phone that does not include an autofocus feature than would be used for camera phones that do include an autofocus feature, because it can be more difficult for a user to obtain as clear an image on using a device that doesn't an autofocus feature.

[0346] In certain embodiments, if the test result exceeded or equaled the threshold, the image passed the test and a determination is made whether there are more tests to be executed (step 2240). If there are more tests to be executed, the next test can be selected and executed on the document snippet (step 2225). Otherwise, if there were not more tests to be executed, the test results, or test messages, or both are output by MDIPE 2100 (step 2270). There can be one or more test messages included with the results if the mobile image failed one more of the tests that were executed on the image.

[0347] In such embodiments, if the test result was less than the threshold, then the mobile image has failed the test. A determination is made whether the test affects the overall status (step 250). If the test affects the overall status of the image, detailed test result messages that explain why the image failed the test can be loaded from the test message data store 134 (step 2255) and the test result messages can be added to the test results (step 2260). The test results and test messages can then be output by the MDIPE 2100 (step 2270).

[0348] Alternatively, if the test did not affect the overall status, the test results can be loaded noted and the transaction can be flagged for review (step 2265). By flagging the transaction for review, a user of a mobile device can be presented with information indicating that a mobile image has failed at least some of the test that were performed on the image, but the image still may be of sufficient quality for use with the mobile application. The user can then be presented with the option to retake the image or to send the mobile image to the mobile application for processing. According to some embodiments, detailed test messages can be loaded from the test message data store 134 for all tests that fail and can be included with the test results, even if the test is not one that affects the overall status of the mobile image.

[0349] According to some embodiments, the mobile IQA test can also be configured to eliminate repeated rejections of a mobile document. For example, if an image of a check is rejected as having too low a contrast by a contrast test, the image is rejected, and the user can retake and resubmit the image via the mobile application, the processing parameters **2107** received with the mobile image can include a flag indicating that the image is being resubmitted. In some embodiments, the thresholds associated with the tests that the image failed can be lowered to see if the image can pass the test with a lower threshold. In some embodiments, the thresholds are only lowered for non-critical tests. According to an embodiment, the processing parameters **2107** can also include a count of the number of times that an image has been resubmitted and the thresholds for a test are only lowered after a predetermined number of times that the image is resubmitted.

[0350] FIG. 27 is a flow diagram of a process for performing mobile image quality assurance on an image of a check captured by a mobile device according to an embodiment. Like the process illustrated in FIG. 26, the process illustrated in FIG. 27 can be performed using the MDIPE **2100** illustrated in FIG. 25. The method illustrated in FIG. 27 can be used where an image of a check is captured in conjunction with a remittance payment. The method illustrated in FIG. 27 can be used to assess the quality of the image of the check.

[0351] The method illustrated in FIG. 27 illustrates how the mobile IQA and MDIPE **2100** can be used with the electronic check processing provided under the Check Clearing for the 21st Century Act. The Check Clearing for the 21st Century Act (also referred to as the “Check 21 Act”) is a United States federal law (Pub.L. 108-100) that was enacted on Oct. 28, 2003. The law allows the recipient of a paper check to create a digital version of the original check called a “substitute check,” which can be processed, eliminating the need to process the original physical document. The substitute check includes an image of the front and back sides of the original physical document. The mobile IQA tests can be used to check the quality of the images captured by a mobile device. The snippets generated by the MDIPE **2100** can then be further tested by one or more Check 21 mobile IQA tests that perform image quality assurance on the snippets to determine whether the images meet the requirements of the Check 21 Act as well.

[0352] The mobile image **2105** captured by a mobile device is received (step **2305**). In an embodiment, image of the front and back sides of the check can be provided. The mobile image **2105** can also be accompanied by one or more processing parameters **2107**. Check data can also be optionally received (step **2307**). The check data can be optionally provided by the user at the time that the check is captured. This check data can include various information from the check, such as the check amount, check number, routing information from the face of the check, or other information, or a combination thereof. In some embodiments, a mobile deposition application requests this information from a user of the mobile device, allows the user to capture an image of a check or to select an image of a check that has already been captured, or both, and the mobile deposit information provides the check image, the check data, and other processing parameters to the MDIPE **2100**.

[0353] Once the mobile image **2105**, the processing parameters **2107**, and the check data have been received, the mobile image is processed to generate a document snippet or snippets (step **2310**). As described above, the preprocessing can produce one or more document snippets that include the portion

of the mobile image in which the document was located. The document snippets can also have additional processing performed on them, such as conversion to a bitonal image or to grayscale, depending on the types of testing to be performed.

[0354] After processing the mobile document image to generate a snippet, the MDIPE **2100** then selects one or more tests to be performed on the snippet or snippets (step **2315**). In an embodiment, the tests to be performed can be selected from test data store **2132**. In an embodiment, the MDIPE **2100** selects the one or more tests based on the processing parameters **2107** that were received with the mobile image **2105**.

[0355] After selecting the tests from the test data store **2132**, test parameters for each of the tests can be selected from the test parameters data store **2134** (step **2320**). As described above, the test parameters can be used to configure or customize the tests to be performed.

[0356] Once the tests and the test parameters have been retrieved and provided to the test execution unit **2130**, a test is selected from tests to be executed, and the test is executed on the document snippet to produce a test result value (step **2325**). In some embodiments, more than one document snippet can be used by a test. For example, a test can be performed that tests whether images of a front and back of a check are actually images of the same document can be performed. The test engine can receive both an image of the front of the check and an image of the back of the check from the preprocessing unit **2110** and use both of these images when executing the test. Step **2325** can be repeated until each of the tests to be executed is performed.

[0357] The test result values obtained by executing each test on the snippet or snippets of the mobile document are then compared to test threshold with that test to determine whether the mobile image passes or fails the test (step **2330**) and a determination can be made whether the mobile image of the check passed the test indicating that image quality of mobile image is acceptable (step **2335**). If the mobile document image of the check passed, the MDIPE **2100** passes then executes one or more Check 21 tests on the snippets (step **2340**).

[0358] The test result values obtained by executing the Check 21 test or tests on the snippet or snippets of the mobile document are then compared to test threshold with that test to determine whether the mobile image passes or fails the test (step **2345**) and a determination can be made whether the mobile image of the check passed the test indicating that image quality of mobile image is acceptable under the requirements imposed by the Check 21 Act (step **2350**). Step **345** can be repeated until each of the Check 21 tests is performed. If the mobile document image of the check passed, the MDIPE **2100** passes the snippet or snippets to the mobile application for further processing (step **2370**).

[0359] If the mobile document image of the check failed one or more mobile IQA or Check **21** tests, detailed test result messages that explain why the image failed the test can be loaded from the test message data store **134** (step **2355**) and the test result messages can be added to the test results (step **2360**). The test results and test messages are then output to the mobile application where they can be displayed to the user (step **2365**). The user can use this information to retake the image of the check in an attempt to remedy some or all of the factors that caused the image of the check to be rejected.

Mobile IQA Tests

[0360] FIGS. 28A-41 illustrate various sample mobile document images and various testing methods that can be performed when assessing the image quality of a mobile document image. As described above, the preprocessing unit 2110 can be configured to extract the document subimage, also referred to herein as the subimage, from the mobile document image. The subimage generally will be non-rectangular because of perspective distortion; however, the shape of the subimage can generally be assumed to be quadrangular, unless the subimage is warped. Therefore, the document can be identified by its four corners.

[0361] In some embodiments, a mobile IQA test generates a score for the subimage on a scale that ranges from 0-1000, where “0” indicates a subimage having very poor quality while a score of “1000” indicates that the image is perfect according to the test criteria.

[0362] Some tests use a geometrically corrected snippet of the subimage to correct view distortion. The preprocessing unit 2110 can generate the geometrically corrected snippet. FIG. 28A illustrates a mobile image where the document captured in the mobile document image exhibits view distortion. FIG. 28B illustrates an example of a grayscale geometrically corrected subimage generated from the distorted image in FIG. 28A.

Image Focus IQA Test

[0363] According to some embodiments, an Image Focus IQA Test can be executed on a mobile image to determine whether the image is too blurry to be used by a mobile application. Blurry images are often unusable, and this test can help to identify such out-of-focus images and reject them. The user can be provided detailed information to assist the user in taking a better quality image of the document. For example, the blurriness may have been the result of motion blur caused by the user moving the camera while taking the image. The test result messages can suggest that the user hold the camera steadier when retaking the image.

[0364] Mobile devices can include cameras that have significantly different optical characteristics. For example, a mobile device that includes a camera that has an auto-focus feature can generally produce much sharper images than a camera that does not include such a feature. Therefore, the average image focus score for different cameras can vary widely. As a result, the test threshold can be set differently for different types of mobile devices. As described above, the processing parameters 2107 received by MDIPE 2100 can include information that identifies the type of mobile device and/or the camera characteristics of the camera used with the device in order to determine what the threshold should be set to for the Image Focus IQA Test.

