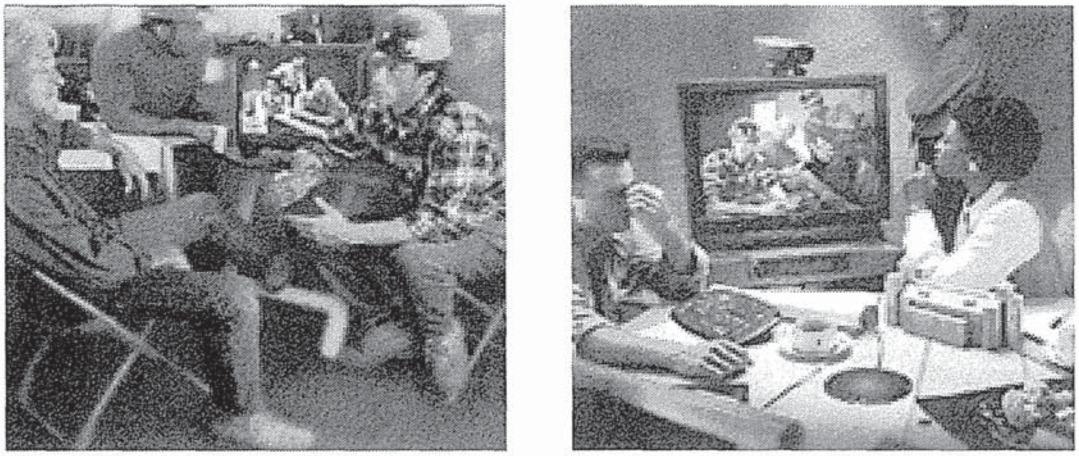

Figure 14.6

The PictureTel Venue 2000 Model 50, a system that supports high-quality videoconferencing, such as this discussion between architects and construction engineers. (Used with permission of PictureTel Corp.)

workspace-alone treatment as it did with the two other treatments. This result reinforces the importance of having a clear voice channel for coordination while users are looking at the objects of interest. The capacity for remote groups to produce work of a quality similar to that of face-to-face groups was demonstrated in studies of audio- and video-supported groups (Olson et al., 1995). Video support improved quality of work over audio only, and users preferred having the video support. The importance of audio and the marginal benefits of video for turn taking or interruption were highlighted in a comparison of three systems, but users expressed desire for video (Sellen, 1994).

Ethnographic observations and field studies reveal actual usage patterns and support competing theories. Kraut and associates (1994) found evidence that "critical mass—the numbers of people one can reach on a system—and social influence—the norms that grow up around a medium" are the key determinants of video-system success. Those people whose jobs involved substantial personnel management used video more than did those with more structured and document-oriented jobs.

The promise of video windows, tunnels, spaces, and so on is that they enable an enriched form of communication compared to a telephone conference or electronic mail, with less disruption than a trip. They enable participants to access the resources of their office environments while affording a chance for successful communication and emotional contact. Successful entrepreneurs will be those who best understand these new media and find the situations for which the media are best suited.

14.5 Face to Face: Same Place, Same Time

Teams of people often work together and use complex shared technology. Pilot and copilot cooperation in airplanes has been designed carefully with shared instruments and displays. Coordination among air-traffic controllers has a long history that has been studied thoroughly (Wiener and Nagel, 1988). Stock-market trading rooms and commodity markets are other existing applications of face-to-face teamwork or negotiations that are computer mediated.

Newer applications in office and classroom environments are attracting more attention because of the large numbers of potential users and the potential for innovative approaches to work and to learning. These applications include:

- *Shared display from lecturer workstation* In this simple form of group computing, a professor or lecturer may use the computer with a large-screen projector to demonstrate a computing application, to show a set of slides with business graphics, to retrieve images, or to run an animation. Fred Hofstetter (1995) of the University of Delaware developed a multimedia lectureware package, PODIUM, that allows instructors to compose illustrated lectures using slides, computer graphics, animations, videos, and audio sequences. Many speakers are happy to use standard commercial packages such as Microsoft PowerPoint, Lotus Freelance, or Adobe Persuasion. User-interface issues include simplicity in moving to the next slide, capacity to jump out of sequence, and ease of making spontaneous changes.
- *Audience response units* Simple keypads have been used effectively in training courses. Students can answer multiple-choice questions at their desks, and results can be shown to the full class on a large display. Similar units have been used by advertising researchers who ask test audiences to respond to commercials shown on a large screen. Votes in parliamentary forums can be rapid and accurate. Promoters claim that this simple technology is easy to learn, is acceptable to most people, is nonthreatening, and heightens attention because of the participatory experience. The National Geographic interactive exhibit gallery in Washington, D.C., has five-button response units that allow visitors to try their hand at answering multiple-choice questions such as "What percentage of the earth is covered by water?" The set of answers is shown on the shared display, but the presentation sequence is unaffected by the audience's selections.
- *Text-submission workstations* By giving each participant a keyboard and simple software, it is possible to create an inviting environment for conversation or brainstorming. Batson (Bruce et al., 1992) at Gallaudet University constructed a highly successful networking program that allows each participant to type a line of text that is shown immediately, with

the author's name, on every participant's display. With 10 people typing, new comments appear a few times per second and lively conversations ensue. Batson's goal was to overcome his frustrated efforts at teaching college-level English writing, and his English Natural Form Instruction (ENFI) network software was spectacularly successful:

> It seems slightly ironic that the computer, which for twenty-five years has been perceived as anti-human, a tool of control and suppression of human instinct and intuition, has really humanized my job. For the first time in a long time, I have real hope that we might make some progress. . . . Freed of having to be the cardboard figure at the front of the classroom, I became a person again, with foibles, feelings and fantasies. As a group, we were more democratic and open with each other than any other writing class I'd had. (Bruce et al., 1992).

The clatter of the keyboards adds to the laughter, groans, cheers, and grimaces to create a good atmosphere.

- *Brainstorming, voting, and ranking* Beyond talking, structured social processes can produce dramatic educational discussions and highly productive business meetings. The University of Arizona was a pioneer in developing the social process, the physical environment, and the software tools (Valacich et al., 1991) to "reduce or eliminate the dysfunctions of the group interaction so that a group reaches or exceeds its task potential" (Fig. 14.7). By allowing anonymous submission of suggestions and ranking of proposals, the authors introduced a wider range of possibilities; also, ideas were valued on their merits, independently of the originator (Fig. 14.8a–c). Because ego investments and conflicts were reduced, groups seemed to be more open to novel suggestions. IBM has built 19 Decision Center rooms based on the Arizona model for its internal use, and another

Figure 14.7

Semicircular classroom with 24 personal computers built into the desks at the University of Arizona. (Group Systems is a registered trademark of Ventana Corporation.)

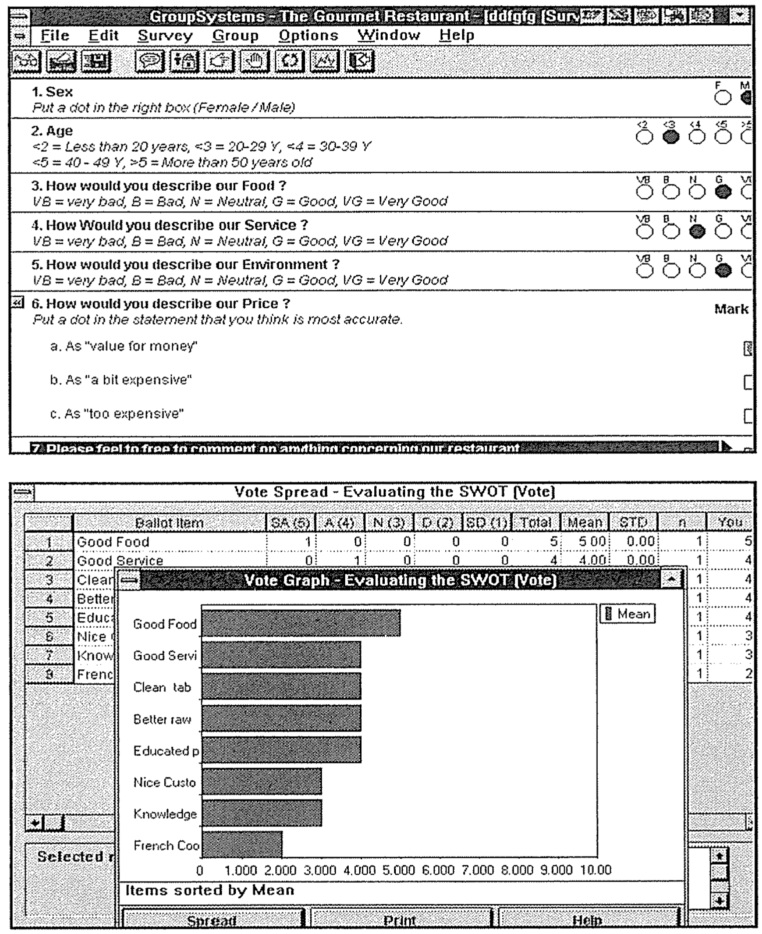

Figure 14.8

Sample screens from GroupSystems Electronic Meeting software. Online restaurant survey (top). Results of a vote in part of the restaurant survey (bottom). (Used with permission from Ventana Corp., Tucson, AZ.) (Group Systems is a registered trademark of Ventana Corporation.)

20 for rental to users under the TeamFocus name. Well-trained facilitators with backgrounds in social dynamics consult with the team leader to plan the decision session and to write the problem statement. In a typical task, 45 minutes of brainstorming by 15 to 20 people can produce hundreds of lines of suggestions for questions such as, "How can we increase sales?" Or, "What are the key issues in technological support for group work?" Then, items can be filtered, clustered into similar groups, and presented to participants for refinement and ranking. Afterward, a printout and electronic-file version of the entire session is immediately available. Numerous studies of electronic meeting systems with thousands of users have demonstrated and explored the benefits (Nunamaker et al., 1991):

- Parallel communication promotes broader input into the meeting process and reduces the chance that a few people dominate the meeting.

- Anonymity mitigates evaluation apprehension and conformance pressure, so issues are discussed more candidly.
- The group memory constructed by participants enables them to pause and reflect on information and on opinions of others during the meeting, and serves as a permanent record of what occurred.
- Process structure helps to focus the group on key issues, and discourages digressions and unproductive behaviors.
- Task support and structure provide information and approaches to analyze that information.

The University of Arizona system is marketed under the name GroupSystems (Ventana Corp.).

- *File sharing* A simple but powerful use of networked computers in a workplace, classroom, or meeting room is to share files. Participants may arrive with sales reports that can be shared with other people in the room rapidly. Alternatively, the group leaders may have agenda or budgets that they wish to broadcast to all participants, who may then annotate or embed these documents in others. Shared files may contain text, programs, spreadsheets, databases, graphics, animations, sound, X-ray images, or video. Presumably, distribution can go beyond the meeting room to allow participants to access the files from their offices and homes.
- *Shared workspace* The complement to each person receiving a personal copy of a file is to have a shared view of a workspace that every user can access. The pioneering Capture Lab at Electronic Data Systems contained an oval desk with eight Macintosh computers built into the desk to preserve the business-meeting atmosphere (Mantei, 1988). The large display in front of the desk is visible to all attendees, who can each take control of the large screen by pressing a button on a machine. At Xerox PARC, the research system Colab has generated the commercial large-screen (167-cm-diagonal) display, LiveBoard (Fig. 14.9), on which users can see the current list of topics or proposals, and can point to, edit, move, or add to under the policy sometimes called WYSIWIS (what you see is what I see) (Stefik et al., 1987). The advantage of a shared workspace is that everyone sees the same display and can work communally to produce a joint and recorded result (Weiser, 1991).
- *Group activities* With the proper networking software among workstations, users can be assigned a problem, and those needing assistance can "raise their hands" to show their display on a large shared display or on the group leader's display. Then, the group leader or other participants can issue commands to resolve the problem. Similarly, if participants have a particularly noteworthy result, graphic, or comment, they can share it with the group either on the large shared display or on individual workstations.

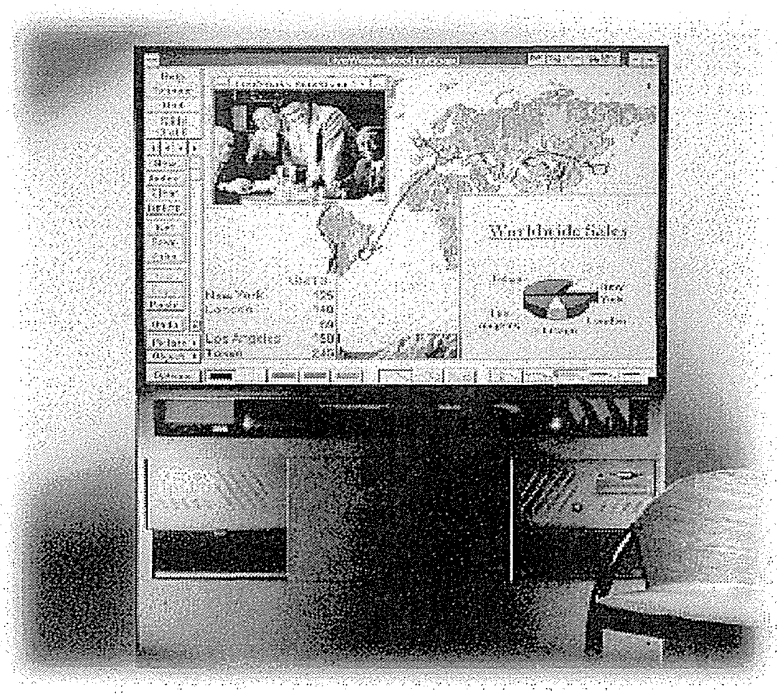

Figure 14.9

The LiveBoard Interactive Meeting System from LiveWorks, Inc., a Xerox company. Team discussions with groups at multiple locations can be facilitated with a 167-cm LiveBoard display. (Used with permission of LiveWorks, Inc.)

14.6 Applying CSCW to Education

The potential for a groupware-mediated paradigm shift in education evokes passion from devotees, but there is ample reason for skepticism and resistance. No single technology will dominate, but successful combinations will have to be suited to the goals of the institution, pedagogic style of the instructor, and availability of equipment for students. The long-promised but slow education revolution is speeding up as use of electronic mail and the web become widespread (Gilbert, 1996). Same-time, same-place electronic classrooms and a rich variety of distance-education strategies are promoted as ways to improve quality or to lower costs, but a change in teaching and learning styles and the inclusion of new students are often the main result (Harasim et al., 1995).

Coordination of students in a *virtual classroom* is a complex process but it can enable a stimulating educational experience for people who cannot travel to a regular classroom (Hiltz, 1992). Multiple trials with sociology, computer-science, and philosophy courses demonstrated the efficacy of a conference format for college courses, complete with homework assignments, projects, tests, and final examinations. Instructors found the constant flow of messages to be a rewarding challenge, and students were generally satisfied with the experience:

> The essence of the Virtual Classroom is an environment to facilitate collaborative learning. For distance education students, the increased ability to be in constant communication with other learners is obvious. But even for campus-based courses the technology provides a means for a rich, collaborative learning environment which exceeds the traditional classroom in its ability to 'connect' students and course materials on a round-the-clock basis. (Hiltz, 1992)

Distance education with broadcast-quality video lectures is common, but interactivity with students is often by telephone, electronic mail, or web exchanges. DTVC has the potential to create livelier two-way interactions for discussion, mentoring, and remediation. The greatest beneficiaries are professionals who can attend courses electronically from their offices or special learning centers, and home-oriented students who cannot commit the time for travel to a traditional campus. Current desktop videoconferencing facilitates communication, but improvements are needed to give instructors better awareness of reactions at multiple sites and ways to manage smoother turn taking (Ramsay et al., 1996). Improved resolution will help to convey gesture, gaze direction, and body language, but seeing detail and context simultaneously at multiple sites is a challenge (Fussel and Benimoff, 1995).

The electronic classrooms at the University of Maryland balance the pursuit of new technologies with the exploration of new teaching and learning styles (Shneiderman et al., 1995). Three classrooms were built with 40 seats and 20 high-resolution monitors partially recessed into the desks to preserve sightlines (Fig. 14.10). The computers were placed in a side room to increase security and room space and to reduce noise and heat. A workstation and two large rear-projected displays enable instructors to show everyone their screen or any student screen. Keys to success included provision of the necessary infrastructure for faculty training and support, and collection of ample evaluation data to guide the process.

Over the first six years, 68 faculty (30 tenured, 16 nontenured, 22 other staff) from 21 departments offered 233 courses with over 6782 students. Courses filled most slots from 8 A.M. to 10 P.M., and were as diverse as "The Role of Media in the American Political Process," "Chinese Poetry into English," "Marketing Research Methods," "Database Design," and "Saving the Bay."

Faculty members who used the electronic classrooms explored novel teaching and learning styles that can create more engaging experiences for students. While traditional lectures with or without discussion remain common, electronic-classroom technologies can enliven lectures (Hofstetter,

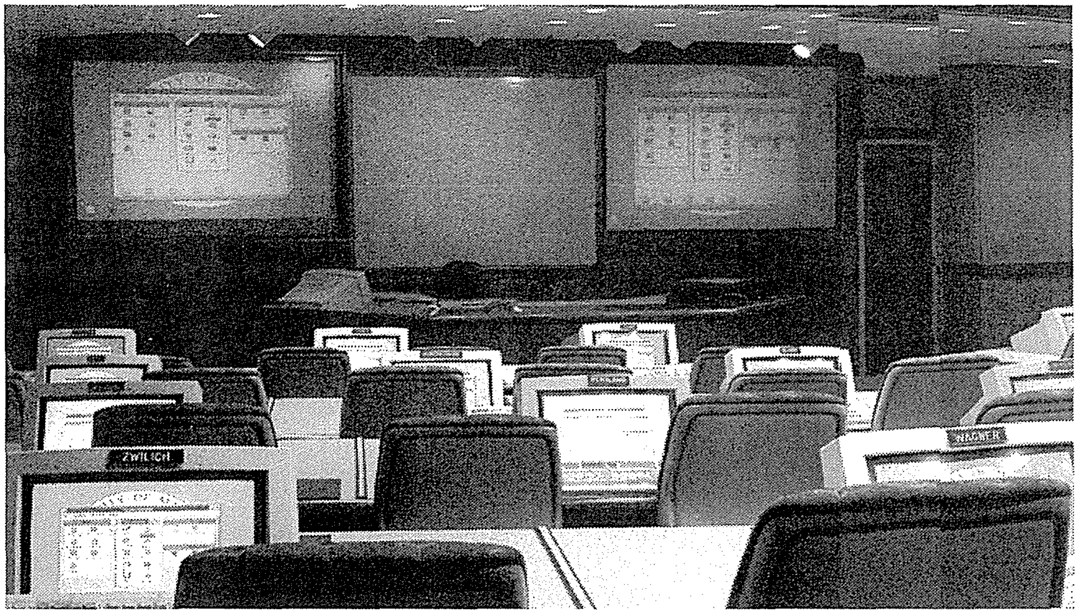

Figure 14.10

AT&T Teaching/Learning Theater at the University of Maryland has 20 high-resolution displays built into custom desks with seats for 40 students.

1995) while enabling active individual learning, small-group collaborative learning, and entire-class collaborative learning. Most faculty acknowledge spending more preparation time to use the electronic classroom especially in their first semester, but one wrote that it is "well worthwhile in terms of greater learning efficiency."

The assumption that improved lectures were the main goal changed as faculty tried collaborative teaching methods and talked about these methods with one another. Faculty who had used paper-based collaborations appreciated the smoothness of showing electronic student submissions to the whole class. Faculty who had not used collaborative methods appreciated the ease and liveliness of an anonymous electronic brainstorming session.