[0365] An in-focus mobile document image, such as that illustrated in FIG. 29A will receive a score of 1000, while an out of focus document, such as that illustrated in FIG. 29B will receive a much lower score, such as in the 50-100 range. Most of the time, images are not completely out of focus. Therefore, a score of 0 is uncommon.

[0366] According to an embodiment, the focus of the image can be tested using various techniques, and the results can then be normalized to the 0-1000 scale used by the MDIPE 2100.

[0367] In an embodiment, the Image Focus Score can be computed using the following technique: The focus measure

is a ratio of maximum video gradient between adjacent pixels, measured over the entire image and normalized with respect to image’s gray level dynamic range and “pixel pitch.” According to an embodiment, the image focus score can be calculated using the following equation described in “The Financial Services Technology Consortium,” Image Defect Metrics, IMAGE QUALITY & USABILITY ASSURANCE: Phase 1. Project, Draft Version 1.0.4. May 2, 2005, which is hereby incorporated by reference:

$$\text{Image Focus Score} = \frac{\text{Maximum Video Gradient}}{[(\text{Gray Level Dynamic Range}) * (\text{Pixel Pitch})]}$$

$$\text{where Video Gradient} = \text{ABS} [(\text{Gray level for pixel "i"}) - (\text{Gray level for pixel "i+1"})]$$

$$\text{Gray Level Dynamic Range} = [(\text{Average of the "N" Lightest Pixels}) - (\text{Average of the "N" Darkest Pixels})]$$

$$\text{Pixel Pitch} = [1 / \text{Image Resolution (in dpi)}]$$

[0368] The variable N is equal to the number of pixels used to determine the average darkest and lightest pixel gray levels in the image. According to one embodiment, the value of N is set to 64. Therefore, the 64 lightest pixels in the image are averaged together and the 64 darkest pixels in the image are averaged together, to compute the “Gray Level Dynamic” range value. The resulting image focus score value is the multiplied by 10 in order to bring the value into the 0-1000 range used for the test results in the mobile IQA system.

[0369] The Image Focus Score determined using these techniques can be compared to an image focus threshold to determine whether the image is sufficiently in focus. As described above, the threshold used for each test may be determined at least in part by the processing parameters 2107 provided to MDIPE 2100. The Image Focus score can be normalized to the 0-1000 range used by the mobile IQA tests and compared to a threshold value associated with the test. If the Image Focus Score meets or exceeds this threshold, then the mobile document image is sufficiently focused for use with the mobile application.

Shadow Test

[0370] Shadows frequently occur on mobile photos taken in bright sunlight, where an object obstructing the direct sunlight causes a deep shadow on part of the document. This problem does not usually appear in an indoor setting, and certainly never on an image scanned in a constrained environment. Undetected or unrepaired shadows result in unusable images, increasing the number of rejected images. With advanced mobile imaging techniques, shadows can not only be detected, but often eliminated, preventing the need to ask the user to take the photo again

[0371] According to some embodiments, a Shadow Test can be executed on a mobile image to determine whether a portion of the image is covered by a shadow. A shadow can render parts of a mobile image unreadable. This test helps to identify whether a shadow coverage a least a portion of a subimage in a mobile document image, and to reject images if the shadow has too much of an effect on the image quality, so that the user can attempt to take a better quality image of the document where the shadow is not present.

[0372] According to an embodiment, the presence of a shadow is measured by examining boundaries in the mobile image that intersect two or more sides of the document sub-

image. FIG. 30 illustrates an example of a shadowed document. The document subimage has been extracted from the mobile document image and converted to a grayscale snippet in this example. The shadow boundary clearly intersects the top and the bottom of the check pictured in the snippet.

[0373] The presence of shadows can be measured using the area and contrast. If a shadow covers the entire image, the result is merely an image that is darker overall. Such shadows generally do not worsen image quality significantly. Furthermore, shadows having a very small surface area also do not generally worsen image quality very much.

[0374] According to an embodiment, the Image Shadowed Score can be calculated using the following formula to determine the score for a grayscale snippet:

Image Shadowed score=1000 if no shadows were found, otherwise

Image Shadowed score=1000-min (Score($S[i]$)),
where Score($S[i]$) is computed for every shadow $S[i]$ detected on the grayscale snippet

[0375] In an embodiment, the Score for each shadow can be computed using the following formula:

Given shadow $S[i]$ in the grayscale image, the score can be calculated Score($S[i]$) as Score($S[i]$) = $2000 * \min(A[i]/A, 1 - A[i]/A) * (\text{Contrast}/256)$, where $A[i]$ is the area covered by shadow $S[i]$ (in pixels), A is the entire grayscale snippet area (in pixels), and Contrast is the difference of brightness inside and outside of the shadow (the maximum value is 256).

[0376] Due to the normalization factor 2000, Score($S[i]$) fits into 0-1000 range. It tends to assume larger values for shadows that occupy about 1/2 of the snippet area and have high contrast. Score($S[i]$) is typically within 100-200 range. In an embodiment, the Image Shadowed score calculated by this test falls within a range of 0-1000 as do the test results from other tests. According to an embodiment, a typical mobile document image with few shadows will have a test result value in a range from 800-900. If no shadows are on or are found the document subimage, then the score will equal 1000. The Image Shadowed score can then be compared to a threshold associated with the test to determine whether the image is of sufficiently high quality for use with the mobile application requesting the assessment of the quality of the mobile document image.

Contrast Test

[0377] According to some embodiments, a Contrast Test can be executed on a mobile image to determine whether the contrast of the image is sufficient for processing. One cause of poor contrast is images taken with insufficient light. A resulting grayscale snippet generated from the mobile document image can have low contrast, and if the grayscale snippet is converted to a binary image, the binarization unit can erroneously white-out part of the foreground, such as the MICR-line of a check, the code line of a remittance coupon, an amount, or black-out part of the background. The Contrast Test measures the contrast and rejects poor quality images, and instructs the user to retake the picture under brighter light to improve the contrast of the resulting snippets.

[0378] FIG. 32 illustrates a method for executing a Contrast IQA Test according to an embodiment. The Contrast IQA Test illustrated in FIG. 32 is performed on a grayscale snippet generated from a mobile document image. The MDIPE 2100

receives the mobile image (step 2805) and generates a grayscale snippet that comprises a grayscale version of the document subimage (step 2810). FIG. 31 is an example of a grayscale snippet generated from a mobile document image of a check. As can be seen from FIG. 27, the contrast of the image is very low.

[0379] A histogram of the grayscale values in the grayscale snippet can then be built (step 2815). In an embodiment, the x-axis of the histogram is divided into bins that each represents a "color" value for the pixel in the grayscale image and the y-axis of the histogram represents the frequency of that color value in the grayscale image. According to an embodiment, the grayscale image has pixel in a range from 0-255, and the histogram is built by iterating through each value in this range and counting the number of pixels in the grayscale image having this value. For example, frequency of the "200" bin would include pixels having a gray value of 200.

[0380] A median black value can then be determined for the grayscale snippet (step 2820) and a median white value is also determined for the grayscale snippet (step 2825). The median black and white values can be determined using the histogram that was built from the grayscale snippet. According to an embodiment, the median black value can be determined by iterating through each bin, starting with the "255" bin that represents pure black and moving progressively toward the "250" bin which represents pure white. Once a bin is found that includes at least 20% of the pixels included in the image, the median black value is set to be the grayscale value associated with that bin. According to an embodiment, the median white value can be determined by iterating through each bin, starting with the "255" bin which represents pure white and moving progressively toward the "0" bin which represents pure black. Once a bin is found that includes at least 20% of the pixels included in the image, the median white value is set to be the color value associated with that bin.

[0381] Once the median black and white values have been determined, the difference between the median black and white values can then be calculated (step 2830). The difference can then be normalized to fall within the 0-1000 test range used in the mobile IQA tests executed by the MDIPE 2100 (step 2835). The test result value can then be returned (step 2840). As described above, the test result value is provided to the test execution unit 2130 where the test result value can be compared to a threshold value associated with the test. See for example, FIG. 26, step 2230, described above. If the mobile image fails the Contrast IQA Test, the MDIPE 2100 can reject the image, and load detailed test messages from the test message data store 134 that include detailed instructions that how the user might retake the image.

Planar Skew Test

[0382] According to some embodiments, a Planar Skew Test can be executed on a mobile image to determine whether the document subimage is skewed within the mobile image. See FIG. 33A for an example of a mobile document image that includes a remittance coupon or check that exhibits significant planar skew. Planar skew does not result in distortion of the document subimage; however, in an embodiment, the subimage detection unit included in the preprocessing unit assumes that the document subimage is nearly horizontal in the mobile document image. If the skew becomes too extreme, for example approaching 45 degrees from horizontal, cropping errors could occur when the document subimage is extracted from the mobile document image.

[0383] According to an embodiment, document skew can be measured by first identifying the corners of the document subimage using one of the techniques described above. The corners of the documents subimage can be identified by the preprocessing unit **130** when performing projective transformations on the subimage, such as that described above with respect to FIGS. **28A** and **28B**. Various techniques for detecting the skew of the subimage can be used. For example, techniques for detecting skew disclosed in the related '071 and '091 Applications, can be used to detect the skew of the subimage. The results from the skew test can then be fall within the 0-1000 test range used in the mobile IQA tests executed by the MDIPE **2100**. The higher the skew of the document subimage, the lower the normalized test value. If the normalized test value falls below the threshold value associated with the test, the mobile document image can be rejected and the user can be provided detailed information from the test result messages data store **136** for how to retake the image and reduce the skew.

View Skew Test

[0384] "View skew" denotes a deviation from direction perpendicular to the document in mobile document image. Unlike planar skew, the view skew can result in the document subimage having perspective distortion. FIG. **33B** illustrates an example of a document subimage that exhibits view skew. View skew can cause problems in processing the subimage if the view skew becomes too great, because view skew changes the width-to-height ratio of the subimage. This can present a problem, since the true dimensions of the document pictured in the subimage are often unknown. For example, remittance coupons and business checks can be various sizes and can have different width-to-height ratios. View skew can result in content recognition errors, such as errors in recognition of the MICR-line data on a check or CAR/LAR recognition (which stands for Courtesy Amount Recognition and Legal Amount Recognition) or errors in recognition of the code line of a remittance coupon. By measuring the view skew, the view skew test can be used to reject images that have too much view skew, which can help reduce false rejects and false accepts rates by addressing an issue that can be easily corrected by a user retaking the mobile document image.

[0385] FIG. **34** is a flow chart illustrating a method for testing for view skew according to an embodiment. The MDIPE **2100** receives the mobile image (step **3005**) and identifies the corners of the document within the subimage (step **3010**). A skew test score can then be determined for the document subimage (step **3015**) and skew test score can then be returned (**3040**). As described above, the test result value can then be provided to the test execution unit **2130** where the test result value can be compared to a threshold value associated with the test.

[0386] According to an embodiment, the view skew of a mobile document can be determined using the following formula:

View Skew score=1000-F(A, B, C, D), where

$$F(A, B, C, D)=500*\max(\frac{\text{abs}(|AB|-|CD|)}{(|DA|+|BC|)}, \frac{\text{abs}(|BC|-|DA|)}{(|AB|+|CD|)}),$$

[0387] where |PQ| denotes the distance from point P to point Q, and the corners of the subimage are denoted as follows: A represents the top-left corner, B represents the top-right corner of the subimage, C represents the

bottom-right corner of the subimage, and D represents the bottom-left corner of the subimage.

[0388] One can see that View Skew score can be configured to fit into [0, 1000] range used in the other mobile IQA tests described herein. In this example, the View Skew score is equal to 1000 when |AB|=|CD| and |BC|=|DA|, which is the case when there is no perspective distortion in the mobile document image and camera-to-document direction was exactly perpendicular. The View Skew score can then be compared to a threshold value associated with the test to determine whether the image quality is sufficiently high for use with the mobile application.

Cut Corner Test

[0389] Depending upon how carefully the user framed a document when capturing a mobile image, it is possible that one or more corners of the document can be cut off in the mobile document image. As a result, important information can be lost from the document. For example, if the lower left-hand corner of a check is cut off in the mobile image, a portion of the MICR-line of a check or the code line of a remittance coupon might be cut off, resulting in incomplete data recognition. FIG. **35** illustrates an example of a mobile document image that features a receipt where one of the corners has been cut off.

[0390] FIG. **36** illustrates a Cut-Off Corner Test that can be used with embodiments of the MDIPE **2100** for testing whether corners of a document in a document subimage have been cut off when the document was imaged. The mobile image including height and width parameters are received (step **3205**). In an embodiment, the height and width of the mobile image can be determined by the preprocessing unit **2110**. The corners of the document subimage are then identified in the mobile document image (step **3210**). Various techniques can be used to identify the corners of the image, including the various techniques described above. In an embodiment, the preprocessing unit **2110** identifies the corners of the document subimage. As illustrated in FIG. **15**, one or more of the corners of a document can be cut off. However, the preprocessing unit **2110** can be configured to determine what the location of the corner should have been had the document not been cut off using the edges of the document in the subimage. FIG. **35** illustrates how the preprocessing unit **2110** has estimated the location of the missing corner of the document by extending lines from the sides of the document out to the point where the lines intersect. The preprocessing unit **2110** can then provide the corners information for the document to the test execution unit **2130** to execute the Cut-Off Corner IQA Test. In an embodiment, test variables and the test results values to be returned by the test are set to default values: the test value V to be returned from the test is set to a default value of 1000, indicating that all of the corners of the document are within the mobile document image, and a maximum cut off variable (MaxCutOff) is set to zero indicating that no corner was cut off.

[0391] A corner of the document is selected (step **3220**). In an embodiment, the four corners are received as an array of x and y coordinates C[I], where I is equal to the values 1-4 representing the four corners of the document.

[0392] A determination is made whether the selected corner of the document is within the mobile document image (step **3225**). The x & y coordinates of the selected corner should be at or between the edges of the image. According to an embodiment, the determination whether a corner is within

the mobile document image can be determined using the following criteria: (1) $C[I].x > 0$ & $C[I].x \leq \text{Width}$, where Width = the width of the mobile document image and $C[I].x$ = the x-coordinate of the selected corner; and (2) $C[I].y > 0$ & $C[I].y \leq \text{Height}$, where Height = the height of the mobile document image and $C[I].y$ = the y-coordinate of the selected corner.

[0393] If the selected corner fails to satisfy the criteria above, the corner is not within the mobile image and has been cut-off. A corner cut-off measurement is determined for the corner (step **3230**). The corner cut-off measurement represents the relative distance to the edge of the mobile document image. According to an embodiment, the corner cut-off measurement can be determined using the following:

[0394] (1) Set $H[I]$ and $V[I]$ to zero, where $H[I]$ represents the horizontal normalized cut-off measure and $V[I]$ represents the vertical normalized cut-off measure.

[0395] (2) If $C[I].x < 0$, then set $H[I] = -1000 * C[I].x / \text{Width}$

[0396] (3) If $C[I].x > \text{Width}$, set $H[I] = 1000 * (C[I].x - \text{Width}) / \text{Width}$, where Width is the width of the mobile image

[0397] (4) If $C[I].y < 0$, set $V[I] = -1000 * C[I].y / \text{Height}$, where Height is the height of the mobile image

[0398] (5) If $C[I].y > \text{Height}$, set $V[I] = 1000 * (C[I].y - \text{Height}) / \text{Height}$

[0399] (6) Normalize $H[I]$ and $V[I]$ to fall within the 0-1000 range used by the mobile IQA tests by setting $H[I] = \min(1000, H[I])$ and $V[I] = \min(1000, V[I])$.

[0400] (7) Set $\text{CutOff}[I] = \min(H[I], V[I])$, which is the normalized cut-off measure of the corner. One can see that the $\text{CutOff}[I]$ lies within [0-1000] range used by the mobile IQA tests and the value increases as the corner moves away from mobile image boundaries.

[0401] An overall maximum cut-off value is also updated using the normalized cut-off measure of the corner (step **3235**). According to an embodiment, the following formula can be used to update the maximum cut-off value: $\text{MaxCutOff} = \max(\text{MaxCutOff}, \text{CutOff}[I])$. Once the maximum cut-off value is determined, a determination is made whether more corners are to be tested (step **3225**).

[0402] If the selected corner satisfies the criteria above, the corner is within the mobile document image and is not cut-off. A determination is then made whether there are additional corners to be tested (step **3225**). If there are more corners to be processed, a next corner to be test is selected (step **3215**). Otherwise, if there are no more corners to be tested, the test result value for the test is computing using the maximum test cut-off measurement. In an embodiment, the test result value $V = 1000 - \text{MaxCutOff}$. One can see that V lies within [0-1000] range for the mobile IQA tests and is equal to 1000 when all the corners are inside the mobile image and decreases as one or more corner move outside of the mobile image.