More active individual learning experiences include using software during class time

- To write essays in English or poems in a foreign language
- To find antecedents of Impressionism in an art-history library of 9000 images
- To run business simulations to increase product quality
- To perform statistical analyses of psychology studies

- To do landscaping with computer-assisted design and graphics packages
- To compose computer programs
- To search the Internet

A common teacher strategy (Norman, 1994) is to assign time-limited (3 to 10 minutes) tasks, and then to use the video switcher to review the students' work, to give individual help when necessary, and to show the students' work to the entire class. The transformational breakthrough lies in opening the learning process by rapidly showing many students' work to the entire class. Doing so at first generates student and faculty anxiety, but quickly becomes normal. Seeing and critiquing exemplary and ordinary work by fellow students provides feedback that inspires better work on subsequent tasks.

Small-group collaborative-learning experiences include having pairs of students work together at a machine on a time-limited task. Pairs often learn better than individuals, because people can discuss their problems, learn from each other, and split their roles into problem solver and computer operator. With paired teams, the variance of completion time for tasks is reduced compared to individual use, and fewer students get stuck in completing a task. Verbalization of problems has often been demonstrated to be advantageous during learning and is an important job skill to acquire for modern team-oriented organizations.

Innovative approaches with larger teams include simulated hostage negotiations with terrorist airplane hijackers in a course on conflict resolution, and business trade negotiations in a United Nations format for a course on commercial Spanish. Teams work to analyze situations, to develop position statements online, and to communicate their positions to their adversaries over the network. In an introductory programming course, 10 teams wrote components and sent them through the network to the lead team, who combined the pieces into a 173-line program, all in 25 minutes. The class performed a walkthrough of the code using the large-screen display, and quickly identified bugs.

Some faculty find that adapting to the electronic-classroom environment changes their styles so much that they teach differently even in traditional classrooms. Other faculty vow that they will never teach in a traditional classroom again. Most faculty users want to continue teaching in these electronic classrooms and discover that more than their teaching styles change—their attitudes about the goals of teaching and about the content of the courses often shift as well. Many faculty develop higher expectations for student projects. Some become evangelists within their disciplines for the importance of teamwork and its accompanying communications skills.

On the negative side, a math professor who used the computers only to do occasional demonstrations returned to teaching in a traditional classroom, where he had much more blackboard space. Some reluctant instructors express resistance to changing their teaching styles and anticipate having to make a large effort to use the electronic classrooms.

Evaluations included standard course evaluations, use of anonymous electronic ratings, and specially prepared questionnaires. A controlled study with 127 students (Alavi, 1994) indicated that electronic-classroom students had higher perceived skill development, self-reported learning, and evaluation of classroom experience than did students in a collaborative-learning traditional classroom. Electronic-classroom students also had statistically significantly higher final-exam grades. Popular features were the electronic note taking, interactivity, idea sharing, and brainstorming.

Evaluations revealed problems with network access from outside the classrooms and with file-sharing methods within the classroom. Students generally were positive, and often were enthusiastic: "Everyone should have a chance to be in here at least once. . . . Great tech. Great education technique. . . . Easy to use, but tends to crash and die at times. . . . the best thing that I could think of to improve the ability to teach interactively. Even though there were a few humps to get over at the beginning—it was well worth the effort (and money)."

Intense interest in educational technology and in new teaching strategies is widespread. Resource-rich universities are investing in teaching–learning theaters; others are making innovative use of electronic mail, listservs, and the web (Gilbert, 1996). Distance learning using CSCW technologies seems likely to expand.

14.7 Practitioner's Summary

Computing has become a social process. The networks and telephone lines have opened up possibilities for cooperation. Electronic mail has made it easy to reach out and touch someone, or thousands of someones. Newsgroups, electronic conferences, and the web have enabled users to be in closer communication. Coordination within projects or between organizations is facilitated by text, graphic, voice, and even video exchanges. Even face-to-face meetings are getting a facelift with new tools for electronic meetings and with teaching–learning theaters. The introspective and isolated style of past computer use is giving way to a lively social environment where training has to include *netiquette* (network etiquette). These collaboration tools are beginning to have a visible effect; it seems that their success will continue spreading. However, as there are in all new technologies, there will be failures and surprising discoveries, because our intuitions about the design of groupware are based on shallow experience (Box 14.1). Thorough testing of new applications is necessary before widespread dissemination.

Box 14.1

Questions for consideration. The novelty and diversity of computer-supported cooperative work means that clear guidelines have not emerged, but these sobering questions might help designers and managers.

Computer-Supported Cooperative Work Questions

- How would facilitating communication improve or harm teamwork?
- Where does the community of users stand on centralization versus decentralization?
- What pressures exist for conformity versus individuality?
- How is privacy compromised or protected?
- What are the sources of friction among participants?
- Is there protection from hostile, aggressive, or malicious behavior?
- Will there be sufficient equipment to support convenient access for all participants?
- What network delays are expected and tolerable?
- What is the user's level of technological sophistication or resistance?
- Who is most likely to be threatened by computer-supported cooperative work?
- How will high-level management participate?
- Which jobs may have to be redefined?
- Whose status will rise or fall?
- What are the additional costs or projected savings?
- Is there an adequate phase-in plan with sufficient training?
- Will there be consultants and adequate assistance in the early phases?
- Is there enough flexibility to handle exceptional cases and special needs (disabilities)?
- What international, national, organizational standards must be considered?
- How will success be evaluated?

14.8 Researcher's Agenda

The opportunities for new products and for refinements of existing products seem great. Even basic products such as electronic mail could be improved dramatically by inclusion of advanced features, such as online directories, filtering, and archiving tools, as well as by universal-access features, such as improved tutorials, better explanations, and convenient assistance. Conferencing methods and cooperative document production will change as bandwidth increases and video is added. The most dramatic projects thus far are the ambitious electronic-meeting systems and teaching–learning theaters. They are costly, but are so attractive that many organizations are likely to spend heavily on these new technologies during the next decade. Although user-interface design of applications will be a necessary component, the larger and more difficult research problems lie in studying the social processes. How will home life and work be changed? How might interfaces differ for games, cooperative work, and conflict-laden online negotiations? Some of the excitement for researchers in computer-supported cooperative work stems from the vast uncharted territory: theories are sparse, controlled studies are difficult to arrange, data analysis is overwhelming, and predictive models are nonexistent (Olson et al., 1993).

World Wide Web Resources WWW

Computer Supported Cooperative Work is naturally a part of the World Wide Web and novel tools are springing up on many websites. You can try various chat services, download special purpose software, or shop for conferencing tools (video, audio, or text-based). Evaluations are also available online.

http://www.aw.com/DTUI

References

Alavi, Maryam, Computer mediated collaborative learning: An empirical evaluation, *MIS Quarterly*, 18, 2 (June 1994), 159–173.

Anderson, Robert H., Bikson, Tora K., Law, Sally Ann, and Mitchell, Bridger M., *Universal Access to Email: Feasibility and Societal Implications*, RAND, Santa Monica, CA (1995), also at http://www.rand.org.

Baecker, Ron, *Readings in Groupware and Computer-Supported Cooperative Work: Assisting Human–Human Collaboration*, Morgan Kaufmann, San Francisco, CA (1993).

Bly, Sara A., Harrison, Steve R., and Irwin, Susan, MediaSpaces: Bringing people together in a video, audio, and computing environment, *Communications of the ACM*, 36, 1 (January 1993), 28–47.

Bruce, Bertram, Peyton, Joy, and Batson, Trent, *Network-Based Classrooms*, Cambridge University Press, Cambridge, U.K. (1992).

Borenstein, Nathaniel S., Multimedia electronic mail: Will the dream become a reality? *Communications of the ACM*, 34, 4 (April 1991), 117–119.

Carroll, John M. and Rosson, Mary Beth, Developing the Blacksburg Electronic Village. *Communications of the ACM*, 39, 12 (December 1996), 69–74.

Chapanis, Alphonse, Interactive human communication, *Scientific American*, 232, 3 (March 1975), 36–42.

Crowley, Terrence, Milazzo, Paul, Baker, Ellie, Forsdick, Harry, and Tomlinson, Raymond, MMConf: An infrastructure for building shared multimedia applications, *Proc. Third Conference on Computer-Supported Cooperative Work*, ACM, New York (1990), 329–355.

Dorcey, Tim, CU-SeeMe desktop videoconferencing software, *Connexions*, 9, 3 (March 1995). Also at http://cu-seeme.cornell.edu/DorceyConnexions.html.

Ellis, C. A., Gibbs, S. J., and Rein, G. L., Groupware: Some issues and experiences, *Communications of the ACM*, 34, 1 (January 1991), 680–689.

Fischer, Gerhard and Stevens, Curt, Information access in complex, poorly structured information spaces, *Proc. ACM CHI '91 Human Factors in Computing Systems*, ACM, New York (1991), 63–70.

Fish, Robert S., Kraut, Robert E., and Chalfonte, Barbara, The VideoWindow System in informal communications, *Proc. Third Conference on Computer-Supported Cooperative Work*, ACM, New York (1990), 1–11.

Flores, F., Graves, M., Hartfield, B., and Winograd, T. Computer systems and the design of organizational interaction, *ACM Transactions on Office Information Systems*, 6, 2 (April 1988), 153–172.

Fussel, S. R. and Benimoff, I., Social and cognitive processes in interpersonal communications: Implications for advanced telecommunications technologies, *Human Factors*, 27, 2 (1995), 228–250.

Gale, Stephen, Human aspects of interactive multimedia communication, *Interacting with Computers*, 2, 2 (1990), 175–189.

Gilbert, Steven, Making the most of a slow revolution, *Change: The Magazine of Higher Learning*, 28, 2 (March/April 1996), 10–23.

Greenberg, Saul, Hayne, Stephen, and Rada, Roy (Editors), *Groupware for Real Time Drawing: A Designer's Guide*, McGraw-Hill, New York (1995).

Greenberg, Saul and Marwood, David, Real time groupware as a distributed system: Concurrency control and its effect on the interface, *Proc. Conference on Computer Supported Cooperative Work '94*, ACM, New York (1994), 207–217.

Grudin, Jonathan, Groupware and social dynamics: Eight challenges for developers, *Communications of the ACM*, 37, 1 (January 1994), 93–105.

Harasim, Linda, Hiltz, Starr Roxanne, Teles, Lucio, and Turoff, Murray, *Learning Networks: A Field Guide to Teaching and Learning Online*, MIT Press, Cambridge, MA (1995).

Hiltz, S. R., *Online Communities: A Case Study of the Office of the Future*, Ablex, Norwood, NJ (1984).

Hiltz, S. R., *The Virtual Classroom*, Ablex, Norwood, NJ (1992).

Hiltz, S. R. and Turoff, M., *The Network Nation: Human Communication via Computer*. Addison-Wesley, Reading, MA (1978).

Hiltz, S. R., and Turoff, M., Structuring computer-mediated communication systems to avoid information overload, *Communications of the ACM*, 28, 7 (July 1985), 680–689.

Hoffert, Eric M. and Gretsch, Greg, The digital news system at EDUCOM: A convergence of interactive computing, newspapers, television and high-speed networks, *Communications of the ACM*, 34, 4 (April 1991), 113–116.

Hofstetter, Fred T., *Multimedia Literacy*, McGraw-Hill, New York (1995).

Isaacs, Ellen, Morris, Trevor, Rodriguez, Thomas K., and Tang, John C., A comparison of face-to-face and distributed presentations, *Proc. CHI '95 Conference: Human Factors in Computing Systems*, ACM, New York (1995), 354–361.

Isaacs, Ellen, Tang, John C., and Morris, Trevor, *Proc. Conference on Computer Supported Cooperative Work '96*, ACM, New York (1996), 325–333.

Ishii, H., Kobayashi, M., and Arita, K., Iterative design of seamless collaboration media: From TeamWorkStation to ClearBoard, *Communications of the ACM*, 37, 8 (1994), 83–97.

Kraut, Robert E., Cool, Colleen, Rice, Ronald E., and Fish, Robert S., Life and death of new technology: Task, utility and social influences on the use of a communications medium, *Proc. Conference on Computer Supported Cooperative Work '94*, ACM, New York (1994), 13–21.

Malone, T., and Crowston, K. What is coordination theory and how can it help design cooperative work systems? In *Proc. Third Conference on Computer-Supported Cooperative Work*, ACM, New York (1990), 357–370.

Malone, T. W., Grant, K. R., Turbak, F. A., Brobst, S. A., and Cohen, M. D., Intelligent information-sharing systems, *Communications of the ACM*, 30 (1987), 390–402.

Mantei, M., Capturing the capture lab concepts: A case study in the design of computer supported meeting environments, *Proc. Second Conference on Computer-Supported Cooperative Work*, ACM, New York (1988), 257–270.

Mantei, Marilyn M., Baecker, Ronald S., Sellen, Abigail J., Buxton, William A. S., and Milligan, Thomas, Experiences in the use of a media space, *Proc. Conference: CHI '91 Human Factors in Computing Systems*, ACM, New York (1991), 203–208.

Mark, Gloria Haake, Jorg M., and Streitz, Norbert A., Hypermedia Structures and the Division of Labor in Meeting Room Collaboration, *Proc. Conference on Computer Supported Cooperative Work '96*, ACM, New York (1996), 170–179.

Mitchell, Alex, Posner, Ilona, and Baecker, Ronald, Learning to write together using groupware, *Proc. CHI '95 Conference: Human Factors in Computing Systems*, ACM, New York (1995), 288–295.

Norman, Kent, Navigating the educational space with HyperCourseware, *Hypermedia*, 6, 1 (January 1994), 35–60.

Nunamaker, J. F., Dennis, Alan R., Valacich, Joseph S., Vogel, Douglas R., and George, Joey F., Electronic meeting systems to support group work, *Communications of the ACM*, 34, 7 (July 1991), 40–61.

Olson, Judith S., Card, Stuart K., Landauer, Thomas K., Olson, Gary M., Malone, Thomas, and Leggett, John, Computer supported co-operative work: Research issues for the 90s, *Behaviour & Information Technology*, 12, 2 (1993), 115–129.

Olson, Judith S., Olson, Gary M., Mack, Lisbeth A., and Wellner, Pierre, Concurrent editing: The group's interface. In Diaper, D., Gilmore, D., Cockton, G., and Shackel, B. (Editors), *Human–Computer Interaction—INTERACT '90*, Elsevier Science Publishers, Amsterdam, The Netherlands (1990), 835–840.

Olson, Judith S., Olson, Gary M., and Meader, David K., What mix of video and audio is useful for small groups doing remote real-time design work?, *Proc. CHI '95 Conference: Human Factors in Computing Systems*, ACM, New York (1995), 362–368.

Patterson, John F., Hill, Ralph D., and Rohall, Steven L., Rendezvous: An architecture for synchronous multi-user applications, *Proc. Third Conference on Computer-Supported Cooperative Work*, ACM, New York (1990), 317–328.

Ramsay, J., Barabesi, A. and Preece, J. (1996) Informal communication is about sharing objects in media, Interacting with Computers, 8, 3, 227–283.

Rein, Gail L. and Ellis, Clarence A., rIBIS: a real-time group hypertext system, *International Journal of Man–Machine Studies*, 34, 3 (1991), 349–367.

Rheingold, Howard, *The Virtual Community: Homesteading on the Electronic Frontier*, Addison-Wesley, Reading, MA (1993).

Roseman, Mark and Greenberg, Saul, TeamRooms: Network places for collaboration, *Proc. Conference on Computer Supported Cooperative Work '96*, ACM, New York (1996), 325–333.

Schuler, Doug, *New Community Networks: Wired for Change*, Addison-Wesley, Reading, MA (1996).

Sellen, Abigail J., Remote conversations: The effects of mediating talk with technology, *Human–Computer Interaction*, 10, 4 (1994), 401–444.

Shneiderman, Ben, Alavi, Maryam, Norman, Kent, and Borkowski, Ellen Y., Windows of opportunity in electronic classrooms, *Communications of the ACM*, 38, 11 (November 1995), 19–24.

Sproull, Lee and Kiesler, Sara, *Connections: New Ways of Working in the Networked Organization*, MIT Press, Cambridge, MA (1991).

Stefik, M., Bobrow, D. G., Foster, G., Lanning, S., and Tartar, D., WYSIWIS revised: Early experiences with multiuser interfaces, *ACM Transactions on Office Information Systems*, 5, 2 (April 1987), 147–186.

Streitz, N., Haake, J., Hannemann, J., Lemke, A., Schuler, W., Schuett, H., and Thuering, M., SEPIA: A cooperative hypermedia authoring environment, In Rada, Roy (Editor), *Groupware and Authoring*, Academic Press, London, U.K. (1996), 241–264.

Turkle, Sherry, *Life on the Screen: Identity in the Age of the Internet*, Simon and Schuster, New York (1995).

Valacich, J. S., Dennis, A. R., and Nunamaker, Jr., J. F., Electronic meeting support: The GroupSystems concept, *International Journal of Man–Machine Studies*, 34, 2 (1991), 261–282.

Vaske, Jerry and Grantham, Charles, *Socializing the Human–Computer Environment*, Ablex, Norwood, NJ (1990).

Watabe, Kazuo, Sakata, Shiro, Maeno, Kazutoshi, Fukuoka, Hideyuki, and Ohmori, Toyoko, Distributed multiparty desktop conferencing system: MERMAID, *Proc. Third Conference on Computer-Supported Cooperative Work*, ACM, New York (1990), 27–28.

Weiser, Mark, The computer for the twenty-first century, *Scientific American*, 265, 3 (September 1991), 94–104.

Wiener, Earl L. and Nagel, David C. (Editors), *Human Factors in Aviation*, Academic Press, New York (1988).