[0403] The test result value is then returned (**3245**). As described above, the test result value is provided to the test execution unit **2130** where the test result value can be compared to a threshold value associated with the test. If the test result value falls below the threshold associated with the test, detailed test result messages can be retrieved from the test result message data store **136** and provided to the user to indicate why the test failed and what might be done to remedy the test. The user may simply need to retake the image with the document corners within the frame.

Cut-Side Test

[0404] Depending upon how carefully the user framed a document when capturing a mobile image, it is possible that one or more sides of the document can be cut off in the mobile document image. As a result, important information can be lost from the document. For example, if the bottom a check is cut off in the mobile image, the MICR-line might be cut off, rendering the image unusable for a Mobile Deposit application that uses the MICR information to electronically deposit checks. Furthermore, if the bottom of a remittance coupon is cut off in the mobile image, the code line may be missing, the image may be rendered unusable by a Remittance Processing application that uses the code information to electronically process the remittance.

[0405] FIG. 37 illustrates an example of a mobile document image that features a receipt where one of the ends of the receipt has been cut off in the image. Unlike the Cut-Corner Test described above which can be configured to allow a document to pass if the amount of cut-off falls is small enough that the document image still receives a test score that meets or exceeds the threshold associated with the test, the Cut-Side Test is either pass or fail. If one or more sides of the document subimage are cut off in the mobile document image, the potential to lose critical information is too high, and mobile document is marked as failing.

[0406] FIG. 38 is a flow diagram of a method for determining whether one or more sides of the document are cut off in the document subimage according to an embodiment. The mobile image is received (step **3405**). In an embodiment, the height and width of the mobile image can be determined by the preprocessing unit **2110**. The corners of the document subimage are then identified in the mobile document image (step **3410**). Various techniques can be used to identify the corners of the image, including the various techniques described above. In an embodiment, the preprocessing unit **2110** identifies the corners of the document subimage.

[0407] A side of the document is selected (step **3420**). In an embodiment, the four corners are received as an array of x and y coordinates $C[I]$, where I is equal to the values 1-4 representing the four corners of the document.

[0408] A determination is made whether the selected corner of the document is within the mobile document image (step **3425**). According to an embodiment, the document subimage has four side and each side $S[I]$ includes two adjacent corners $C1[I]$ and $C2[I]$. A side is deemed to be cut-off if the corners comprising the side are on the edge of the mobile image. In an embodiment, a side of the document is cut-off if any of the following criteria are met:

[0409] (1) $C1[I].x = C2[I].x = 0$, where x = the x-coordinate of the corner

[0410] (2) $C1[I].x = C2[I].x = \text{Width}$, where Width = the width of the mobile image

[0411] (3) $C1[I].y = C2[I].y = 0$, where y = the y-coordinate of the corner

[0412] (4) $C1[I].y = C2[I].y = \text{Height}$, where Height = the height of the mobile image

[0413] If the side does not fall within the mobile image, the test result value is set to zero indicating that the mobile image failed the test (step **3430**), and the test results are returned (step **3445**).

[0414] If the side falls within the mobile image, a determination is made whether there are more sides to be tested (step **3425**). If there are more sides to be tested, an untested side is selected (step **3415**). Otherwise, all of the sides were within

the mobile image, so the test result value for the test is set to 1000 indicating the test passed (step 3440), and the test result value is returned (step 3445).

Warped Image Test

[0415] In real life, paper documents are often warped (folded) in various, irregular ways due to long and/or careless handling. Traditional scanners deal with this situation by physically smoothing out the paper during scanning by pressing it between two flat surfaces. However, this is not the case with a mobile photo of a warped paper document. Failure to de-warp results in an unreadable document. Without advanced de-warping techniques, a large number of all document images will be rejected by the bank's processing system (or flagged for manual processing), since the information on them cannot be extracted automatically. This leads to a large proportion of rejected or failed payments and increased labor costs, frustrated users and damage to the bank's reputation and business

[0416] The warped image test identifies images where document is warped. FIG. 39 illustrates an example of a mobile document image where the document is warped. In some embodiments, the preprocessing unit 2110 can be configured to include de-warping functionality for correcting warped images. However, in some embodiments, a Warped Image Test is provided to detect and reject warped images. One solution for correcting warped images is to instruct the user to retake the image after flattening the hardcopy of the document being imaged.

[0417] FIG. 40 is a flow diagram of a method for identifying a warped image and for scoring the image based on how badly the document subimage is warped according to an embodiment. A warped image test score value is returned by the test, and this value can be compared with a threshold value by the test execution unit 2130 to determine whether the image warping is excessive.

[0418] The mobile image is received (step 3605). In an embodiment, the height and width of the mobile image can be determined by the preprocessing unit 2110. The corners of the document subimage are then identified in the mobile document image (step 3610). Various techniques can be used to identify the corners of the image, including the various techniques described above. In an embodiment, the preprocessing unit 2110 identifies the corners of the document subimage.

[0419] A side of the document is selected (step 3615). According to an embodiment, the document subimage has four side and each side S[I] includes two adjacent corners C1[I] and C2[I].

[0420] A piecewise linear approximation is built for the selected side (step 3620). According to an embodiment, the piecewise-linear approximation is built along the selected side by following the straight line connecting the adjacent corners C1[I] and C2[I] and detecting position of the highest contrast starting from any position within [C1[I], C2[I]] segment and moving in orthogonal direction.

[0421] After the piecewise linear approximation is built along the [C1[I], C2[I]] segment, the [C1[I], C2[I]] segment is walked to compute the deviation between the straight line and the approximation determined using piecewise linear approximation (step 3625). Each time the deviation is calculated, a maximum deviation value (MaxDev) is updated to reflect the maximum deviation value identified during the walk along the [C1[I], C2[I]] segment.

[0422] The maximum deviation value for the side is then normalized to generate a normalized maximized deviation value for the selected size of the document image (step 3630). According to an embodiment, the normalized value can be determined using the following formula:

$$\text{NormMaxDev}[I] = 1000 * \text{MaxDev}[I] / \text{Dim}, \text{ where Dim} \\ \text{is the mobile image dimension perpendicular to} \\ \text{side SM.}$$

[0423] An overall normalized maximum deviation value is then updated using the normalized deviation value calculated for the side. According to an embodiment, the overall maximum deviation can be determined using the formula:

$$\text{OverallMaxDeviation} = \max(\text{OverallMaxDeviation}, \\ \text{NormMaxDev}[I])$$

[0424] A determination is then made whether there are anymore sides to be tested (step 3640). If there are more sides to be tested, an untested side is selected for testing (step 3615). Otherwise, if no untested sides remain, the warped image test value is computed. According to an embodiment, the warped image test value can be determined using the following formula:

$$V = 1000 - \text{OverallMaxDeviation}$$

[0425] One can see that V lies within [0-1000] range used by the image IQA system and is equal to 1000 when the sides S[I] are straight line segments (and therefore no warp is present). The computed test result is then returned (step 3650). As described above, the test result value is provided to the test execution unit 2130 where the test result value can be compared to a threshold value associated with the test. If the test result value falls below the threshold associated with the test, detailed test result messages can be retrieved from the test result message data store 136 and provided to the user to indicate why the test failed and what might be done to remedy the test. For example, the user may simply need to retake the image after flattening out the hardcopy of the document being imaged in order to reduce warping.

Image Size Test

[0426] The Image Size Test detects the actual size and the effective resolution of the document subimage. The perspective transformation that can be performed by embodiments of the preprocessing unit 2110 allows for a quadrangle of any size to be transformed into a rectangle to correct for view distortion. However, a small subimage can cause loss of detail needed to process the subimage.

[0427] FIG. 41 illustrates an example of a document subimage within a mobile document image that is relatively small. Small size of the subimage can cause the loss of important foreground information. This effect is similar to digital zooming in a digital camera where image of an object becomes larger, but the image quality of object can significantly degrade due to loss of resolution and important details can be lost.

[0428] FIG. 42 is a flow diagram of a process that for performing an Image Size Test on a subimage according to an embodiment. The mobile image is received (step 3805). In an embodiment, the height and width of the mobile image can be determined by the preprocessing unit 2110. The corners of the document subimage are then identified in the mobile document image (step 3810). Various techniques can be used to identify the corners of the image, including the various techniques described above. In an embodiment, the preprocessing

unit **2110** identifies the corners of the document subimage. In the method the corners of the subimage are denoted as follows: A represents the top-left corner, B represents the top-right corner of the subimage, C represents the bottom-right corner of the subimage, and D represents the bottom-left corner of the subimage.

[0429] A subimage average width is computed (step **3815**). In an embodiment, the subimage average width can be calculated using the following formula:

$$\text{Subimage average width as AveWidth} = (|AB| + |CD|) / 2,$$

where |PQ| represents the Euclidian distance from point P to point Q.

[0430] A subimage average height is computed (step **3820**). In an embodiment, the subimage average height can be calculated using the following formula:

$$\text{AveHeight} = (|BC| + |DA|) / 2$$

[0431] The average width and average height values are then normalized to fit the 0-1000 range used by the mobile IQA tests (step **3822**). The following formulas can be used determine the normalized average width and height:

$$\text{NormAveWidth} = 1000 * \text{AveWidth} / \text{Width}$$

$$\text{NormAveHeight} = 1000 * \text{AveHeight} / \text{Height}$$

[0432] A minimum average value is then determined for the subimage (step **3825**). According to an embodiment, the minimum average value is the smaller of the normalized average width and the normalized average height values. The minimum average value falls within the 0-1000 range used by the mobile IQA tests. The minimum average value will equal 1000 if the document subimage fills the entire mobile image.

[0433] The minimum average value is returned as the test result (step **3865**). As described above, the test result value is provided to the test execution unit **2130** where the test result value can be compared to a threshold value associated with the test. If the test result value falls below the threshold associated with the test, detailed test result messages can be retrieved from the test result message data store **2136** and provided to the user to indicate why the test failed and what might be done to remedy the test. For example, the user may simply need to retake the image by positioning the camera closer to the document.

Code Line Test

[0434] The Code Line Test can be used to determine whether a high quality image of a remittance coupon front has been captured using the mobile device according to an embodiment. The Code Line Test can be used in conjunction with a Remittance Processing application to ensure that images of remittance coupon captures for processing with the Remittance Processing information are of a high enough quality to be processed so that the remittance can be electronically processed. Furthermore, if a mobile image fails the Code Line Test, the failure may be indicative of incorrect subimage detections and/or poor overall quality of the mobile image, and such an image should be rejected anyway.

[0435] FIG. 43 is a flow chart of a method for executing a Code Line Test according to an embodiment. A mobile image of a remittance coupon is received (step **3955**) and a bitonal image is generated from the mobile image (step **3960**). In an embodiment, preprocessor **110** extracts the document subimage from the mobile image as described above, including preprocessing such as geometric correction. The extracted

subimage can then be converted to a bitonal snippet by the preprocessor **110**. The code line is then identified in the bitonal snippet (step **3965**). According to an embodiment, a code line recognition engine is then applied to identify the code line and to compute character-level and overall confidence values for the image (step **3970**). These confidences can then be normalized to the 0-1000 scale used by the mobile IQA tests where 1000 means high quality and 0 means poor code line quality. The confidence level is then returned (step **3975**). As described above, the test result value is provided to the test execution unit **2130** where the test result value can be compared to a threshold value associated with the test. If the test result value falls below the threshold associated with the test, detailed test result messages can be retrieved from the test result message data store **136** and provided to the user to indicate why the test failed and what might be done to remedy the test. For example, the user may simply need to retake the image to adjust for geometrical or other factors, such as poor lighting or a shadowed document. In some instances, the user may not be able to correct the errors. For example, if the code line on the document is damaged or incomplete and the document will continue to fail the test even if the image were retaken.

Aspect Ratio Tests

[0436] The width of a remittance coupon is typically significantly longer than the height of the document. According to an embodiment, an aspect ratio test can be performed on a document subimage of a remittance coupon to determine whether the aspect ratio of the document in the image falls within a predetermined ranges of ratios of width to height. If the document image falls within the predetermined ranges of ratios, the image passes the test. An overall confidence value can be assigned to different ratio values or ranges of ratio values in order to determine whether the image should be rejected.

[0437] According to some embodiments, the mobile device can be used to capture an image of a check in addition to the remittance coupon. A second aspect ratio test is provided for two-sided documents, such as checks, where images of both sides of the document may be captured. According to some embodiments, a remittance coupon can also be a two-sided document and images of both sides of the document can be captured. The second aspect ratio test compares the aspect ratios of images that are purported to be of the front and back of a document to determine whether the user has captured images of the front and back of the same document according to an embodiment. The Aspect Ratio Test could be applied to various types two-sided or multi-page documents to determine whether images purported to be of different pages of the document have the same aspect ratio.

[0438] FIG. 44 illustrates a method for executing an Aspect Ratio Test for two-sided documents according to an embodiment. In the embodiment illustrated in FIG. 40, the test is directed to determining whether the images purported to be of the front and back side of a document have the same aspect ratio. However, the method could also be used to test whether two images purported to be from a multi-page and/or multi-sided document have the same aspect ratio.

[0439] A front mobile image is received (step **4005**) and a rear mobile image is received (step **4010**). The front mobile image is supposed to be of the front side of a document while the rear mobile image is supposed to be the back side of a document. If the images are really of opposite sides of the

same document, the aspect ratio of the document subimages should match. Alternatively, images of two different pages of the same document may be provided for testing. If the images are really of pages of the same document, the aspect ratio of the document subimages should match.

[0440] The preprocessing unit **2110** can process the front mobile image to generate a front-side snippet (step **4015**) and can also process the back side image to generate a back-side snippet (step **4020**).

[0441] The aspect ratio of the front-side snippet is then calculated (step **4025**). In an embodiment, the $\text{AspectRatioFront} = \text{Width}/\text{Height}$, where Width = the width of the front-side snippet and Height = the height of the front-side snippet.

[0442] The aspect ratio of the back-side snippet is then calculated (step **4030**). In an embodiment, the $\text{AspectRatioBack} = \text{Width}/\text{Height}$, where Width = the width of the back-side snippet and Height = the height of the back-side snippet.

[0443] The relative difference between the aspect ratios of the front and rear snippets is then determined (step **4035**). According to an embodiment, the relative difference between the aspect ratios can be determined using the following formula:

$$\text{RelDiff} = 1000 * \text{abs}(\text{AspectRatioFront} - \text{AspectRatioBack}) / \max(\text{AspectRatioFront}, \text{AspectRatioBack})$$

[0444] A test result value is then calculated based on the relative difference between the aspect ratios (step **4040**). According to an embodiment, the test value V can be computed using the formula $V = 1000 - \text{RelDiff}$.

[0445] The test results are then returned (step **4045**). As described above, the test result value is provided to the test execution unit **2130** where the test result value can be compared to a threshold value associated with the test. If the test result value falls below the threshold associated with the test, detailed test result messages can be retrieved from the test result message data store **136** and provided to the user to indicate why the test failed and what might be done to remedy the test. For example, the user may have mixed up the front and back images from two different checks having two different aspect ratios. If the document image fails the test, the user can be prompted to verify that the images purported to be the front and back of the same document (or images of pages from the same document) really are from the same document.

Form Identification

[0446] Various embodiments of the present invention may utilize a novel technique of form identification in order to expeditiously identify key features of a captured mobile image. The form identification can be provided by a user, or it can be automatically determined by reading a captured mobile image. This captured mobile image may include any type of document including, without limitation: remittance coupons, employment forms, store receipts, checks, bills or sales invoices, business cards, medical and dental records, store coupons, educational information such as progress reports and report cards, birth and death certificates, insurance policies, legal documents, magazine and newspaper clippings, forms of personal identification such as passports and driver licenses, police records, real estate records, etc. In the form identification step, a template is identified that is associated with a document that has been captured in a mobile image. The template identifies the layout of information con-

tained within the document. This layout information can be used to improve data capture accuracy because data should be in known locations on the document.

[0447] Form identification can be helpful in a number of different situations. If the layout of the document is known, capturing the data from known locations on the document can be more accurate than relying on a dynamic data capture technique to extract the data from the document. Additionally, according to some embodiments, the identification of a pre-requisite minimum number of data fields associated with only one type of document can enable a faster lookup of data from other data fields as soon as the specific type of document has been identified.

[0448] Form identification can also be used for documents that lack keywords that could otherwise be used to identify key data on the document. For example, if a document does not include an "Account Number" label for an account number field, the dynamic data capture may misidentify the data in that field. Misidentification can become even more likely if multiple fields have similar formats. Form identification can also be used for documents having ambiguous data. For example, a document might include multiple fields that include data having a similar format. If a document includes multiple unlabeled fields having similar formats, dynamic data capture may be more likely to misidentify the data. However, if the layout of the document is known, the template information can be used to extract data from known positions in the document image.

[0449] According to some embodiments, form identification can also be used for documents having a non-OCR friendly layout. For example, a document may use fonts where identifying keywords and/or form data is printed using a non-OCR friendly font. Form identification can also be used to improve the chance of correctly capturing data when a poor quality image is presented. A poor quality image of a document can make it difficult to locate and/or read data.