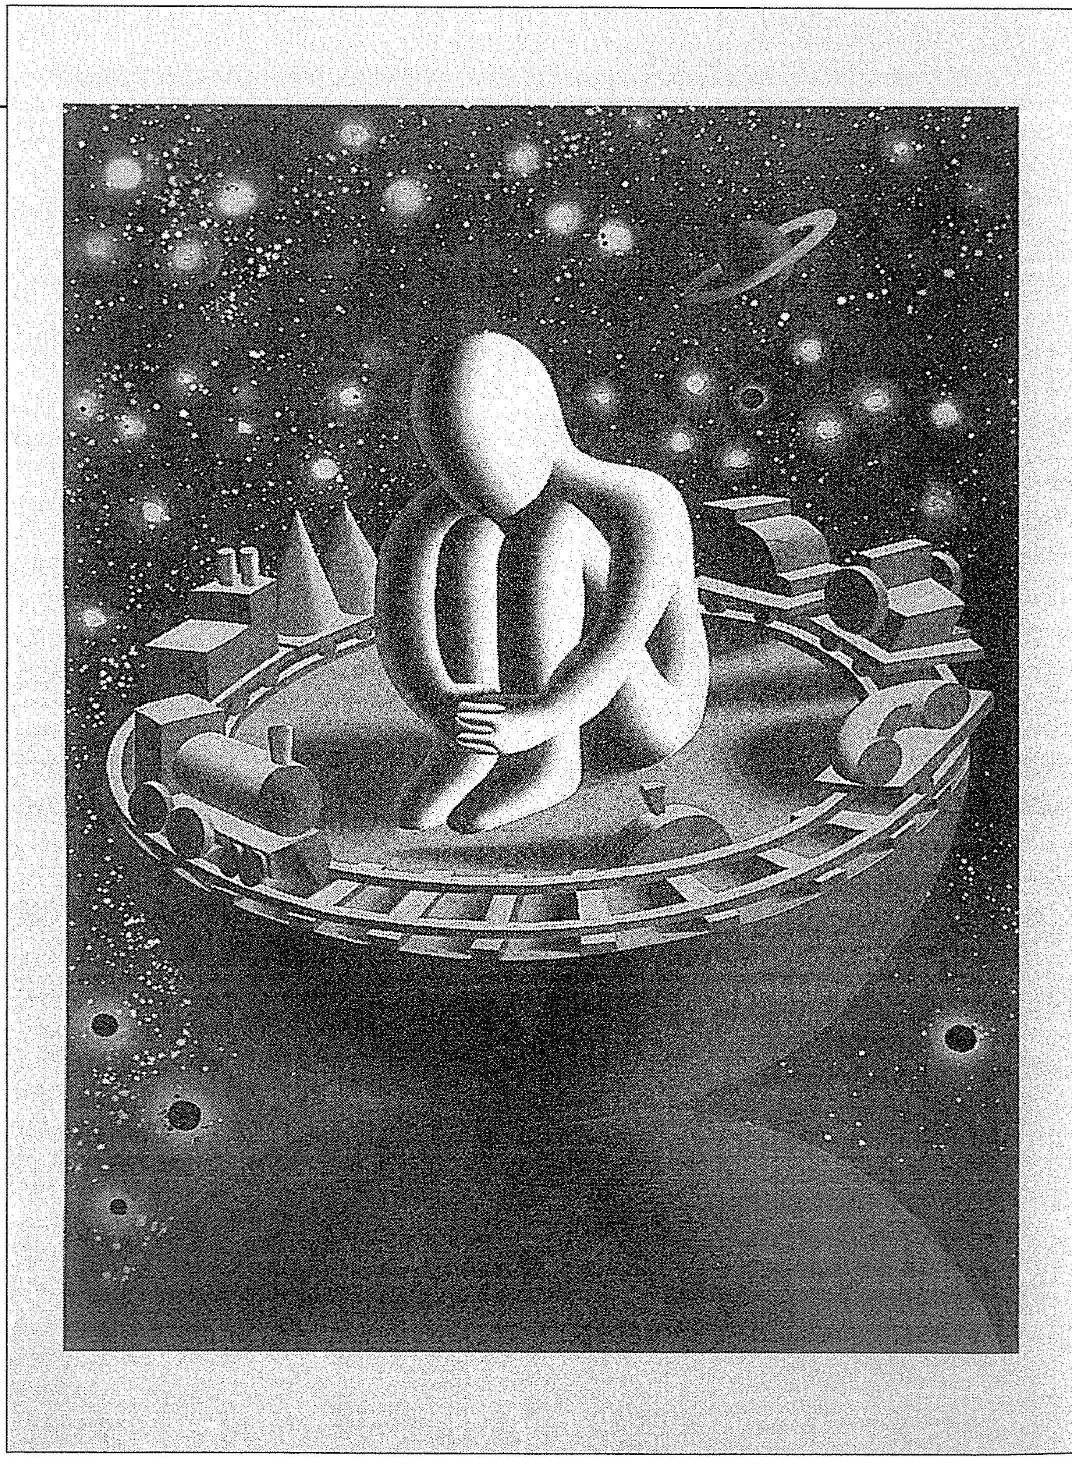

Mark Kostabi, *The Industrial Revolution (Luminary)*, 1994

CHAPTER

15

Information Search and Visualization

Everything points to the conclusion that the phrase "the language of art" is more than a loose metaphor, that even to describe the visible world in images we need a developed system of schemata.

E. H. Gombrich, *Art and Illusion,* 1959 (p. 76)

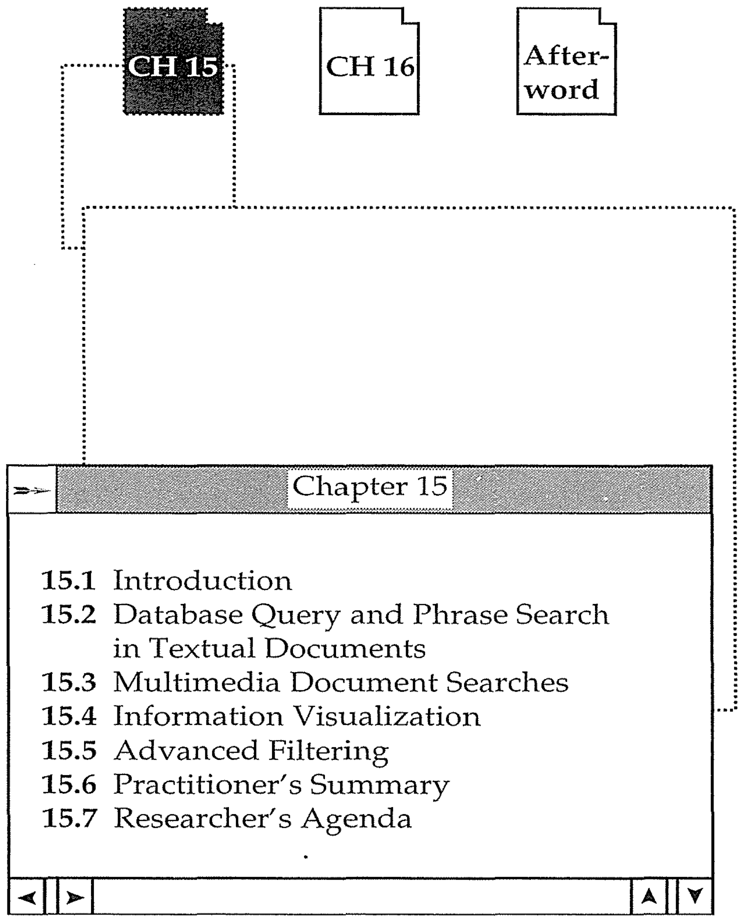

15.1 Introduction

Information exploration should be a joyous experience, but many commentators talk of information overload and anxiety (Wurman, 1989). However, there is promising evidence that the next generation of digital libraries will enable convenient exploration of growing information spaces by a wider range of users. User-interface designers are inventing more powerful search and visualization methods, while offering smoother integration of technology with task.

The terminology swirl in this domain is especially colorful. The older terms of *information retrieval* (often applied to bibliographic and textual document systems) and *database management* (often applied to more structured relational database systems with orderly attributes and sort keys), are being pushed aside by newer notions of *information gathering, seeking, filtering,* or *visualization*. Business-oriented developers focus on the huge volumes of data when they talk of *data mining* and *warehousing*, while expert-system visionaries talk about *knowledge networks*. The distinctions are subtle; the common goals range from finding a narrow set of items in a large collection that satisfy a well-understood information need (known-item search) to browsing to discover unexpected patterns within the collection (Marchionini, 1995).

Exploring information collections becomes increasingly difficult as the volume and diversity grows. A page of information is easy to explore, but when the information representation becomes the size of a book, or library, or even larger, it may be difficult to locate known items or to browse to gain an overview. The strategies to focus and narrow are well understood by librarians and information-search specialists, and now these strategies are being implemented for widespread use. The computer is a powerful tool for searching, but traditional user interfaces have been a hurdle for novice users (complex commands, Boolean operators, unwieldy concepts) and an inadequate tool for experts (difficulty in repeating searches across multiple databases, weak methods for discovering where to narrow broad searches, poor integration with other tools) (Borgman, 1986). This chapter suggests novel possibilities for first-time or intermittent versus frequent computer users, and also for task novices versus experts. Improvements on traditional text and multimedia searching seem possible as a new generation of visualization strategies for query formulation and information presentation emerges.

Designers are just discovering how to use rapid and high-resolution color displays to present large amounts of information in orderly and user-controlled ways. Perceptual psychologists, statisticians, and graphic designers (Bertin, 1983; Cleveland, 1993; Tufte, 1983, 1990) offer valuable guidance about presenting static information, but the opportunity for dynamic displays takes user-interface designers well beyond current wisdom.

The objects–actions interface (OAI) model (see Fig. 2.2) helps by separating task concepts (do you think of your organization as a hierarchy or a matrix?) from interface concepts (is your hierarchy can best represented as an outline, node–link diagram, or a treemap?). The OAI model also separates high-level interface issues (are overview diagrams necessary for navigation?) from low-level interface issues (will color or size coding be used to represent salary levels?).

First-time users of an information-exploration system (whether they have little or much task knowledge) are struggling to understand what they see on the display while keeping in mind their information needs. They would be distracted if they had to learn complex query languages or elaborate shape-coding rules. They need the low cognitive burdens of menu and direct-manipulation designs and simple visual-coding rules. As users gain experience with the interface, they can request additional features by adjusting control panels. Knowledgeable and frequent users want a wide range of search tools with many options that allow them to compose, save, replay, and revise increasingly elaborate query plans.

To facilitate discussion, we need to define a few terms. *Task objects,* such as Leonardo's notebooks or sports-video segments from the Olympics, are represented by *interface objects* in structured relational databases, textual document libraries, or multimedia document libraries. A *structured relational database* consists of *relations* and a *schema* to describe the relations. Relations have *items* (usually called *tuples* or *records*), and each item has multiple *attributes* (often called *fields*), which each have *attribute values.* In the relational model, items form an unordered set (although one attribute can contain sequencing information or be a unique key to identify or sort the other items) and attributes are *atomic.*

A *textual document library* consists of a set of *collections* (typically up to a few hundred collections per library) plus some *descriptive attributes* about the library (for example, name, location, owner). Each collection has a *name* plus some descriptive attributes about the collection (for example, location, media type, curator, donor, dates, geographic coverage), and a set of items (typically 10 to 100,000 items per collection). Items in a collection may vary greatly, but usually a moderate-sized superset of attributes exists that covers all the items. Attributes may be blank, have single values, have multiple values, or be lengthy texts. A collection is owned by a single library, and an item belongs to a single collection, although exceptions are possible. A *multimedia document library* consists of collections of documents that can contain images, sound, video, animations, and so on.

Task actions such as *fact finding* are decomposed into *browsing* or *searching,* and are represented by *interface actions* such as scrolling, zooming, joining, or linking. Users begin by formulating their information needs in the task domain. Tasks can range from specific fact finding, where there is a single readily identifiable outcome, to more extended fact finding, with uncertain but replicable outcomes. Relatively unstructured tasks include open-ended browsing of known collections and exploration of the availability of information on a topic:

Specific fact finding (known-item search)

Find the Library of Congress call number of "Future Shock."
Find the telephone number of Bill Clinton.
Find the highest-resolution LANDSAT image of College Park at noon on Dec. 13, 1997.

Extended fact finding

What other books are by the author of Jurassic Park?
What genres of music is Sony publishing?
Which satellites took images of the Persian Gulf War?

Open-ended browsing

Does the Mathew Brady Civil War photo collection show the role of women in that war?
Is there new work on voice recognition being reported from Japan?
Is there a relationship between carbon-monoxide levels and desertification?

Exploration of availability

What genealogy information is available at the National Archives?
What information is available on the Grateful Dead band members?
Do NASA datasets demonstrate acid-rain damage to soy crops?

Once users have clarified their information needs, the first step in satisfying those needs is to decide where to search (Marchionini, 1995). The conversion of information needs, stated in task-domain terminology, to interface actions is a large cognitive step, but it must be accomplished before expression of these actions in a query language or via a series of mouse selections can begin.

Supplemental *finding aids* can help users to clarify and pursue their information needs. Examples include tables of contents or indexes in books, descriptive introductions, concordances, key-word-in-context (KWIC) lists, and subject classifications. Careful understanding of previous and potential search requests, and of the task analysis, can improve search results by allowing the system to offer hot-topic lists and useful classification schemes. For example, the U.S. Congressional Research Service has a list of approximately 80 hot topics covering current bills before Congress, and has 5000 terms in its Legislative Indexing Vocabulary. The National Library of Medicine maintains the Medical Subject Headings (MeSH), with 14,000 items in a seven-level hierarchy.

This chapter covers database query and textual-document searches briefly, then suggests innovative directions for multimedia-document searches and introduces a four-phase framework. The main contribution is a taxonomy of information-visualization strategies based on data types and user tasks. The final section explores advanced filtering methods.

15.2 Database Query and Phrase Search in Textual Documents

Searching in structured relational database systems is a well-established task for which the SQL language has become a widespread standard (Reisner, 1988). Users write queries that specify matches on attribute values, such as author, date of publication, language, or publisher. Each document has values for the attributes, and database-management methods enable rapid retrieval even with millions of documents. For example, an SQL-like command might be

```
SELECT DOCUMENT#
FROM JOURNAL-DB
WHERE   (DATE >= 1994 AND DATE <= 1997)
   AND  (LANGUAGE = ENGLISH OR FRENCH)
   AND  (PUBLISHER = ASIS OR HFES OR ACM).
```

SQL has powerful features, but using it requires training (2 to 20 hours), and even then users make frequent errors for many classes of queries (Welty, 1985). Alternatives such as *query-by-example* can help users to formulate simpler queries, such as requesting all English-language ACM articles published during or after 1994:

```
JOURNAL | DOCUMENT# | DATE   | AUTHOR | LANGUAGE | PUBLISHER
-----------------------------------------------------------
        |  P._X     | >=1994 |        | ENGLISH  | ACM
        |           |        |        |          |
```

The full set of Boolean expressions, however, is difficult to express except inside a special *condition box.*

Form-fillin queries can substantially simplify many queries, and, if the user interface permits, some Boolean combinations (usually a conjunction of disjuncts (ORs) within attributes with ANDs between attributes) can be easy to express:

```
JOURNAL DATABASE
  DOCUMENT#:
       DATE: 1994..1997
     AUTHOR:
   LANGUAGE: ENGLISH, FRENCH
  PUBLISHER: ASIS, HFES, ACM
```

Although SQL is a standard, many form-fillin variants for expressing relational database queries have been proposed to aid novice searchers. The diversity is itself an impediment to easy use, but designers assume that users are willing to invest minutes or hours to learn each interface. This assumption is not valid for walk-up kiosks or for web pages offering textual-document library searches, in which users are often invited to type keywords or *natural-language queries* in a box, and to click on a `run` button. This presentation is meant to be appealing, but the computer's capacity for responding to the natural-language query is often limited to eliminating frequent terms or commands ("please list the documents that deal with") and searching for remaining words. A ranked list of documents is usually presented, and users must do their best in choosing relevant items from the list.

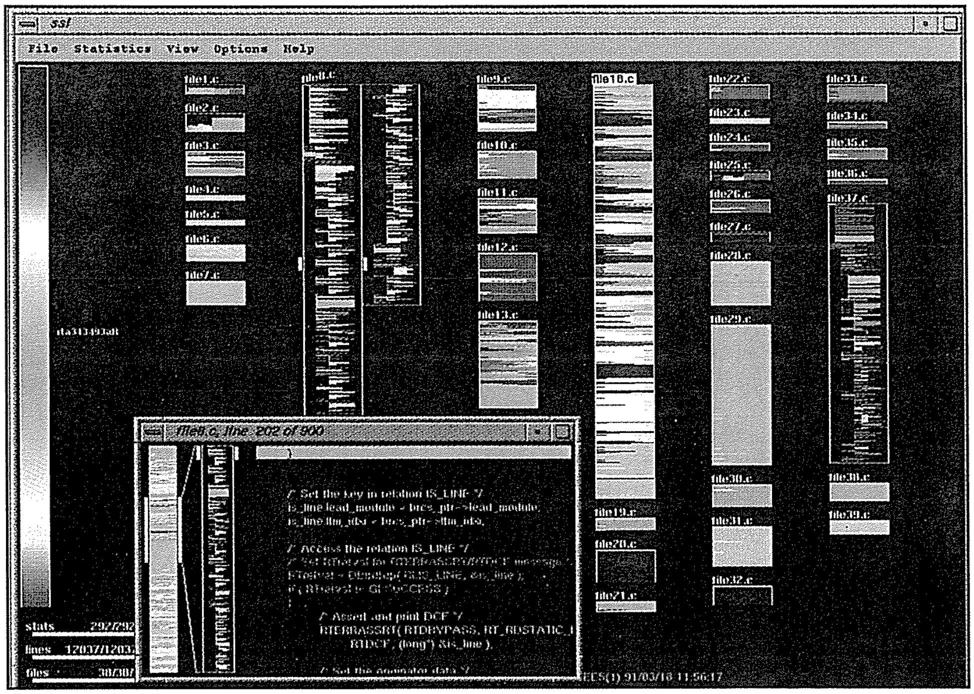

Plate B1: A computer program with 4000 lines of code. The newest lines are in red; the oldest are in blue. The smaller browser window shows a code overview and detail view. (Used with permission of ATT Bell Labs, Naperville, IL.)

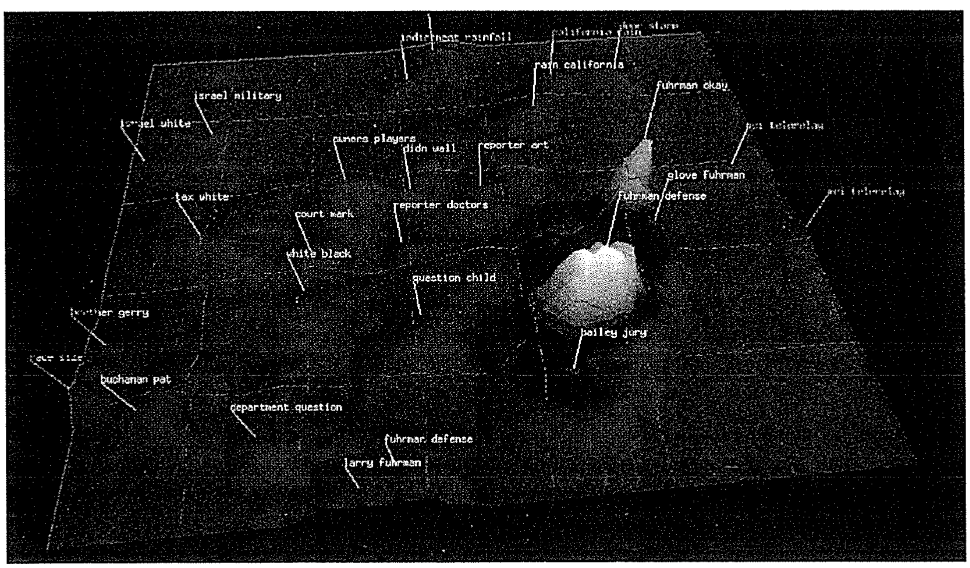

Plate B2: Information-retrieval themescape, showing a multidimensional information space pressed down into a two-dimensional topological map. Some clustering of points can be interesting, but they carry the danger of misinterpretations of the meaning of adjacency. (Wise et al., 1995.) (Used with permission of Battelle Pacific Northwest National Laboratory, Richland, WA.)

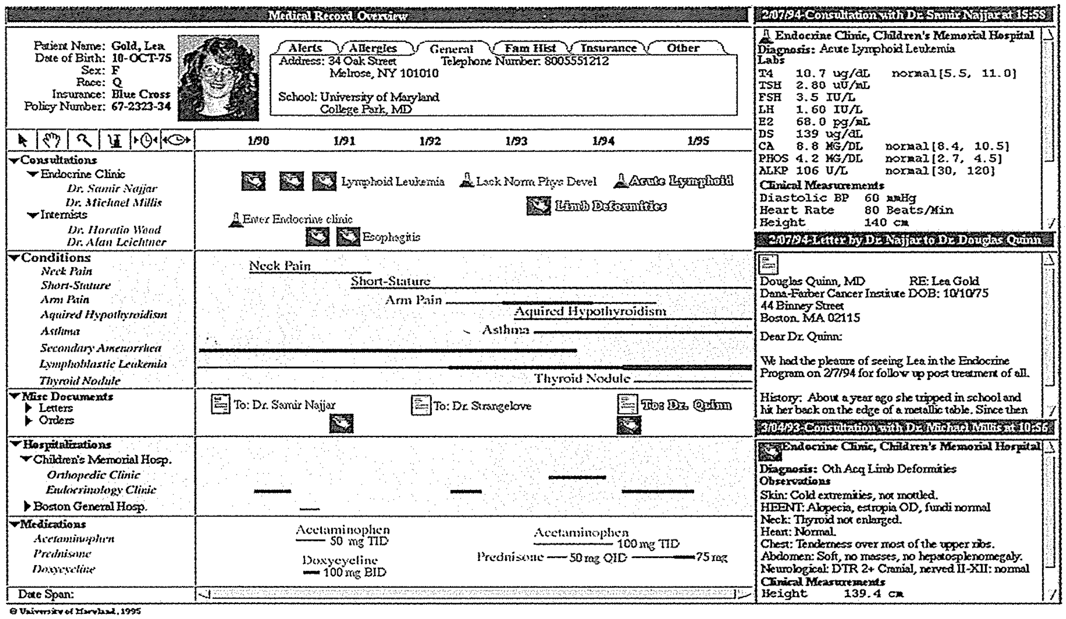

Plate B3: Medical version of LifeLines. Physician visits, conditions, hospitalizations, and medications are shown. Every item in the record is seen as a line or icon in the overview, with color coding by doctor. Line thickness indicates severity and dosage. Windows on the right side show the details. (Plaisant et al., 1997.)