[0450] FIG. 45 is a flow chart of a method for processing an image using form identification according to an embodiment. At step **4205**, a binarized/bi-tonal document image is received. Various techniques for creating a bi-tonal subimage from a mobile image are provided above. For example, step **1225** of FIG. 12 describes binarization of a document subimage. FIG. 14 also illustrates a method of binarization that can be used to generate a bi-tonal image according to one embodiment of the present invention.

[0451] A matching algorithm is executed on the bi-tonal image of the document in an attempt to find a matching template (step **4210**). According to an embodiment, one or more computing devices can include a template data store that can be used to store templates of the layouts of various types of documents. Various matching techniques can be used to match a template to a document image. For example, optical character recognition can be used to identify and read text content from the image. The types of data identified and the positions of the data on the document can be used to identify a matching template. According to another embodiment, a document can include a unique symbol or identifier that can be matched to a particular document template. In yet other embodiments, the image of the document can be processed to identify "landmarks" on the image that may correspond to labels and/or data. In some embodiments, these landmarks can include, but are not limited to: positions of horizontal and/or vertical lines on the document, the position and/or size of boxes and/or frames on the document, and/or the location

of pre-printed text. The position of these landmarks on the document may be used to identify a template from the plurality of templates in the template data store. According to some embodiments, a cross-correlation matching technique can be used to match a template to an image of a document. In some embodiments, the positions of frames/boxes found on image and/or other such landmarks, can be cross-correlated with landmark information associated a template to compute the matching confidence score. If the confidence score exceeds a predetermined threshold, the template is considered to be a match and can be selected for use in extracting information from the mobile image.

[0452] A determination is made whether a matching template has been found (step 4215). If no matching template is found, a dynamic data capture can be performed on the image of the document (step 4225). Dynamic data capture is described in detail below and an example method for dynamic data capture is illustrated in the flow chart of FIG. 46.

[0453] If a matching template is found, data can be extracted from the image of the document using the template (step 4220). The template can provide the location of various data within the document, such as the document's author(s), the document's publication date, the names of any corporate, governmental, or educational entities associated with the document, an amount due, an account holder name, an account number, a payment due date, etc. In some embodiments, various OCR techniques can be used to read text content from the locations specified by the template. Since the location of various data elements is known, ambiguities regarding the type of data found can be eliminated. That is, use of the template enables the system to distinguish among data elements which have a similar data type.

Dynamic Data Capture

[0454] FIG. 46 is a flow chart of a dynamic data capture method for extracting data from an image according to an embodiment. The dynamic data capture method illustrated in FIG. 46 can be used if a form ID for identifying a particular format of a document is not available. The method illustrated in FIG. 46 can also be used if the form ID does not match any of the templates stored in the template data store. The method begins with receiving a binarized/bi-tonal document image (step 4305). Various optical character recognition techniques can then be used to locate and read fields from the bi-tonal image (step 4310). Some example OCR techniques are described below. Once data fields have been located, the data can be extracted from the bi-tonal image (step 4315). In some embodiments, steps 4310 and 4315 can be combined into a single step where the field data is located and the data extracted in a combined OCR step. Once the data has been extracted from the image, the data can be analyzed to identify what data has been extracted (step 4320). The data can also be analyzed to determine whether any additional data is required in order to be able to process the image.

[0455] According to an embodiment, a keyword-based detection technique can be used to locate and read the data from the bi-tonal image in steps 4310 and 4315 of the method of FIG. 46. The method uses a set of field-specific keywords to locate fields of interest in the bitonal image. For example, if the captured image is an image of a remittance coupon, the keywords "Account Number," "Account #," "Account No.," "Customer Number," and/or other variations can be used to identify the customer's account number. According to an embodiment, text located proximate to the keyword can be

associated with the keyword. For example, text located within a predetermined distance to the right of or below an "Account Number" keyword may be identified and extracted from the image using OCR and the text found in this location can then be treated as the account number. According to an embodiment, the distance and directions in relation to the keyword in which the field data can be located can be configured based on the various parameters, such as locale or language. The position of the keyword in relation to field that includes the data associated with the keyword may vary based on the language being used, e.g. written right to left versus left to right.

[0456] According to an embodiment, a format-based detection technique can be used to locate and read the data from the bi-tonal image in steps 4310 and 4315. For example, an OCR technique can be used to recognize text in the document image. A regular expression mechanism can then be applied to the text extracted from the bitonal image. A regular expression can be used to formalize the format description for a particular field, such as "contains 7-12 digits," "may start with 1 or 2 uppercase letters," or "contains the letter 'U' in the second position." According to an embodiment, multiple regular expressions may be associated with a particular field, such as an account number, in order to increase the likelihood of a correct match.

[0457] According to yet another embodiment, a combination of keyword-based and format-based matching can be used to identify and extract field data from the bi-tonal image (steps 4310 and 4315). This approach can be particularly effective where multiple fields of the same or similar format are included within the image. A combination of keyword-based and format-based matching can be used to identify field data can be used to disambiguate the data extracted from the bi-tonal image.

[0458] According to an embodiment, a code-line validation technique can be used to locate and read the data from the bi-tonal image of in steps 4310 and 4315. One or more fields may be embedded into a code-line. In some embodiments, the code-line characters may be cross-checked against fields recognized in other parts of the document. In the event that a particular field is different from a known corresponding value in the code line, the value in the code line may be selected over the field value due to the relative difference in the reliabilities of reading the code line versus reading the field value.

[0459] According to an embodiment, a cross-validation technique can be used where multiple bi-tonal images of the same document have been captured, and one or more OCR techniques are applied the each of the bi-tonal images (such as by any of the techniques described above). The results from the one or more OCR technique from one bi-tonal image can be compared to the results of OCR techniques applied one or more other bitonal images in order to cross-validate the field data extracted from the images. If conflicting results are found, a set of results having a higher confidence value can be selected to be used for document image processing.

Recurring Payment Scheduling

[0460] According to various embodiments, a user of the mobile device application can set up one or more recurring payment schedules. A recurring payment schedule may have a variety of advantages over a series of single payments, including: i.) utilizing persistent data in order to make the process of paying a bill more expeditious for the user (i.e., less input may be required from the user before each bill is submitted), ii.) enabling a fast lookup of a remittance coupon

template associated with a specified payee (thereby decreasing search time); and iii.) enabling the remittance application to send one or more payment reminders to the user so as to safeguard against a payment default.

[0461] FIG. 47 is a flow diagram illustrating an exemplary method for configuring a recurring bill payment schedule according to one embodiment. At block 4702, a user launches a remittance application. In some embodiments, the remittance application is resident within the mobile device (see FIG. 1). In other embodiments, the remittance application is resident within a remote computing device, such as a remote server (see FIG. 1). Once the remittance application is launched, a splash screen may appear (block 4704) indicating the name and/or software version of the remittance application.

[0462] At block 4706, a login screen can then be displayed, prompting the user to input one or more security credentials (e.g., username and a password). In some embodiments, the security credentials of all users of the remittance application may be encrypted and stored locally, for example, within a non-volatile storage device associated with the mobile device 350. In other embodiments, the security credentials may be encrypted and stored in a non-volatile device present at a remote location.

[0463] Once the credentials have been validated, a main menu is then displayed (block 4708). The main menu may list a number of functions associated with the remittance application, including the option to “pay a bill” or to “view the last bill paid.” An option to “configure recurring payments” is also presented to the user as one of the options, and the application will listen for the user’s selection of this option at decision block 4710.

[0464] At block 4712, a listing of all recurring payment schedules associated with the user is then displayed. For example, if the user had previously set up a recurring payment with Time Warner Cable and San Diego Gas and Electric, these two entries will be selectable within this listing. However, if no entries had been previously entered and saved by the user, a message such as: “No recurring payments have been scheduled” may appear in the display window in the alternative. An additional option to “set up a new recurring payment” is also presented to the user, for example, at the bottom of the display screen.

[0465] At blocks 4714 and 4716, the user will decide whether to update an existing recurring bill payment or to otherwise set up a new recurring payment. In the event that the user selected a preexisting recurring payment entry, previously stored data regarding this entry will be loaded at block 4718 (such as the name of the recurring payment entry, the payor, the payee, the selected payment method, a bank account or check routing number, a credit card number, and any other preferred payment options). Otherwise, in the event that the user had selected to set up a new recurring payment, these data fields may be blank by default.

[0466] At block 4720, a sub-menu is then displayed including various data fields associated with this recurring payment entry. In some embodiments, the user may have an option to auto-populate at least some of these fields by instructing the system to extract data from a bill that has already been paid. Other fields can be modified, for example, by a keyboard, touchpad, mouse, or other such input device.