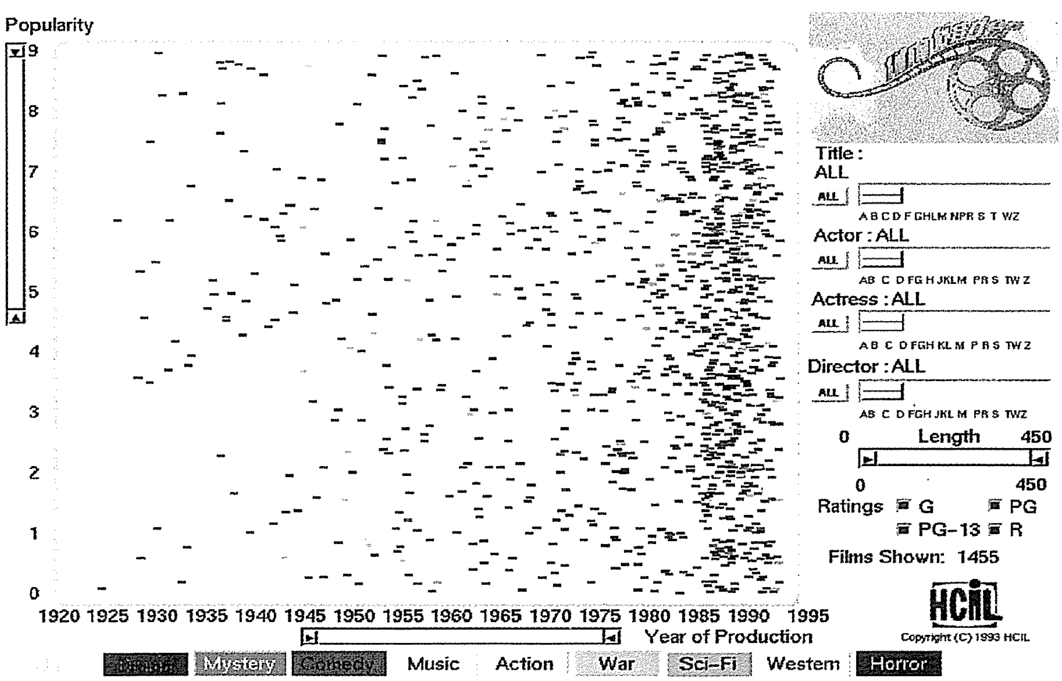

Plate B4 (a): FilmFinder showing 1500 films in a starfield display, where the location of each point is determined by the year of the film (x axis) and by the film's popularity in video store rentals (y axis). The color encodes the film type.

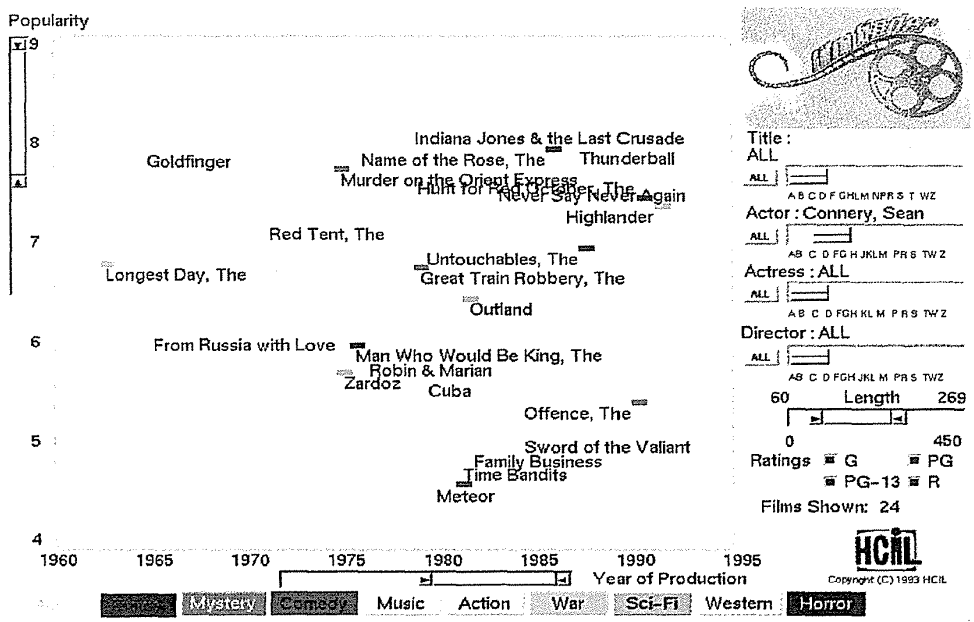

Plate B4 (b): FilmFinder after zooming in on recent popular films. When less than 25 films remain, the titles appear automatically.

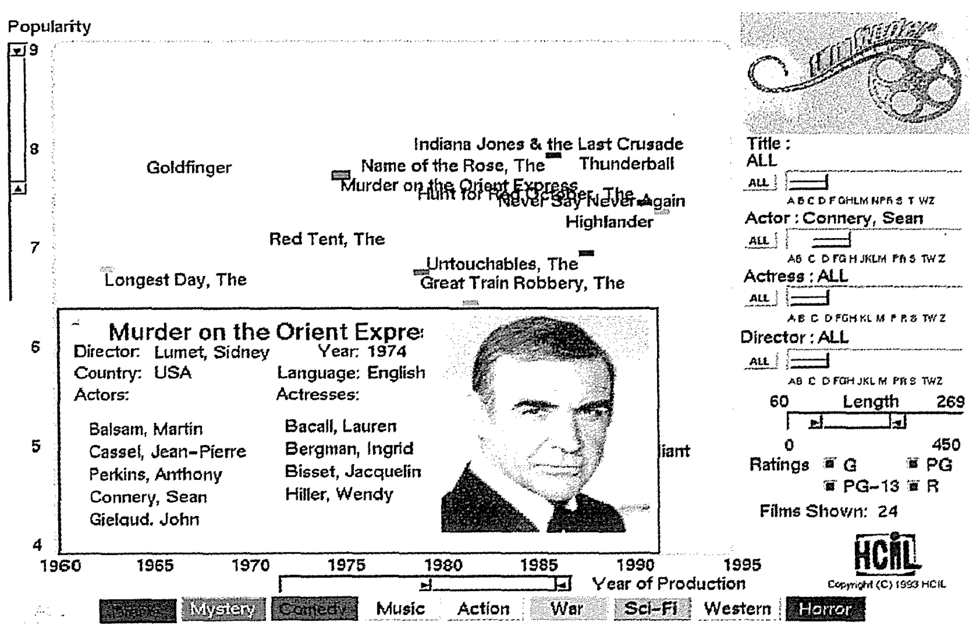

Plate B4 (c): FilmFinder after selection of a single film. The info card pops up with details on demand. (Ahlberg and Shneiderman, 1997.)

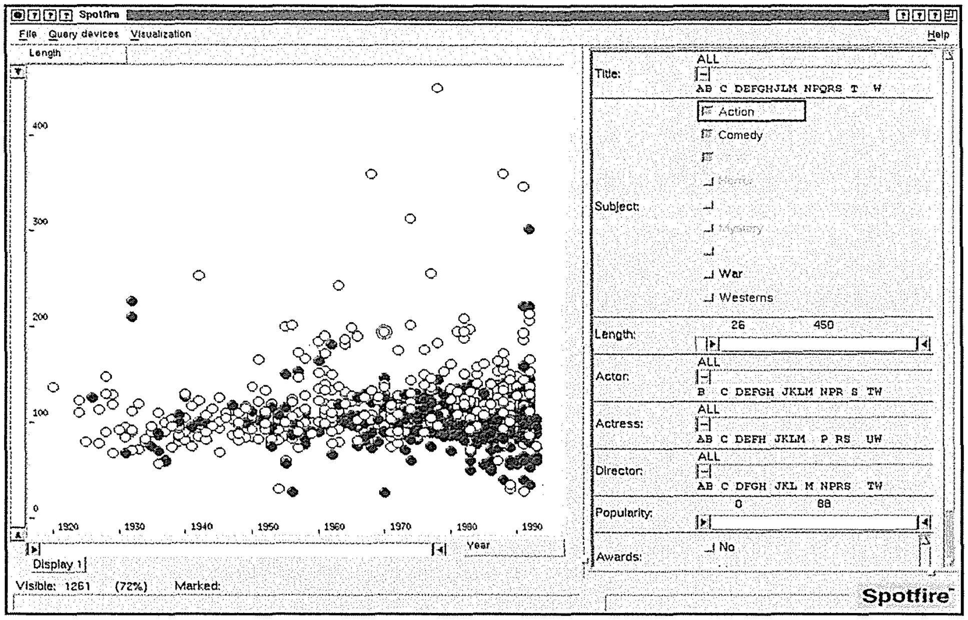

Plate B5: Spotfire version of FilmFinder, which provides increased user controls. Users can set axes (set to length in minutes and year) and glyph attributes (color is set to subject, and larger size indicates award-winning film). (Used with permission of IVEE Development, Goteborg, Sweden.) (http://www.ivee.com)

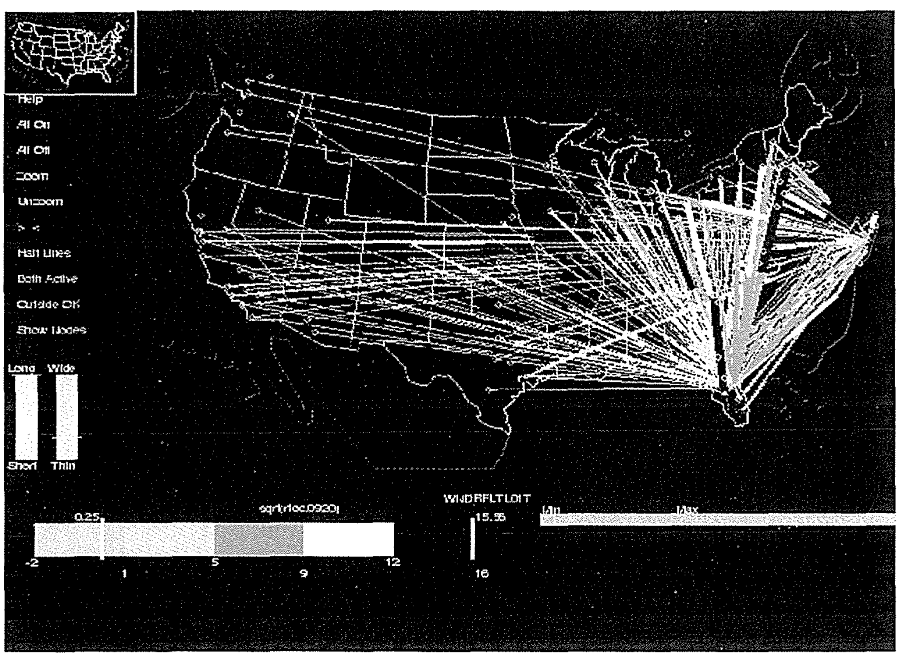

Plate B6: Telephone network traffic represented by thickness and color of the half-line segments between cities. (Used with permission of ATT Bell Labs, Naperville, IL.)

Although an easy-to-use interface is a good idea, if users cannot express their intentions or are uncertain about the meaning of the results, then the interface may need improvement. Finding a way to provide powerful search without overwhelming novice users is a current challenge. Existing interfaces often hide important aspects of the search (by poor design or to protect proprietary relevance-ranking schemes), or make query specification so difficult and confusing that they discourage use. Evidence from empirical studies shows that users perform better and have higher subjective satisfaction when they can view and control the search (Koenemann and Belkin, 1996).

An analogy to the evolution of automobile user interfaces might clarify the goals. Early competitors offered a profusion of controls, and each manufacturer had a distinct design. Some designs—such as having a brake pedal that was far from the gas pedal—were dangerous. Furthermore, if you were accustomed to driving a car with the brake to the left of the gas pedal, and your neighbor's car had the reverse design, it might be risky to trade cars. It took a half-century to achieve good design and appropriate consistency in automobiles; let's hope that we can make the transition faster for text-search user interfaces.

Improved designs and consistency across multiple systems can bring faster performance, reduce mistaken assumptions, and increase success in finding relevant items. For example, with the variety of web search systems, such as Lycos, Infoseek, and AltaVista, users might expect that the search string `direct manipulation` would produce one of the following:

- Search on the exact string `direct manipulation`
- Probabilistic search for `direct` and `manipulation`
- Probabilistic search for `direct` and `manipulation`, with some weighting if the terms are in close proximity
- Boolean search on `direct` AND `manipulation`
- Boolean search on `direct` OR `manipulation`
- Error message indicating missing AND/OR operator or other delimiters

In many systems, there is little or no indication regarding which interpretation has been chosen and whether stemming, case matching, stop words, or other transformations are being applied. Often, the results are displayed in a relevance-ranked manner that is a mystery to many users (and sometimes is a proprietary secret).

To coordinate design practice, we might use a *four-phase framework* to satisfy the needs of first-time, intermittent, and frequent users who are accessing a variety of textual and multimedia libraries (Shneiderman et al., 1997). Finding common ground will be difficult; not finding it will be tragic. Although early adopters of technology are willing to overcome difficulties, the middle and late adopters are not so tolerant. The future of search services on the World Wide Web and elsewhere may depend on the degree to which user frustration and confusion are reduced, while the ability to find reliably sought items in the surging sea of information is increased.

The four-phase framework (Box 15.1) gives great freedom to designers to offer features in an orderly and consistent manner. The phases are

1. *Formulation:* expressing the search
2. *Initiation of action:* launching the search
3. *Review of results:* reading messages and outcomes
4. *Refinement:* formulating the next step

Formulation includes the *source* of the information, the *fields* for limiting the source, the *phrases*, and the *variants*. Even if technically and economically feasible, searching all libraries or collections in a library is not always the preferred approach. Users often prefer to limit the sources to a specific library, collection in a library, or subcollection range of items (users may choose date ranges, languages, media types, publishers, and so on). Users may wish to limit their search to specific *fields* (for example, the title, abstract, or full text of a scientific article) of items within a collection. Typically, users searching on common phrases would prefer to retrieve only those documents whose title contains those phrases. Sources may also be restricted by structured fields (year of publication, volume number, and so on).

In textual databases, users often seek items that contain meaningful *phrases* (`Civil War`, `Environmental Protection Agency`, `George Washington`, `air pollution`, `carbon monoxide`), and multiple entry windows should be provided to allow for multiple phrases. Searches on phrases have proved to be more accurate than are searches on words. Since some relevant items may be missed by a phrase approach, users should have the option to expand a search by breaking the phrases into separate words. Phrases also facilitate searching on names (for example, search on `George Washington` should not turn up George Bush or Washington, D.C.). If Boolean operations, proximity restrictions, or other combining strategies are specifiable, then the users should be able to express them. Users or service providers should have control over stop lists (common words, single letters, obscenities).

When users are unsure of the exact value of the field (subject term, or spelling or capitalization of a city name), they may want to relax the search constraints by allowing *variants* to be accepted. In structured databases, the variants may include a wider range on a numeric attribute. In a textual-document search, interfaces should allow user control over variant capitalization (case sensitivity), stemmed versions (the keyword `teach` retrieves variant suffixes such as teacher, teaching, or teaches), partial matches (the keyword `biology` retrieves sociobiology and astrobiology), phonetic variants from soundex methods (the keyword `Johnson` retrieves Jonson, Jansen, Johnsson), synonyms (the keyword `cancer` retrieves oncology), abbrevia-

Box 15.1

Four-phase framework to clarify user interfaces for textual search.

1. *Formulation*
   - Search the appropriate sources in libraries and collections
   - Use *fields* for limiting the source: structured fields such as year, media, or language, and text fields such as titles or abstracts of documents
   - Recognize *phrases* to allow entry of names such as `George Washington` or `Environmental Protection Agency`, and concepts such as `abortion rights reform` or `gallium arsenide`.
   - Permit *variants* to allow relaxation of search constraints, such as case sensitivity, stemming, partial matches, phonetic variations, abbreviations, or synonyms from a thesaurus.
2. *Action*
   - Include *explicit actions* initiated by buttons with consistent labels (such as "Search"), location, size, and color.
   - Include *implicit actions* initiated by changes to a parameter of the formulation phase that immediately produce a new set of search results.
3. *Results*
   - Read explanatory messages.
   - View textual lists.
   - Manipulate visualizations.
   - Control what the size of the result set is and which fields are displayed.
   - Change sequencing (alphabetical, chronological, relevance ranked, etc.).
   - Explore clustering (by attribute value, topics, etc.).
4. *Refinement*
   - Use meaningful messages to guide users in progressive refinement; for example, if the two words in a phrase are not found near each other, then offer easy selection of individual words or variants.
   - Make changing of search parameters convenient.
   - Allow search results and the setting of each parameter to be saved, sent by email, or used as input to other programs, such as visualization or statistical tools.

tions (the keyword `IBM` retrieves International Business Machines, and vice versa), and broader or narrower terms from a thesaurus (the keyphrase `New England` retrieves Vermont, Maine, Rhode Island, New Hampshire, Massachusetts, and Connecticut).

The second phase is the initiation of *action,* which may be explicit or implicit. Most current systems have a search button for explicit initiation, or for delayed or regularly scheduled initiation. The button label, size, and color should be consistent across versions. An appealing alternative is *implicit initiation,* in which each change to a component of the formulation phase immediately produces a new set of search results. *Dynamic queries*—in which users adjust query widgets to produce continuous updates—have proved to be effective and satisfying. They require adequate screen space and rapid processing, but their advantages are great.

The third phase is the review of *results,* in which the users read messages, view textual lists, or manipulate visualizations. Users may be given control over what the size of the result set is, which fields are displayed, how results are sequenced (alphabetical, chronological, relevance ranked), and how results are clustered (by attribute value, by topics) (Pirolli et al., 1996).

The fourth phase is the *refinement.* Search interfaces should provide meaningful messages to explain search outcomes and to support progressive refinement. For example, if a stop word, obscenity, or misspelling is eliminated from a search input window, or if stemmed terms, partial matches, or variant capitalizations are included, users should be able to see these changes to their query. If two words in a keyphrase are not found proximally, then feedback should be given about the occurrence of the words individually. If multiple phrases are input, then items containing all phrases should be shown first and identified, followed by items containing subsets; but if no documents are found with all phrases, that failure should be indicated. There is a fairly elaborate decision tree (maybe 60 to 100 branches) of search outcomes and messages that needs to be specified. Another aspect of feedback is that, as searches are made, the system should keep track of them in a *history buffer* to allow review of earlier searches. Progressive refinement, in which the results of a search are refined by changing of the search parameters, should be convenient. Search results and the settings of all parameters should be objects that can be saved, sent by electronic mail, or used as input to other programs—for example, for visualization or statistical tools.

The four-phase framework can be applied by designers to make the search process more visible, comprehensible, and controllable by users. This approach is in harmony with movement toward direct manipulation, in which the state of the system is made visible and is placed under user control. Novices may not want to see all the components of the four phases initially, but, if they are unhappy with the search results, they should be able to view the settings and change their queries easily. A revised interface for the Library of Congress' THOMAS system (Fig. 15.1) shows how the framework might be applied to full-text searching of proposed legislation.