[0467] At block 4722, the user may then update these fields accordingly. In some embodiments, a “save” or “apply changes” option enables the user to save his input after the

recurring payment entry has been updated. In other embodiments, the remittance application automatically saves the recurring payment entry after any data field has been modified by the user. Also, according to some embodiments, the remittance application may prevent the user from saving changes to the recurring bill payment entry if a certain minimum number of prerequisite data fields have not been filled out, or otherwise, if the data entered within any of these fields is of an invalid format.

[0468] According to some embodiments, the user may be presented the option of how he wishes to schedule recurring payments with the payee. FIG. 48 is a flow diagram illustrating this process. At block 4802, the user may be prompted to select among the options of: “Immediately,” “Manually,” “By Schedule,” or “Return to Previous Menu.” The remittance application may then check which option was selected at respective decision blocks 4810, 4820, 4830, and 4840.

[0469] If the user selected to schedule bill payments with the payee “Immediately,” then at block 4812, the remittance application configures itself to attempt to make a payment soon after receiving an image of a check and/or remittance coupon from the user. The document images can be preprocessed by the mobile device 350 and/or processed by the remote server in any of the manners already described above. After the images have been successfully processed, one or more of the image quality assurance tests already described can then be run in real-time in order to ensure that the user has taken an image with a quality sufficient to process a payment.

[0470] If the user selected to schedule bill payments with the payee “Manually,” then at block 4822, the remittance application configures itself to attempt to make a payment only upon a specific input from the user. This input might be, for example, a “Pay Bill” button located in one or more menus or sub-menus of the remittance application. Images of any remittance coupons/checks received from the user may then be persistently stored within a non-volatile storage device until the user acknowledges he is ready to pay a certain bill by providing the specific input required by the remittance application.

[0471] If the user selected to schedule payments with the payee “By Schedule,” then at block 4832, a submenu may appear prompting the user to specify certain scheduling options. In some embodiments, the user may specify how many days he wishes the application to submit the payment before (or after) a certain payment due date. For example, if a utility bill is always due the 15th of every month, the user may elect to have these recurring bills paid on the 10th of every month. Images of any remittance coupons/checks received from the user may then be persistently stored within a non-volatile storage device until the scheduled date of payment. In some embodiments, any preprocessing, processing, or image quality and assurance tests are run on the document images soon after they are received from the user. This enables the user to detect and correct any defects with the image documents well before the scheduled date of payment.

[0472] Irrespective of the option selected, the user will be returned to scheduling menu after providing the input from the recurring payment sub-menu. If the user selected to “Return to Previous Menu,” then at block 4842 the user will be directed to the previous menu and the process will end.

[0473] According to some embodiments, the user may be presented the option of whether he wishes to have the remittance application send him one or more reminders about upcoming payment due dates. The reminders may thus serve

to assist the user in preventing a payment default due to inattention, inadvertence, or neglect.

[0474] FIG. 49 is a flow diagram illustrating an exemplary process of enabling a user to set one or more reminders associated with a recurring bill payment according to one embodiment of the present invention. At block 4902, a menu is displayed to the user, the menu including an option (such as a hyperlink or selectable button) for setting one or more payment reminders associated with a recurring payment schedule.

[0475] Once this option is selected at block 4904, then at block 4906, a sub-menu may then be displayed to the user. In some embodiments, the sub-menu presents the user with a number of configurable options regarding payment reminders. For example, the user may decide whether to set up a single reminder or a series of periodic reminders. Additionally, the user may specify when the reminders are to be sent (for example, on a regularly occurring day each month, such as on the 5th, or instead on a day that is always measured relative to the payment due date, such as 7 days before the bill is due). In some embodiments, the user may also specify how frequently the reminders are to be sent (e.g., daily, every third day, weekly, bi-weekly, etc.).

[0476] Additionally, according to some embodiments, the user may specify the type of reminders to be provided to the user by the remittance application. Any number of mechanisms for informing the user about an upcoming payment may be used according to embodiments of the present invention (including, but not limited to: e-mail, popup windows, SMS messages, “push”/PAP messaging, calendar alerts, scheduled printing, and phone messages/voicemail). Once the user has finished inputting preferred options at block 4908, the options are saved at block 4910, and the process then ends. Subsequently, the remittance application can provide payment reminders to the user in any manner or manner (s) that the user has specified.

Exemplary Hardware Embodiments

[0477] FIG. 50 is an exemplary embodiment of a mobile device 4400 according to an embodiment. Mobile device 4400 can be used to implement the mobile device 340 of FIG. 1. Mobile device 4200 includes a processor 4410. The processor 4410 can be a microprocessor or the like that is configurable to execute program instructions stored in the memory 4420 and/or the data storage 4440. The memory 4420 is a computer-readable memory that can be used to store data and or computer program instructions that can be executed by the processor 4410. According to an embodiment, the memory 4420 can comprise volatile memory, such as RAM and/or persistent memory, such as flash memory. The data storage 4440 is a computer readable storage medium that can be used to store data and or computer program instructions. The data storage 4440 can be a hard drive, flash memory, a SD card, and/or other types of data storage.

[0478] The mobile device 4400 also includes an image capture component 4430, such as a digital camera. According to some embodiments, the mobile device 4400 is a mobile phone, a smart phone, or a PDA, and the image capture component 4430 is an integrated digital camera that can include various features, such as auto-focus and/or optical and/or digital zoom. In an embodiment, the image capture component 4430 can capture image data and store the data in memory 4220 and/or data storage 4440 of the mobile device 4400.

[0479] Wireless interface 4450 of the mobile device can be used to send and/or receive data across a wireless network. For example, the wireless network can be a wireless LAN, a mobile phone carrier’s network, and/or other types of wireless network.

[0480] I/O interface 4460 can also be included in the mobile device to allow the mobile device to exchange data with peripherals such as a personal computer system. For example, the mobile device might include a USB interface that allows the mobile to be connected to USB port of a personal computer system in order to transfers information such as contact information to and from the mobile device and/or to transfer image data captured by the image capture component 4430 to the personal computer system.

[0481] As used herein, the term unit might describe a given unit of functionality that can be performed in accordance with one or more embodiments of the present invention. As used herein, a unit might be implemented utilizing any form of hardware, software, or a combination thereof. For example, one or more processors, controllers, ASICs, PLAs, logical components, software routines or other mechanisms might be implemented to make up a module. In implementation, the various modules described herein might be implemented as discrete modules or the functions and features described can be shared in part or in total among one or more modules. In other words, as would be apparent to one of ordinary skill in the art after reading this description, the various features and functionality described herein may be implemented in any given application and can be implemented in one or more separate or shared modules in various combinations and permutations. Even though various features or elements of functionality may be individually described or claimed as separate modules, one of ordinary skill in the art will understand that these features and functionality can be shared among one or more common software and hardware elements, and such description shall not require or imply that separate hardware or software components are used to implement such features or functionality.

[0482] Where components or modules of processes used in conjunction with the operations described herein are implemented in whole or in part using software, in one embodiment, these software elements can be implemented to operate with a computing or processing module capable of carrying out the functionality described with respect thereto. One such example-computing module is shown in FIG. 51, which illustrates a computer system that can be used to implement mobile remittance server according to an embodiment.

[0483] Various embodiments are described in terms of this example-computing module 1900. After reading this description, it will become apparent to a person skilled in the relevant art how to implement the invention using other computing modules or architectures.

[0484] Referring now to FIG. 51, computing module 1900 may represent, for example, computing or processing capabilities found within desktop, laptop and notebook computers; mainframes, supercomputers, workstations or servers; or any other type of special-purpose or general-purpose computing devices as may be desirable or appropriate for a given application or environment. Computing module 1900 might also represent computing capabilities embedded within or otherwise available to a given device. For example, a computing module might be found in other electronic devices. Computing module 1900 might include, for example, one or more processors or processing devices, such as a processor

1904. Processor 1904 might be implemented using a general-purpose or special-purpose processing engine such as, for example, a microprocessor, controller, or other control logic.

[0485] Computing module 1900 might also include one or more memory modules, referred to as main memory 1908. For example, random access memory (RAM) or other dynamic memory might be used for storing information and instructions to be executed by processor 1904. Main memory 1908 might also be used for storing temporary variables or other intermediate information during execution of instructions by processor 1904. Computing module 1900 might likewise include a read only memory (“ROM”) or other static storage device coupled to bus 1902 for storing static information and instructions for processor 1904.

[0486] The computing module 1900 might also include one or more various forms of information storage mechanism 1910, which might include, for example, a media drive 1912 and a storage unit interface 1920. The media drive 1912 might include a drive or other mechanism to support fixed or removable storage media 1914. For example, a hard disk drive, a floppy disk drive, a magnetic tape drive, an optical disk drive, a CD or DVD drive (R or RW), or other removable or fixed media drive. Accordingly, storage media 1914 might include, for example, a hard disk, a floppy disk, magnetic tape, cartridge, optical disk, a CD or DVD, or other fixed or removable medium that is read by, written to or accessed by media drive 1912. As these examples illustrate, the storage media 1914 can include a computer usable storage medium having stored therein particular computer software or data.