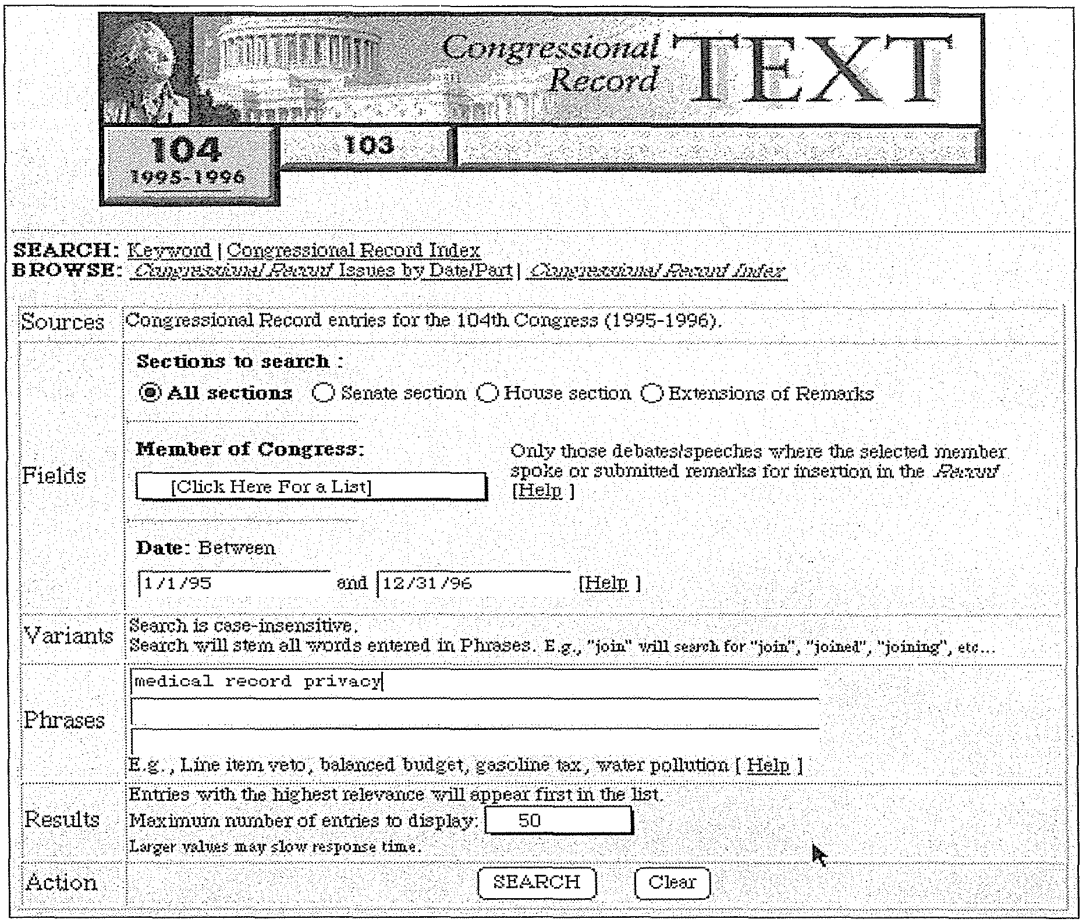

Figure 15.1

A revised interface for the Library of Congress' THOMAS system. The display shows how the four-phase framework might be applied to text searching on Congressional Record articles. Implemented by Bryan Slavin at the University of Maryland Human–Computer Interaction Laboratory. (Shneiderman et al., 1997.)

15.3 Multimedia Document Searches

Interfaces to search structured databases and textual-document libraries are good and getting better, but searching in multimedia document libraries is still in a primitive stage. Current approaches to locating images, videos, sound, or animation depend on a parallel database or document search to locate the items. For example, searches in photo libraries can be done by date, photographer, medium, location, or text in captions, but finding photos showing a ribbon-cutting ceremony or videos of a sunset is difficult. In the near term, those people who must search multimedia documents should push for ambitious captioning and attribute recording. Classification according to useful search categories (agriculture, music, sports, personalities) is helpful, although costly and imperfect.

Recent advances in computer algorithms may enable greater flexibility in locating items in multimedia libraries. User-interface designs to specify the permissible matches are varied. Some systems have elaborate textual commands, but most are moving toward graphical specification of query components:

- *Photo search* Finding photos with images such as the Statue of Liberty is a substantial challenge for image-analysis researchers, who describe this task as *query by image content (QBIC).* Lady Liberty's distinctive profile might be identifiable if the orientation, lens focal length, and lighting were held constant, but the general problem is difficult in large and diverse collections of photos. Two promising approaches are to search for distinctive features such as the torch or the seven spikes in the crown, or to search for distinctive colors, such as the faded green copper verdigris. Users can specify features or color patterns with standard drawing tools, and even can indicate where in the image to search. For example, users could specify red, white, and blue in the upper third of an image to look for an American flag flying above a building. Of course, separating out the British, French, or other flags is not easy.

  More success is attainable with restricted collections, such as of glass vases, for which users could draw a desired profile and retrieve vases with long narrow necks. Other candidate collections include photos of constellations, subatomic particle tracks, or red blood cells. Users could specify their requests by selecting from a set of templates and adjusting the templates to describe their query. For critical applications, such as fingerprint matching, current successes depend on human identification of as many as 20 distinct features, but automatic recognition is improving. Even if completely automatic recognition is not possible, it will still be useful to have computers perform filtering, such as finding all the portraits with neutral backgrounds in a photo library.

- *Map search* Computer-generated maps are increasingly available online. Locating a map by latitude and longitude is the structured-database solution, but search by features is becoming possible because the tools used to build maps preserve the structural aspects and the multiple layers in maps. For example, users might specify a search for all port cities with a population greater than 1 million and an airport within 10 miles. Search on simpler maps such as airline routes might find flights to a given destination with no more than two connections on the same airline. Another candidate is weather maps, in which structured data—such as temperature, winds, or barometric pressure—make the search specification convenient.

- *Design or diagram search* Some computer-assisted design packages offer users limited search capabilities within a single design or across design collections. Finding red circles inside blue squares may help in some cases, but more elaborate strategies for finding engine designs with pistons smaller than 6 centimeters could prove more beneficial. Diagramming tools for making flowcharts or organization charts can add search capabilities to locate organizations that have more than five levels of management or situations where vice presidents are managing more than seven projects. Newspaper-layout packages could allow search for all occasions of headlines using fonts larger than 48 points, or headlines that span the front page.
- *Sound search* Imagine a music database system that would respond when users hum a few notes by producing a list of symphonies that contain that string of notes. Then, with a single touch, users could listen to the full symphony. Implementing this idea in the unstructured world of analog-encoded or even digitally encoded music is difficult, but imagine that the score sheets of symphonies were stored with the music and that string search over the score sheets was possible. Then, the application becomes easier to conceive. Identification of the users' hummed input might not be reliable, but if visual feedback were provided or if users entered the notes on a staff, then the fantasy would become feasible. Finding a spoken word or phrase in databases of telephone conversations is still difficult, but is becoming possible, even on a speaker-independent basis.
- *Video search* Searching a video or film involves more than simply searching through each of the frames. Users may wish to have a video segmented into scenes or cuts, and to identify zooming in or out and panning left or right. Gaining an overview of a 2-hour video by a time line of scenes would enable better understanding, editing, or selection. Combinations of structured databases and textual documents with video libraries lead to powerful services. Television news or sports libraries maintain structured databases and textual documents to support search for presidential appearances, disasters, or football highlights, carefully indexed for rapid future retrieval.
- *Animation search* Animation-authoring tools are still in early stages of development, but it might be possible to specify searches for certain kinds of animation—for example, spinning globes, moving banners, bouncing balls, or morphing faces. Although it might be less useful, it should be relatively easy to search for slides in a presentation that have moving text that comes in from the left, or in which the transition from one slide to another is by a barndoor animation.

15.4 Information Visualization

Grasping the whole is a gigantic theme. Arguably, intellectual history's most important. Ant-vision is humanity's usual fate; but seeing the whole is every thinking person's aspiration.

David Gelernter, *Mirror Worlds*, 1992

Visualization is a method of computing. It transforms the symbolic into the geometric, enabling researchers to observe their simulations and computations. Visualization offers a method for seeing the unseen. It enriches the process of scientific discovery and fosters profound and unexpected insights. In many fields it is already revolutionizing the way scientists do science.

McCormick et al., 1987

The success of direct-manipulation interfaces is indicative of the power of using computers in a more visual or graphic manner. A picture is often said to be worth a thousand words and, for some tasks, a visual presentation—such as a map or photograph—is dramatically easier to use or comprehend than is a textual description or a spoken report. As computer speeds and display resolution increase, information visualization and graphical interfaces are likely to have an expanding role. If a map of the United States is displayed, then it should be possible to point rapidly at one of 1000 cities to get tourist information. Of course, a foreigner who knows a city's name (for example, New Orleans), but does not know its location, may do better with a scrolling alphabetical list. Visual displays become even more attractive to provide orientation or context, to enable selection of regions, and to provide dynamic feedback for identifying changes (for example, on a weather map). Scientific visualization has the power to make visible and comprehensible atomic, cosmic, and common three-dimensional phenomena (such as heat conduction in engines, airflow over wings, or ozone holes). Abstract-information visualization has the power to reveal patterns, clusters, gaps, or outliers in statistical data, stock-market trades, computer directories, or document collections.

Overall, the bandwidth of information presentation is potentially higher in the visual domain than it is for media reaching any of the other senses. Humans have remarkable perceptual abilities that are greatly underutilized in current designs. Users can scan, recognize, and recall images rapidly, and can detect subtle changes in size, color, shape, movement, or texture. They can point to a single pixel, even in a megapixel display, and can drag one object to another to perform an action. User interfaces thus far have been largely text oriented, so as visual approaches are explored, appealing new opportunities are emerging.

There are many visual design guidelines. The central principle might be summarized as this *visual-information-seeking mantra:*

Overview first, zoom and filter, then details on demand
Overview first, zoom and filter, then details on demand
Overview first, zoom and filter, then details on demand
Overview first, zoom and filter, then details on demand
Overview first, zoom and filter, then details on demand
Overview first, zoom and filter, then details on demand
Overview first, zoom and filter, then details on demand
Overview first, zoom and filter, then details on demand
Overview first, zoom and filter, then details on demand
Overview first, zoom and filter, then details on demand
Overview first, zoom and filter, then details on demand
Overview first, zoom and filter, then details on demand

Each line represents one project in which I found myself rediscovering this principle and therefore wrote it down as a reminder. The mantra proved to be a good starting point when I was trying to characterize the multiple information-visualization innovations occurring at university, government, and industry research laboratories. To sort out the numerous prototypes and to guide researchers to new opportunities, Box 15.2 gives a *data type by task taxonomy (TTT)* of information visualizations.

As in the case of search, users are assumed to be viewing collections of items, where items have multiple attributes. In all seven data types (one-, two-, three-dimensional data; temporal and multi-dimensional data; and tree and network data) the items have one or more attributes. A basic search task is to select all items that match target attributes—for example, find all divisions in an company that have a budget greater than $500,000.

The data types of the TTT characterize the task-domain information objects and are organized by the problems that users are trying to solve. For example, in two-dimensional information such as maps, users are trying to grasp adjacency or to navigate paths, whereas in tree-structured information, users are trying to understand parent–child–sibling relationships. The tasks in the TTT are task-domain information actions that users wish to perform.

The seven tasks are at a high level of abstraction. Refinements and additions to these tasks would be natural next steps in expanding this taxonomy. The seven tasks are overview, zoom, filter, details-on-demand, relate, history, extract. Further discussion of the seven tasks follows the descriptions of the seven data types.

Box 15.2

Data Type by Task Taxonomy (TTT) to identify visualization data types and the tasks that need to be supported.

Data Type by Task Taxonomy (TTT)

Data Types

1-D Linear	Document Lens, SeeSoft Information Mural
2-D Map	GIS, Arcinfo, ThemeMap LyberWorld, InfoCrystal
3-D World	Desktops, WebBook, VRML CAD, Medical, Molecules
Temporal	Perspective Wall, ESDA MSProjects, LifeLines
Multi-Dimensional	Parallel Coordinates, Starfield, Visage Influence Explorer, TableLens
Tree	Outliners, Superbook, FileManager Cone/Cam/Hyperbolic, TreeBrowser, Treemaps
Network	Netmap, SemNet, SeeNet, Butterfly

Tasks

Overview	Gain an overview of the entire collection.
Zoom	Zoom in on items of interest.
Filter	Filter out uninteresting items.
Details-on-demand	Select an item or group and get details when needed.
Relate	View relationships among items.
History	Keep a history of actions to support undo, replay, and progressive refinement.
Extract	Allow extraction of subcollections and of the query parameters.

1-D Linear Data

Linear data types include textual documents, program source code, and alphabetical lists of names, all of which are all organized in a sequential manner. Each item in the collection is a line of text containing a string of characters. Additional attributes might be the date of most recent update or author name. Interface-design issues include what fonts, color, size to use, and what overview, scrolling, or selection methods can be used. User tasks might be to find the number of items, to see items having certain attributes (show only lines of a document that are section titles, lines of a program that were changed from the previous version, or people in a list who are older than 21 years), or to see an item with all its attributes.

An early approach to dealing with large one-dimensional data sets was the *bifocal display,* which provided detailed information in the focus area and less information in the surrounding context area (Spence and Apperley, 1982). A selected issue of a scientific journal had details about each article; the older and newer issues of the journal were to the left and right on the bookshelf with decreasing space. Another effort to visualize one-dimensional data showed the attribute values of thousands of items in a fixed-sized space using a scrollbar-like display called *value bars* (Fig. 15.2) (Chimera, 1992). Even greater compressions were accomplished in compact displays of tens of thousands of lines of program source code in See Soft (Color Plate B1) (Eick et al., 1992) or lines in Hamlet (Fig. 15.3). Other examples of one-dimensional data include large textual documents in Document Lens (Fig. 15.4) (Robertson and Mackinlay, 1993) and historical

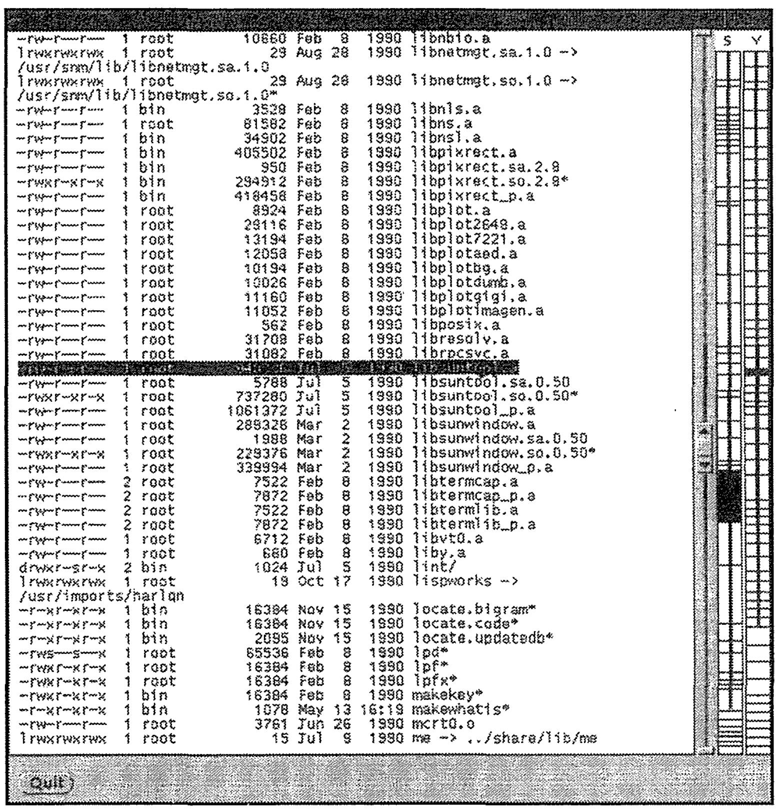

Figure 15.2

Each value bar shows one attribute of the linear list of items. In this Unix directory example, the two value bars on the right represent the file size (S) and file modification recency, or youth (Y). The currently selected file is one of the biggest in size and is moderately youthful. (Chimera, 1992.)

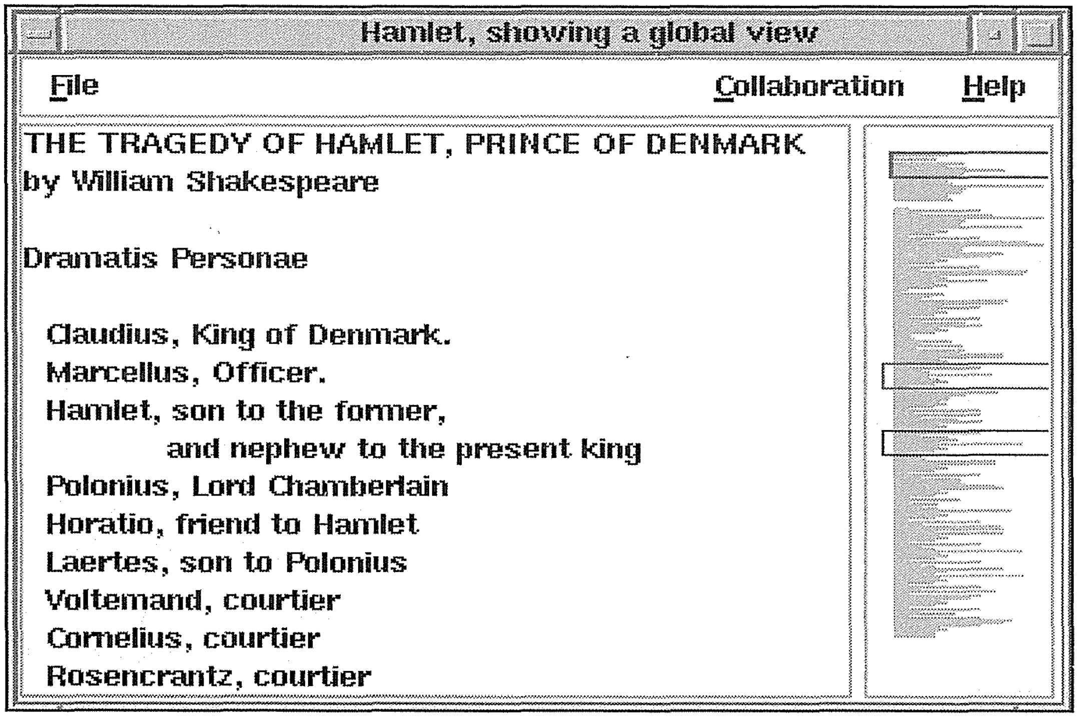

Figure 15.3

Shakespeare's Hamlet, viewed with a one-dimensional overview on the right side showing where three users are reading the document. Each person's field-of-view box shows their location. This user is at the start. (Used with permission of the University of Calgary, Alberta, Canada.)

data about sunspots using the information-mural algorithms (Fig. 15.5) (Jerding and Stasko, 1995).

2-D Map data

Planar or map data include geographic maps, floorplans, and newspaper layouts. Each item in the collection covers some part of the total area and may or may not be rectangular. Each item has task-domain attributes, such as name, owner, and value, and interface-domain features, such as size, color, and opacity. Many systems adopt a multiple-layer approach to dealing with map data, but each layer is two-dimensional. User tasks are to find adjacent items, containing items and paths between items, and to perform the seven basic tasks.

Examples include geographic-information systems, which are a large research and commercial domain (Laurini and Thompson, 1992; Egenhofer and Richards, 1993) with numerous systems available (see Fig. 6.5). Information-visualization researchers have used spatial dis-

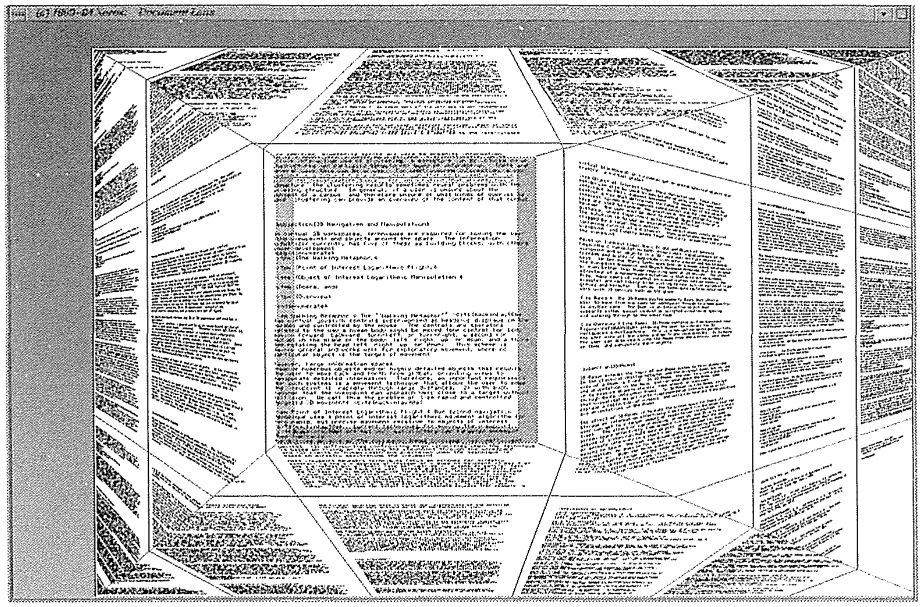

Figure 15.4

Document Lens showing many pages of a document in miniature form. Users can zoom in on any page easily and quickly. (Used with permission from Xerox PARC, Palo Alto, CA.)

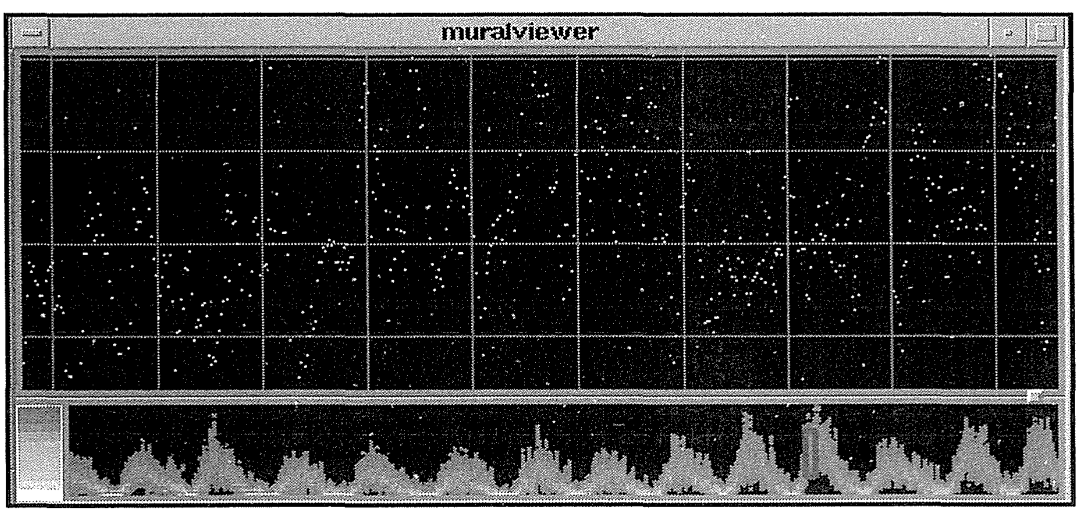

Figure 15.5

The information-mural overview at the bottom uses an antialiasing algorithm to show 52,000 readings of sun spots from 1850 to 1993. The field-of-view box at the bottom shows the context for the detail view on top. (Jerding & Stasko, 1995.) (Used with permission of Georgia Tech University, Atlanta, GA.)

plays of document collections (Color Plate B2) (Korfhage, 1991; Hemmje et al., 1993; Wise et al., 1995) organized proximally by term co-occurrences.

3-D World

Real-world objects such as molecules, the human body, and buildings have items with volume and with potentially complex relationships with other items. Computer-assisted design systems for architects, solid modelers, and mechanical engineers are built to handle complex three-dimensional relationships. Users' tasks deal with adjacency plus above–below and inside–outside relationships, as well as the seven basic tasks. In three-dimensional applications, users must cope with their position and orientation when viewing the objects, plus must handle the serious problems of *occlusion*. Solutions are proposed in many prototypes with techniques such as overviews, landmarks, perspective, stereo display, transparency, and color coding.

Examples of three-dimensional computer graphics and computer-assisted design are numerous, but information-visualization work in three dimensions is still novel. Some virtual-environment researchers have sought to present information in three-dimensional structures (see Section 6.8). Navigating high-resolution images of the human body is the challenge in the National Library of Medicine's Visible Human project (Fig. 15.6) (North et al., 1996). Architectural walkthroughs or flythroughs can give users an idea of what a finished building will look like. A three-dimensional desktop is thought to be appealing to users, but disorientation, navigation, and hidden data problems remain (Fig. 15.7) (Card et al., 1996).

Temporal data

Time lines are widely used and are sufficiently vital for medical records, project management, or historical presentations that researchers have created a data type that is separate from one-dimensional data. The distinctions of *temporal data* are that items have a start and finish time, and that items may overlap. Frequent tasks include finding all events before, after, or during some time period or moment, plus the seven basic tasks.

Many project-management tools exist; novel visualizations of time include the perspective wall (Fig. 15.8) (Robertson et al., 1993) and LifeLines (see Fig. 1.5 and Color Plate B3) (Plaisant et al., 1996). LifeLines shows a youth's history keyed to the needs of the Maryland Department of Juvenile Justice, but is intended to present medical patient histories as a compact overview with selectable items that allow users to get details-on-demand. Temporal-data visualizations appear in systems for editing video data, composing music, or preparing animations, such as Macromedia Director.

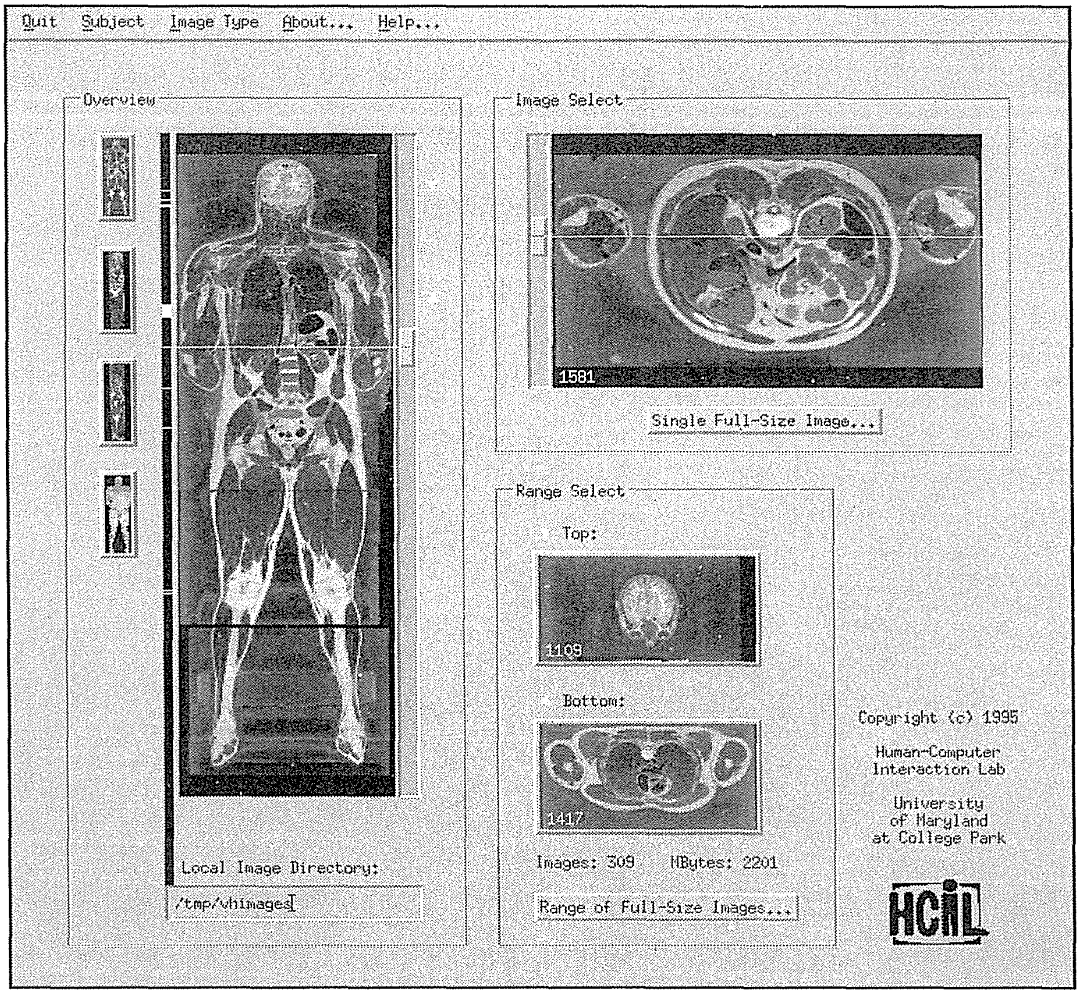

Figure 15.6

Visible Human Explorer user interface, showing a reconstructed coronal section overview (on the left) and an axial preview image of the upper abdominal region (on the upper right). Dragging the sliders animates the cross-sections through the body. (North et al., 1996.) (Available at http://www.nlm.nih.gov)

Multidimensional data

Most relational- and statistical-database contents are conveniently manipulated as multidimensional data, in which items with *n* attributes become points in a *n-dimensional space.* The interface representation can be dynamic two-dimensional scattergrams, with each additional dimension controlled by a slider (Ahlberg and Shneiderman, 1994). Buttons can used

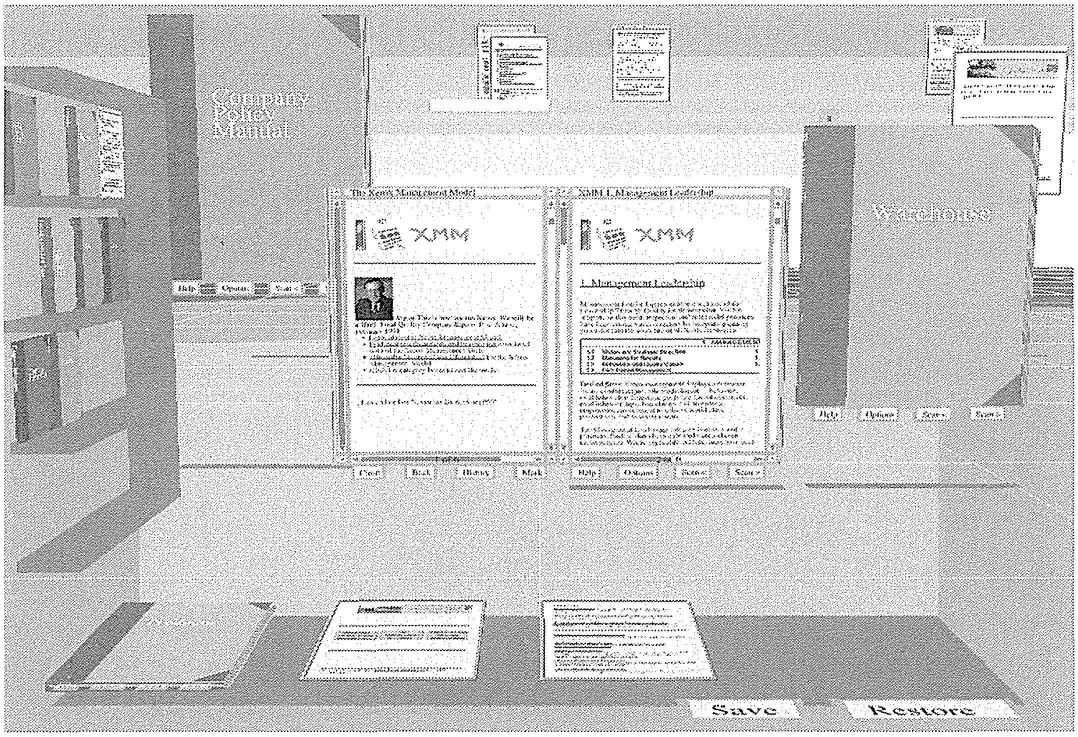

Figure 15.7

WebBook and WebForager. These three-dimensional worlds are used for browsing and recording web pages. (Used with permission from Xerox PARC, Palo Alto, CA.)

for attribute values when the cardinality is small—say, less than 10. Tasks include finding patterns, clusters, correlations among pairs of variables, gaps, and outliers. Multidimensional data can be represented by a three-dimensional scattergram, but disorientation (especially if the user's point of view is from inside the cluster of points) and occlusion (especially if close points are represented as being larger) can be problems. The technique of using parallel coordinates (Fig. 15.9) is a clever innovation that makes certain tasks easier, but takes practice for users to comprehend (Inselberg, 1985).

The early HomeFinder (Fig. 15.10) developed dynamic queries and sliders for user-controlled visualization of multidimensional data (Williamson and Shneiderman, 1992). The successor FilmFinder (Color Plate B4a–c) refined the techniques (Ahlberg and Shneiderman, 1994) for starfield displays (zoomable, color-coded, user-controlled scattergrams), and laid the basis for the commercial product Spotfire (Color Plate B5) (Ahlberg and Wistrand, 1995). Extrapolations include the Aggregate Manipulator (Goldstein and Roth, 1994), movable filters (Fishkin and

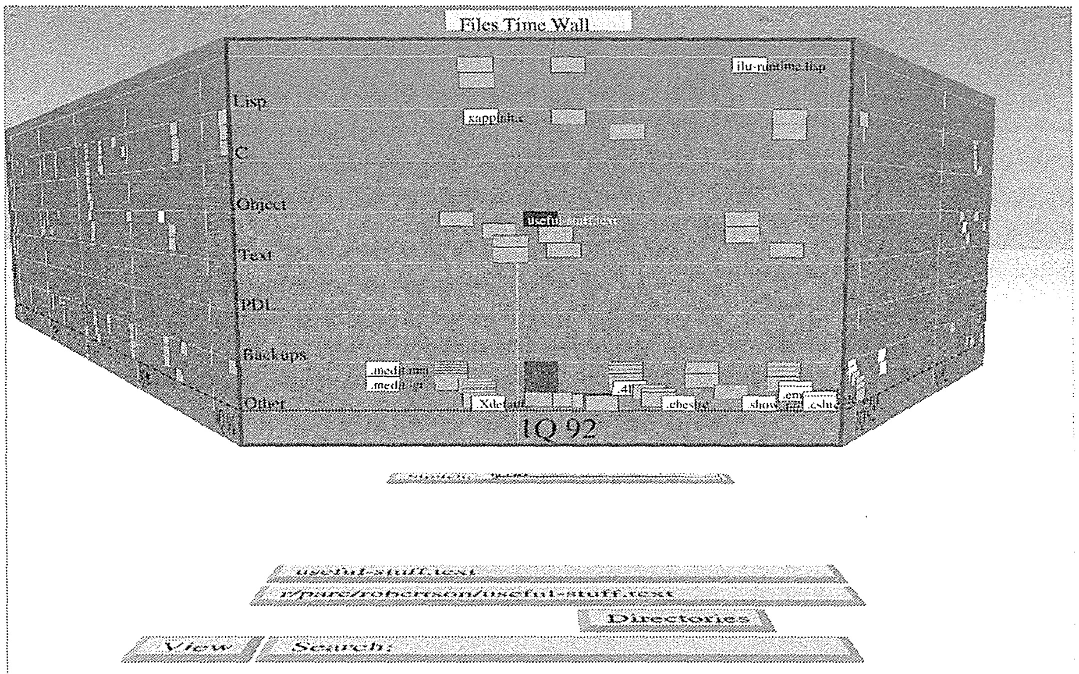

Figure 15.8

A perspective wall, showing time moving from left to right, with the focus in the center. Different categories of programs are shown on each level of the wall. Color or size coding can be used. (Used with permission from Xerox PARC, Palo Alto, CA.)

Stone, 1995), and Selective Dynamic Manipulation (Chuah et al., 1995). Related works include VisDB for multidimensional database visualization (Keim and Kreigal, 1994), the spreadsheet-like Table Lens (Fig. 15.11) (Rao and Card, 1994), and the multiple linked histograms in the Influence Explorer (Tweedie et al., 1996).

Tree data

Hierarchies or tree structures are collections of items, in which each item (except the root) has a link to one parent item. Items and the links between parent and child can have multiple attributes. The basic tasks can be applied to items and links, and tasks related to structural properties become interesting—for example, how many levels are in the tree, or how many children does an item have? While it is possible to have similar items at leaves and internal nodes, it is also common to find different items at each level in a tree. Fixed-level trees, with all leaves equidistant from the root, and fixed-fanout trees, with the same number of children for every parent, are easier to handle. High-fanout (broad) and small-fanout (deep) trees are important special cases. Inter-

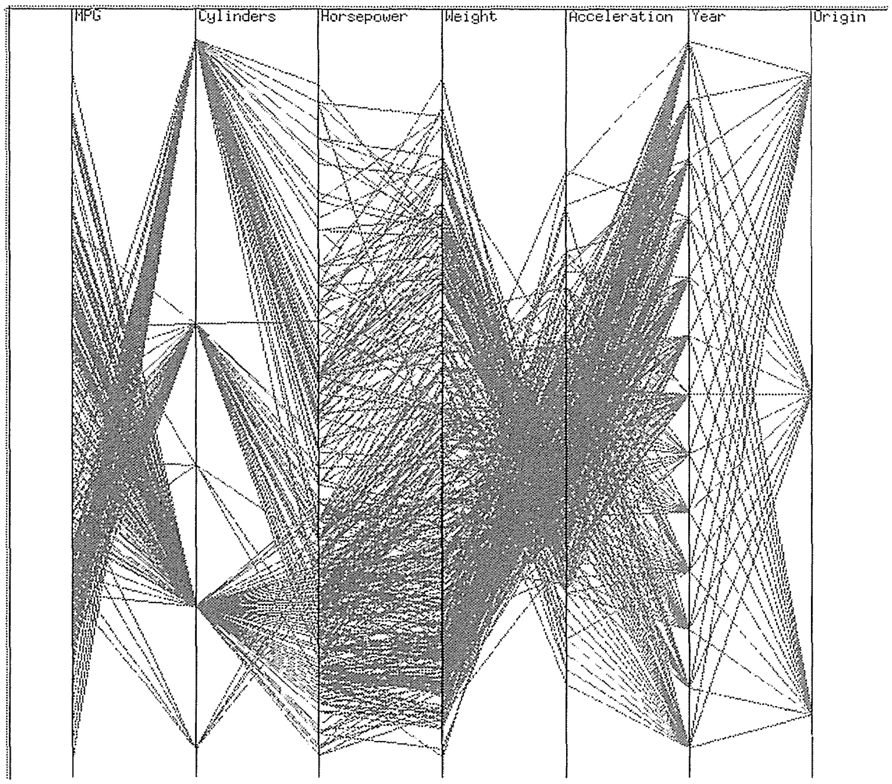

Figure 15.9

Parallel coordinate plot of seven dimensions of automobile data (CARS dataset obtained from StatLib at Carnegie Mellon University). There is a range of MPG (miles per gallon) values, but clear clusters of two-, four-, and six-cylinder cars are visible. More cylinders generally produce higher horsepower. Also notable is the generally inverse relationship of weight to acceleration. (Used with permission of Matt Ward, Worcester Polytechnic Institute, Worcester, MA.)

face representations of trees can use the outline style of indented labels used in tables of contents (Chimera and Shneiderman, 1993); a node-and-link diagram; or a *treemap,* in which child items are rectangles nested inside parent rectangles.