[0487] In alternative embodiments, information storage mechanism 1910 might include other similar instrumentalities for allowing computer programs or other instructions or data to be loaded into computing module 1900. Such instrumentalities might include, for example, a fixed or removable storage unit 1922 and an interface 1920. Examples of such storage units 1922 and interfaces 1920 can include a program cartridge and cartridge interface, a removable memory (for example, a flash memory or other removable memory module) and memory slot, a PCMCIA slot and card, and other fixed or removable storage units 1922 and interfaces 1920 that allow software and data to be transferred from the storage unit 1922 to computing module 1900.

[0488] Computing module 1900 might also include a communications interface 1924. Communications interface 1924 might be used to allow software and data to be transferred between computing module 1900 and external devices. Examples of communications interface 1924 might include a modem or softmodem, a network interface (such as an Ethernet, network interface card, WiMedia, IEEE 802.XX or other interface), a communications port (such as for example, a USB port, IR port, RS232 port Bluetooth® interface, or other port), or other communications interface. Software and data transferred via communications interface 1924 might typically be carried on signals, which can be electronic, electromagnetic (which includes optical) or other signals capable of being exchanged by a given communications interface 1924. These signals might be provided to communications interface 1924 via a channel 1928. This channel 1928 might carry signals and might be implemented using a wired or wireless communication medium. These signals can deliver the software and data from memory or other storage medium in one computing system to memory or other storage medium in computing system 1900. Some examples of a channel might include a phone line, a cellular link, an RF link, an

optical link, a network interface, a local or wide area network, and other wired or wireless communications channels.

[0489] Computing module 1900 might also include a communications interface 1924. Communications interface 1924 might be used to allow software and data to be transferred between computing module 1900 and external devices. Examples of communications interface 1924 might include a modem or softmodem, a network interface (such as an Ethernet, network interface card, WiMAX, 802.XX or other interface), a communications port (such as for example, a USB port, IR port, RS232 port, Bluetooth interface, or other port), or other communications interface. Software and data transferred via communications interface 1924 might typically be carried on signals, which can be electronic, electromagnetic, optical or other signals capable of being exchanged by a given communications interface 1924. These signals might be provided to communications interface 1924 via a channel 1928. This channel 1928 might carry signals and might be implemented using a wired or wireless medium. Some examples of a channel might include a phone line, a cellular link, an RF link, an optical link, a network interface, a local or wide area network, and other wired or wireless communications channels.

[0490] In this document, the terms “computer program medium” and “computer usable medium” are used to generally refer to physical storage media such as, for example, memory 1908, storage unit 1920, and media 1914. These and other various forms of computer program media or computer usable media may be involved in storing one or more sequences of one or more instructions to a processing device for execution. Such instructions embodied on the medium, are generally referred to as “computer program code” or a “computer program product” (which may be grouped in the form of computer programs or other groupings). When executed, such instructions might enable the computing module 1900 to perform features or functions of the present invention as discussed herein.

[0491] While various embodiments of the present invention have been described above, it should be understood that they have been presented by way of example only, and not of limitation. The breadth and scope of the present invention should not be limited by any of the above-described exemplary embodiments. Where this document refers to technologies that would be apparent or known to one of ordinary skill in the art, such technologies encompass those apparent or known to the skilled artisan now or at any time in the future. In addition, the invention is not restricted to the illustrated example architectures or configurations, but the desired features can be implemented using a variety of alternative architectures and configurations. As will become apparent to one of ordinary skill in the art after reading this document, the illustrated embodiments and their various alternatives can be implemented without confinement to the illustrated example. One of ordinary skill in the art would also understand how alternative functional, logical or physical partitioning and configurations could be utilized to implement the desired features of the present invention.

[0492] Furthermore, although items, elements or components of the invention may be described or claimed in the singular, the plural is contemplated to be within the scope thereof unless limitation to the singular is explicitly stated. The presence of broadening words and phrases such as “one or more,” “at least,” “but not limited to” or other like phrases

in some instances shall not be read to mean that the narrower case is intended or required in instances where such broadening phrases may be absent.

What is claimed is:

1. A method of capturing an image of a document on a mobile device, comprising:

initializing an image sensor on the mobile device to detect at least one image of the document;

measuring a value of at least one parameter related to quality of the at least one image being detected by the image sensor; and

capturing at least one image of the document when the value of the at least one measured parameter reaches or exceeds a threshold value.

2. The method of claim 1, wherein the at least one measured parameter is selected from at least one of: an angle of the mobile device with respect to the document; a distance of the mobile device from the document, movement of the mobile device, a lighting level of the image being detected, detection of all four corners of the document, a level of focus of the image being detected, a level of reflection of the document, and a size of the document within the image being detected.

3. The method of claim 1, further comprising adjusting at least one parameter if its value does not reach the threshold value.

4. The method of claim 1, wherein the threshold value of one parameter is changed based on the measured value of a separate parameter.

5. The method of claim 1, wherein a user of the mobile device is provided with feedback prior to the capturing of the at least one image which provides the user with at least one instruction for adjusting the at least one measured parameter.

6. The method of claim 5, wherein the feedback is visual feedback displayed on a display screen of the mobile device that is represented as a visual "icon" or "symbol," representing how optimized the image is during the capture process.

7. The method of claim 1, further comprising displaying the at least one image being detected by the image sensor in real-time on a display screen of the mobile device.

8. The method of claim 7, further comprising displaying a rectangular-shaped frame on the display screen while the image being detected by the image sensor is being displayed on the display screen.

9. The method of claim 8, wherein the size of the rectangular-shaped frame corresponds to a threshold size of the document needed before the at least one image can be captured.

10. The method of claim 9, wherein the rectangular-shaped frame changes color when edges of the document fit within the rectangular-shaped frame.

11. The method of claim 9, wherein the rectangular-shaped frame animates when edges of the document fit within the rectangular-shaped frame.

12. The method of claim 7, wherein a document-shaped outline is displayed over the probable document.

13. The method of claim 7, wherein an icon or image is displayed over the probable document, tracking the probable document per user movement.

14. The method of claim 7, wherein dynamic user feedback is also displayed on the display screen as text information.

15. The method of claim 7, wherein dynamic user feedback is also displayed on the display screen as non-text information.

16. The method of claim 7, wherein dynamic user feedback is given as audio cues.

17. The method of claim 7, wherein dynamic user feedback is also given as haptic cues.

18. The method of claim 7, wherein a centering image is displayed instead of a rectangular shaped frame.

19. A mobile device for capturing an image of a document, comprising:

an image sensor which detects at least one image of the document;

a measurement unit which measures a value of at least one parameter related to the quality of the at least one image detected by the image sensor; and

an image capture unit which captures the at least one image of the document when the value of the at least one measured parameter reaches a threshold value.

20. The mobile device of claim 19, wherein the at least one measured parameter is selected from at least one of: an angle of the mobile device with respect to the document; a distance of the mobile device from the document, movement of the mobile device, a lighting level of the image being detected, a contrast level of the image being taken, detection of all four corners of the document, a level of focus of the image being detected, a level of reflection of the document, and a size of the document within the image being detected.

21. The mobile device of claim 19, wherein the at least one parameter is adjusted if its value does not reach the threshold value.

22. The mobile device of claim 19, wherein the threshold value of one parameter is changed based on the measured value of a separate parameter.

23. The mobile device of claim 19, wherein a user of the mobile device is provided with feedback prior to the capturing of the at least one image which provides the user with at least one instruction for adjusting the at least one measured parameter.

24. The mobile device of claim 23, wherein the feedback is visual feedback displayed on a display screen of the mobile device.

25. The mobile device of claim 23, wherein the feedback is audio played through the speaker output of the mobile device.

26. The mobile device of claim 23, wherein the feedback is haptic actuated through the vibration actuator of the mobile device.

27. The mobile device of claim 24, wherein the at least one image being detected by the image sensor is displayed in real-time on a display screen of the mobile device.

28. The mobile device of claim 27, wherein a document size detection unit displays a rectangular-shaped frame on the display screen while the image being detected by the image sensor is being displayed on the display screen.

29. The mobile device of claim 28, wherein the size of the rectangular-shaped frame corresponds to a threshold size of the document needed before the at least one image can be captured.

30. The mobile device of claim 29, wherein the rectangular-shaped frame changes color or animates when edges of the document fit within the rectangular-shaped frame.

* * * * *