Tree-structured data has long been displayed with indented outlines (Egan et al., 1989), or with connecting lines, as in many computer-directory file managers. Attempts to show large tree structures as node-and-link diagrams in compact forms include the three-dimensional cone (Fig. 15.12) and cam trees (Robertson et al., 1993; Carriere and Kazman, 1995), dynamic pruning in the TreeBrowser (Fig. 15.13) (Kumar et al., 1997), and the appealingly animated hyperbolic trees (Fig. 15.14) (Lamping et al.,

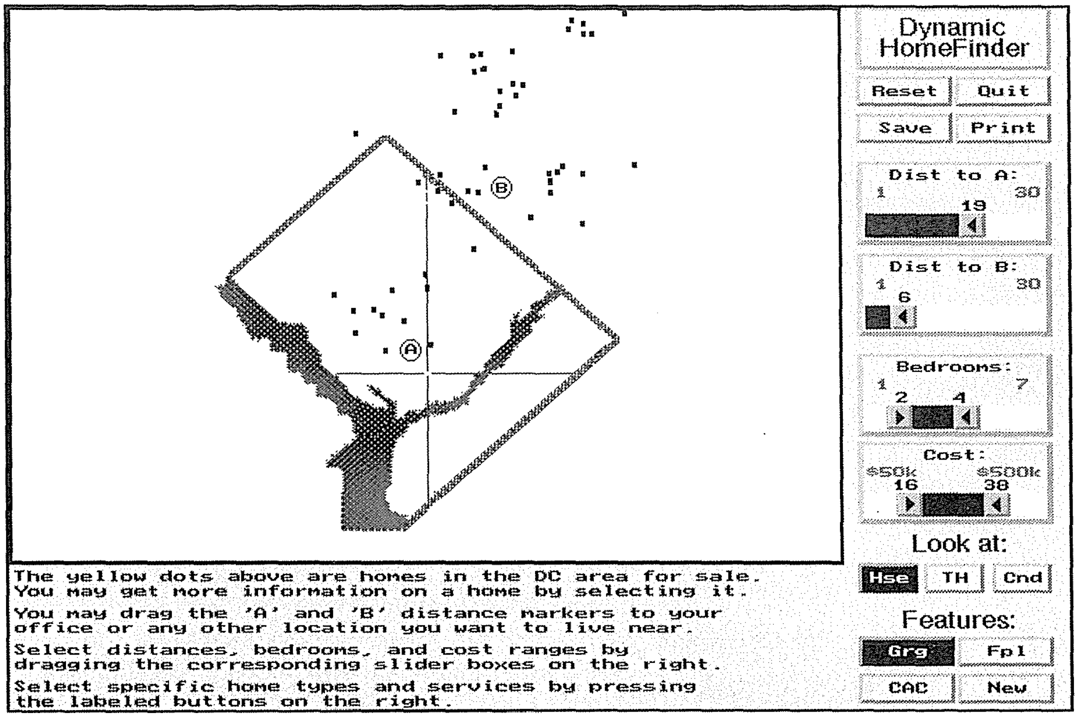

Figure 15.10

Dynamic HomeFinder, an early application of dynamic queries. Homes for sale in the Washington, D.C., area were shown as 1100 points of light. As users moved the sliders on the right, the screen was updated immediately to show the points matching the current query. By clicking on any point, users could get a detailed description. (Williamson and Shneiderman, 1992.)

1995). The space-filling mosaic approach, treemaps, shows an arbitrary-sized tree in a fixed rectangular space (Shneiderman, 1992; Johnson and Shneiderman, 1991). The treemap approach was applied successfully to libraries (Fig. 15.15), computer directories (Fig. 15.16), sales data, business decision making (Asahi et al., 1995), and web browsing (Mitchell et al., 1995; Mukherjea et al., 1995), but first-time users take 10 to 20 minutes to accommodate to treemaps.

Network data

Sometimes, relationships among items cannot be captured conveniently with a tree structure, and it is useful to have items linked to an arbitrary number of other items. Although many special cases of networks exist (acyclic, lattices, rooted versus unrooted, directed versus undirected), it is

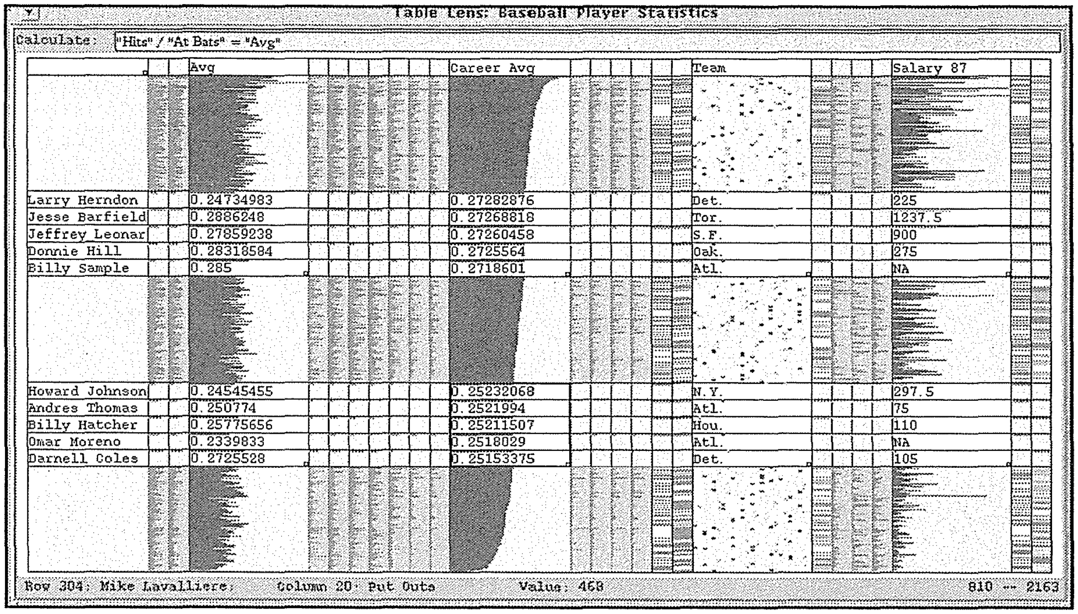

Figure 15.11

Table Lens, a program that provided a spreadsheetlike world that also supported information-visualization methods to find rankings and correlations among baseball players. (Used with permission from Xerox PARC, Palo Alto, CA.)

convenient to consider them all as one data type. In addition to performing the basic tasks applied to items and links, network users often want to know about shortest or least costly paths connecting two items or traversing the entire network. Interface representations include a node-and-link diagram, and a square matrix of the items with the value of a link attribute in the row and column representing a link.

Network visualization is an old but still imperfect art because of the complexity of relationships and user tasks. Commercial packages can handle small networks or simple strategies, such as Netmap's layout of nodes on a circle with links criss-crossing the central area. Specialized visualizations can be designed to be more effective for a given task, such as a network diagram showing heavy telephone traffic on holidays (Color Plate B6). An ambitious three-dimensional approach allowed users to fly into a network and control the visualization (Fairchild et al., 1988). New interest in this topic has been spawned by attempts to visualize the World Wide Web (Andrews, 1995; Hendley et al., 1995).

The seven data types that we have discussed reflect an abstraction of the reality. There are many variations on these themes (two-and-one-half

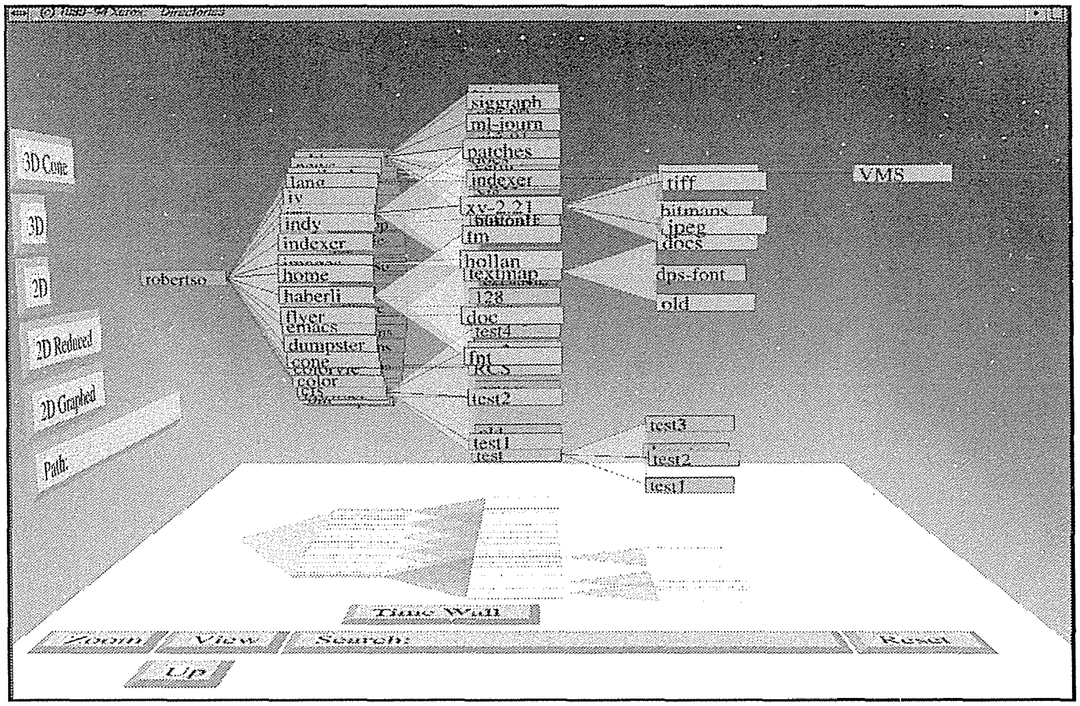

Figure 15.12

Cam-tree design, showing a hierarchical directory from left to right side. Users can rotate the trees smoothly in this three-dimensional viewer. When the root of the tree is shown at the top, this representation is called a cone-tree. (Used with permission from Xerox PARC, Palo Alto, CA.)

or four-dimensional data, multitrees) and many prototypes use combinations of these data types. This taxonomy is useful only if it facilitates discussion and leads to useful discoveries. We can get an idea of missed opportunities by looking at the tasks and data types in depth.

Overview task

We can gain an overview of the entire collection. Overview strategies (Section 13.5) include zoomed-out views of each data type that allow the user to see the entire collection plus an adjoining detail view. The overview contains a movable field-of-view box with which the user controls the contents of the detail view, allowing zoom factors of 3 to 30. Replication of this strategy with intermediate views enables users to reach larger zoom factors. Another popular approach is the fisheye strategy (Furnas, 1986), which has been applied most commonly for network browsing (Fig.

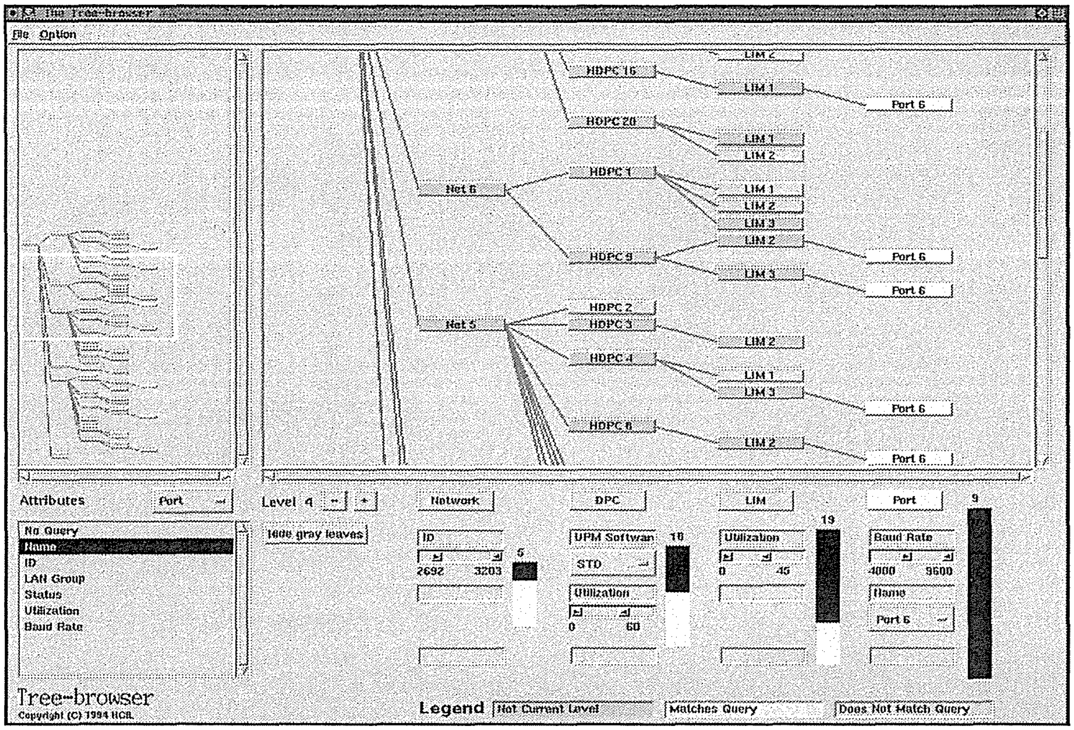

Figure 15.13

PDQ TreeBrowser, which supports pruning of nodes at every level of a tree. A user has pruned an 1100-node tree of a satellite network, using dynamic query sliders at four levels; only nine possible ports (leaf nodes) remain in the result set. (Kumar et al., 1997)

15.17) (Sarkar and Brown, 1994; Bartram et al., 1995; Schaffer et al., 1996). The fisheye distortion magnifies one or more areas of the display, but zoom factors in prototypes are limited to about 5. Although query-language facilities made it difficult to gain an overview of a collection, information-visualization interfaces support some overview strategy—or should do so. Adequate overview strategies are a useful criterion to judge such interfaces. In addition, look for navigation tools to pan or scroll through the collection.

Zoom task

We can zoom in on items of interest. Users typically have an interest in some portion of a collection, and they need tools to enable them to control the zoom focus and the zoom factor. Smooth zooming helps users to pre-

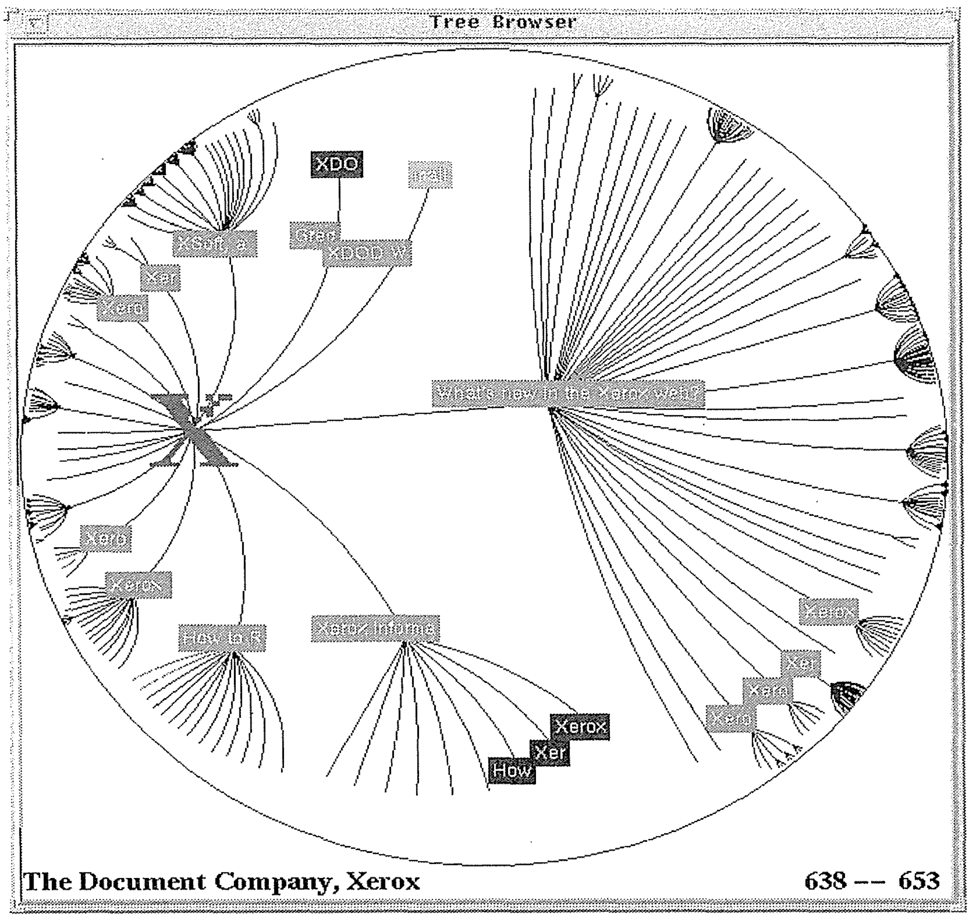

Figure 15.14

A hyperbolic tree browser that allows 10 to 30 nodes near the center to be seen clearly; branches are reduced gradually as they get closer to the periphery. This display technique guarantees that large trees can be accommodated in a fixed screen size. As the focus is shifted among nodes, the display updates smoothly, producing a satisfying animation. Landmarks or other features can be introduced to reduce the disorienting effect of movement (Lamping, Rao et al., 1995). (Used with permission of InXight Software, Palo Alto, CA.)

serve their sense of position and context (Schaffer et al., 1996). A user can zoom on one dimension at a time by moving the zoombar controls or by adjusting the size of the field-of-view box. A satisfying way to zoom in is to point to a location and to issue a zooming command, usually by holding down a mouse button (Bederson and Hollan, 1993). Zooming in one dimension has proved useful in starfield displays (Jog and Shneiderman, 1995).

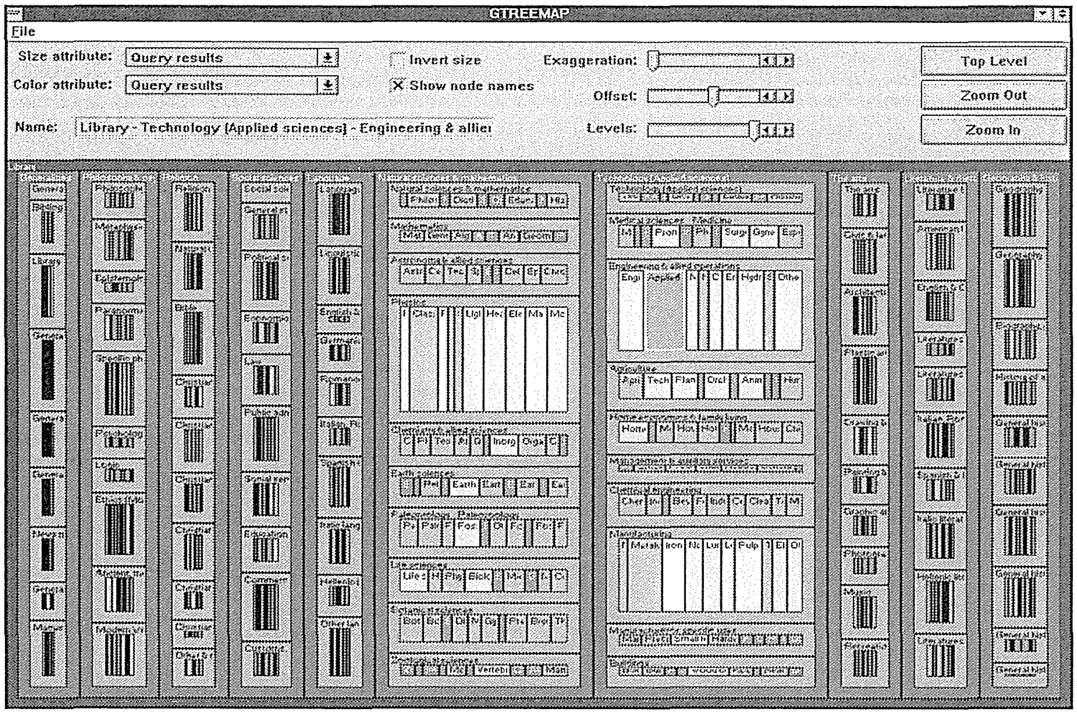

Figure 15.15

The first three levels of the Dewey Decimal System shown as a treemap in which size indicates the number of books held in each of the 1000 categories. Color indicates frequency of utilization, with darker indicating high utilization (hot) and lighter indicating low utilization. Implemented by Marko Teittinen at the University of Maryland Human–Computer Interaction Laboratory.

Filter task

We can filter out uninteresting items. Dynamic queries applied to the items in the collection constitute one of the key ideas in information visualization (Ahlberg et al., 1992; Williamson and Shneiderman, 1992; Kumar et al., 1997). When users control the contents of the display, they can quickly focus on their interests by eliminating unwanted items. Sliders, buttons, or other control widgets coupled to rapid (less than 100 milliseconds) display update is the goal, even when there are tens of thousands of displayed items.

Details-on-demand task

We can select an item or group to get details. Once a collection has been trimmed to a few dozen items, it should be easy to browse the details about the group or individual items. The usual approach is to simply click on an item to get a pop-up window with values of each of the attributes.

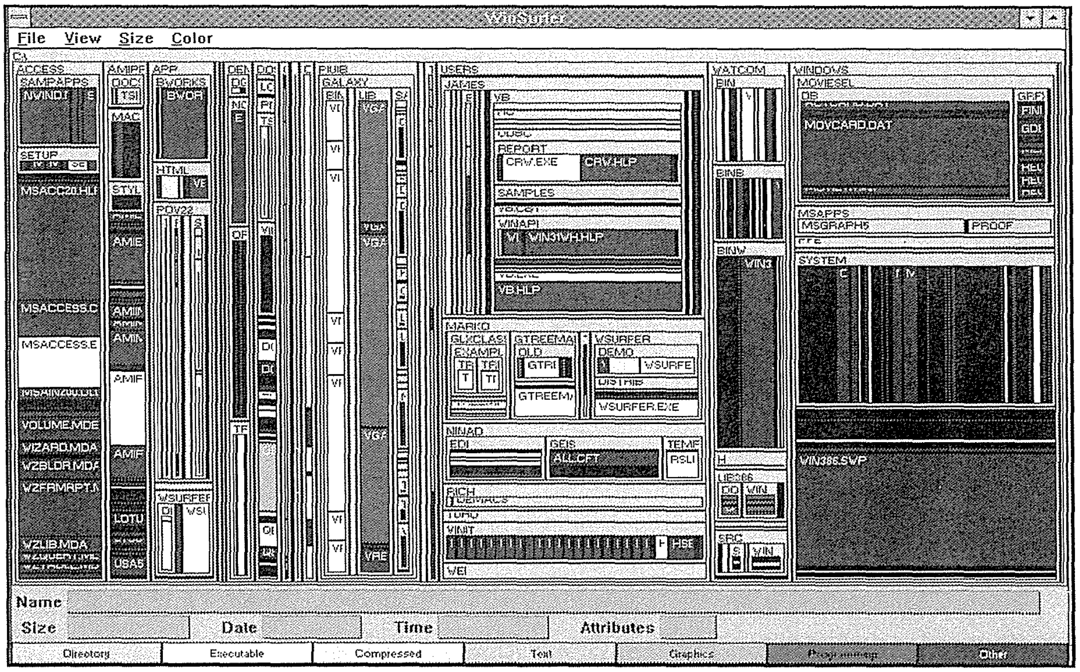

Figure 15.16

Winsurfer treemap that shows the 4900 files at several levels on a hard disk. Area is set to be proportional to file size and color to file type. Moving the cursor over an area produces an immediate display of attribute values on the bottom. Developed by Marko Teittinen at the University of Maryland Human–Computer Interaction Laboratory.

In Spotfire (Color Plate B5), the details-on-demand window can contain HTML text with links to further information.

Relate task

We can view relationships among items. In the FilmFinder details-on-demand window (Ahlberg and Shneiderman, 1994), users could select an attribute, such as the film's director, and cause the director alphaslider to be reset to the director's name, thereby displaying only films by that director. Similarly, in SDM (Chuah et al., 1995), users can select an item and then highlight items with similar attributes. In LifeLines (Color Plate B3) (Plaisant et al., 1996), users can click on a medication and see the related visit report, prescriptions, and laboratory test results. Designing user-interface actions to specify which relationship is to be manifested is still a challenge. The Influence Explorer (Tweedie et al., 1996) emphasizes exploration of relationships among attributes. The Table Lens (Fig. 15.11) emphasizes finding correlations among pairs of numerical attributes (Rao and Card, 1994).

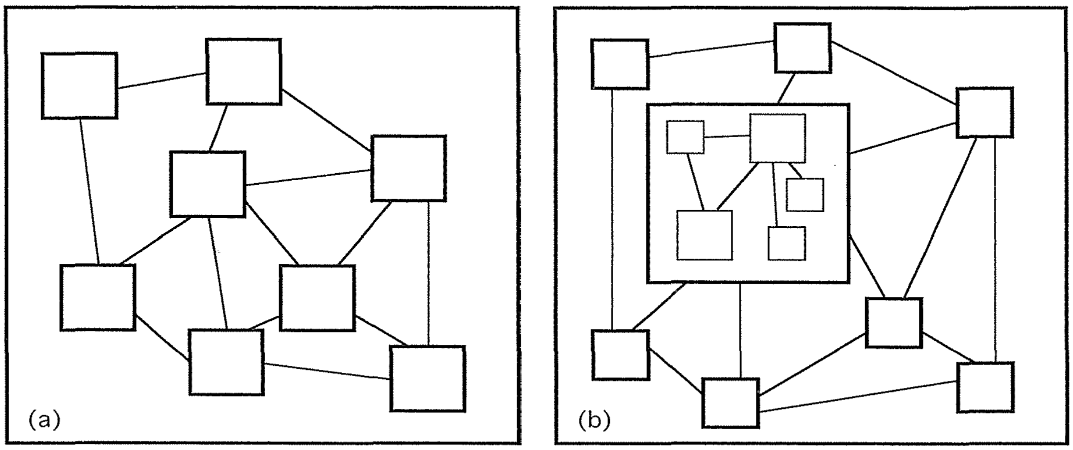

Figure 15.17

A fisheye view or variable zooming on a hierarchical network diagram. These techniques can help to focus attention on details while preserving context. (a) A central node has been selected for zooming. (b) The node is expanded by a zoom factor of 3, exposing five nodes at the next level. (Schaffer et al., 1996.)

History task

We can keep a history of actions to support undo, replay, and progressive refinement. It is rare that a single user action produces the desired outcome. Information exploration is inherently a process with many steps, so keeping the history of actions and allowing users to retrace their steps is important. However, most prototypes fail to deal with this requirement. Maybe they are reflecting the current state of GUIs, but designers would do better to model information-retrieval systems, which typically preserve the sequence of searches so that these searches can be combined or refined.

Extract task

We can allow extraction of subcollections and of the query parameters. Once users have obtained the item or set of items that they desire, it would be useful for them to be able to extract that set and to save it to a file in a format that would facilitate other uses, such as sending by electronic mail, printing, graphing, or insertion into a statistical or presentation package. As an alternative to saving the set, they might want to save, send, or print the settings for the control widgets. Few prototypes support such actions, although Roth's recent work on Visage provides an elegant capability to extract sets of items and simply drag-and-drop them into the next application window (Roth et al., 1996).

The attraction of visual displays, when compared to textual displays, is that they make use of the remarkable human perceptual ability for visual information. Within visual displays, there are opportunities for showing relationships by proximity, by containment, by connected lines, or by color coding. Highlighting techniques (for example, boldface text or brightening, inverse video, blinking, underscoring, or boxing) can be used to draw attention to certain items in a field of thousands of items. Pointing to a visual display can allow rapid selection, and feedback is apparent. The eye, the hand, and the mind seem to work smoothly and rapidly as users perform actions on visual displays.

15.5 Advanced Filtering

Users have highly varied needs for filtering features. The dynamic-queries approach of adjusting numeric range sliders, alphasliders for names or categories, or buttons for small sets of categories is appealing to many users for many tasks (Shneiderman, 1994). Dynamic queries might be called *direct-manipulation queries,* since they share the same concepts of visual display of actions (the sliders or buttons) and objects (the query results in the task-domain display); the use of rapid, incremental, and reversible actions; and the immediate display of feedback (less than 100 milliseconds). Additional benefits are the prevention of syntax errors and an encouragement of exploration.

Dynamic queries can reveal global properties, as well as assist users in answering specific questions. As the database grows, it is more difficult to update the display fast enough, and specialized data structures or parallel computation is required. Dynamic queries have attracted attention, although many user-interface problems remain; for example, we need to discover how to perform these tasks:

- Select a set of sliders from a large set of attributes.
- Specify greater than, less than, or greater than and less than.
- Deal with Boolean combinations of slider settings.
- Choose among highlighting by color, by points or light, by regions, by blinking, and so on.
- Cope with tens of thousands of points.
- Permit weighting of criteria.

The dynamic-query approach to the chemical table of elements was tested in an empirical comparison with a form-fillin query interface (Ahlberg et al., 1991). The counterbalanced-ordering within-subjects design with 18 chemistry students showed strong advantages for the dynamic queries, in terms of faster performance and lower error rates (Ahlberg et al., 1991).

Commercial information-retrieval systems, such as DIALOG or First-Search, permit complex Boolean expressions with parentheses, but their widespread adoption has been inhibited by their difficulty of use. Numerous proposals have been put forward to reduce the burden of specifying complex Boolean expressions (Reisner, 1988). Part of the confusion stems from informal English usage, in which a query such as "List all employees who live in New York and Boston" usually would result in an empty list because the "and" would be interpreted as an intersection; only employees who live in both cities would qualify! In English, "and" usually expands the options; in Boolean expressions, AND is used to narrow a set to the intersection of two others. Similarly, in the English "I'd like Russian or Italian salad dressing," the "or" is exclusive, indicating that you want one or the other but not both; in Boolean expressions, an OR is inclusive, and is used to expand a set.

The desire for *full Boolean expressions,* including nested parentheses and NOT operators, has led to novel metaphors for query specification. *Venn diagrams* (Michard, 1982) and *decision tables* (Greene et al., 1990) have been used, but these representations become clumsy as query complexity increases. To support arbitrarily complex Boolean expressions with a graphical specification, we applied the metaphor of water flowing from left to right through a series of filters, where each filter lets through only the appropriate documents, and the flow paths indicate AND or OR (Young and Shneiderman, 1993).

In this filter–flow model, ANDs are shown as a linear sequence of filters, suggesting the successive application of required criteria. As the flow passes through each filter, it is reduced, and the visual feedback shows a narrower stream of water. In Fig. 15.18(a) a journal database containing 6741 articles passes through the Date filter, about one-half of the articles satisfy the Date requirements of 94 to 97 (years 1994 to 1997). Only about one quarter of those articles pass through the Language filter, which selects English OR French. Users can also specify ORs across attributes, by putting filters in parallel flow paths (Fig. 15.18b). When the parallel flow paths converge, the width reflects the size of the union of the document sets.

Negation is handled by a NOT operator that, when selected, inverts all currently selected items in a filter (Fig. 15.18b). In the example, NOT 91 allows about 80 percent of the articles to pass the Date filter. Clusters of filters and flow paths (with one ingoing and one outgoing flow) can be made into a single labeled filter. Creation of clusters ensures that the full query can be shown on the display at once, and allows named clusters to be saved in a library for later reuse.

The filter-flow approach has been shown to help novices and intermittent users to specify complex Boolean expressions and to learn Boolean concepts. A usability study was conducted with 20 subjects who had had little experience using Boolean algebra. The prototype filter-flow interface was preferred over textual interface by all 20 subjects, and statistically significant advantages emerged on comprehension and composition tasks.

Another form of filtering is to apply a user-constructed set of keywords to dynamically generated information, such as incoming electronic-mail mes-

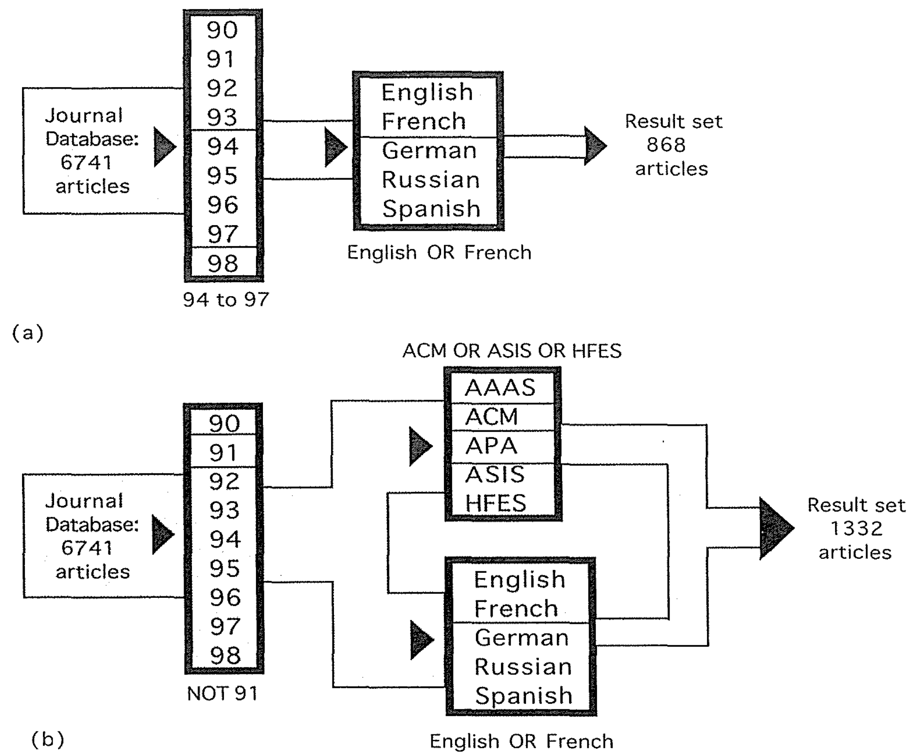

Figure 15.18

(a) Filter-flow model for the query `(Date between 94 to 97) AND (Language is English OR French)`.
(b) Filter-flow model for query `(Date NOT 91) AND (Publisher is ACM OR ASIS OR HFES) OR (Language is English OR French))`.

sages, newspaper stories, or scientific journal articles (Belkin and Croft, 1992). The users create and store their profiles, which are evaluated each time that a new document appears. Users can be notified by electronic mail that a relevant document has appeared, or the results can be simply collected into a file until the users seek them out. These approaches are a modern version of traditional information-retrieval strategy called *selective dissemination of information (SDI)*, which was used in the earliest days of magnetic-tape distribution of document collections. Elaborate strategies for using the user-supplied set of keywords include latent semantic indexing, use of thesauri for find narrower or broader terms, and natural-language parsing techniques (Foltz and Dumais, 1992). Use of these strategies and term-frequency data can produce relevance rankings of retrieved documents that are appealing to many users and are successful in increasing the recall and precision of searches. A series of text-retrieval conferences (TREC) organized by Donna Harman at the National Institute for Standards and

Technology (http://potomac.ncsl.nist.gov/TREC) has allowed developers of research and commercial products to compare their strategies against a large test collection of textual documents.

A social form of filtering is *collaborative filtering*, in which groups of users combine their evaluations to help one another find interesting items in a large document collection (Resnick et al., 1994). Each user rates documents in terms of their interest. Then, the system can suggest unread articles that are close to the user's interests, as determined by matches with other people's interests. This method can also be applied to movies, music, restaurants, and so on. For example, if you rate six restaurants high, the algorithms will provide you with other restaurants that were rated high by people who liked your six restaurants. This strategy has an inherent appeal, and dozens of systems have been built for organizational databases, news files, music groups, and World Wide Web pages.

15.6 Practitioner's Summary

Improved user interfaces to traditional database-query and text- or multimedia-document search will spawn appealing new products. Flexible queries against complex text, sound, graphics, image, and video databases are emerging. Novel graphical and direct-manipulation approaches to query formulation and information visualization are now possible. Whereas research prototypes have typically dealt with only one data type (one-, two-, and three-dimensional data; temporal and multidimensional data; and tree and network data), successful commercial products will have to accommodate several. These products will need to provide smooth integration with existing software and to support the full task list: overview, zoom, filter, details-on-demand, relate, history, and extract. These methods are attractive because they present information rapidly and allow user-controlled exploration. If they are all to be fully effective, we will require advanced data structures, high-resolution color displays, fast data retrieval, and novel forms of user training. Many user interfaces for specifying advanced filtering are being built and are worthy of evaluation for commercial projects.

15.7 Researcher's Agenda

Although the computer contributes to the information explosion, it is potentially the magic lens for finding, sorting, filtering, and presenting the relevant items. Search in complex structured documents, graphics, images, sound, or video presents grand opportunities for the design of advanced user interfaces and powerful search engines to find the needles in the haystacks and the forests beyond the trees. The novel information-exploration tools—such as dynamic queries, treemaps, fisheye views, parallel coordinates, starfields, and perspective walls—are but a few of the inventions that will have to be tamed and validated by user-interface researchers. A better integration with perceptual psychology (understanding preattentive processes and the impact of varied coding or highlighting techniques) and with business decision making (identifying tasks and procedures that occur in realistic situations) is needed, as are theoretical foundations and practical benchmarks for choosing among the diverse emerging visualization techniques. Empirical studies would help to sort out the specific situations in which visualization was most helpful. Finally, software toolkits for building innovative visualizations would facilitate the exploration process.

World Wide Web Resources

WWW

The search services such as Alta Vista, Excite, Infoseek, and Lycos provide remarkable but flawed access to the World Wide Web. Other information retrieval topics such as collaborative filtering, document summarization, and indexing methods are covered. Information visualization tools are growing more effective for a wider range of tasks.

http://www.aw.com/DTUI

References

Ahlberg, Christopher and Shneiderman, Ben, Visual information seeking: Tight coupling of dynamic query filters with starfield displays, *Proc. CHI '94 Conference: Human Factors in Computing Systems*, ACM, New York (1994), 313–321 and color plates.

Ahlberg, Christopher and Shneiderman, Ben, AlphaSlider: A compact and rapid selector, *Proc. of ACM CHI '94 Conference Human Factors in Computing Systems*, ACM, New York (1994), 365–371.

Ahlberg, Christopher, Williamson, Christopher, and Shneiderman, Ben, Dynamic queries for information exploration: An implementation and evaluation, *Proc. ACM CHI '92: Human Factors in Computing Systems*, ACM, New York (1992), 619–626.

Ahlberg, Christopher and Wistrand, Erik, IVEE: An information visualization and exploration environment, *Proc. IEEE Information Visualization '95*, IEEE Computer Press, Los Alamitos, CA (1995), 66–73.

Andrews, Keith, Visualising cyberspace: Information visualisation in the Harmony internet browser, *Proc. IEEE Information Visualization '95*, IEEE Computer Press, Los Alamitos, CA (1995), 97–104.

Asahi, T., Turo, D., and Shneiderman, B., Using treemaps to visualize the analytic hierarchy process, *Information Systems Research*, 6, 4 (December 1995), 357–375.

Bartram, Lyn, Ho, Albert, Dill, John, and Henigman, Frank, The continuous zoom: A constrained fisheye technique for viewing and navigating large information spaces, *Proc. User Interface Software and Technology '95*, ACM, New York (1995), 207–215.

Becker, Richard A., Eick, Stephen G., and Wilks, Allan R. Visualizing network data, *IEEE Transactions on Visualization and Computer Graphics*, 1, 1 (March 1995), 16–28.

Bederson, Ben B. and Hollan, James D., PAD++: A zooming graphical user interface for exploring alternate interface physics, *Proc. User Interfaces Software and Technology '94* (1994), 17–27.

Belkin, Nick J. and Croft, Bruce W. Information filtering and information retrieval: Two sides of the same coin?, *Communications of the ACM*, 35, 12 (1992), 29–38.

Bertin, Jacques, *Semiology of Graphics*, University of Wisconsin Press, Madison, WI (1983).

Borgman, Christine, L., Why are online catalogs hard to use? Lessons learned from information-retrieval studies, *Journal of the American Society for Information Science*, 37, 6 (1986), 387–400.

Card, Stuart K., Robertson, George G., and York, William, The WebBook and the WebForager: An information workspace for the World-Wide Web, *Proc. CHI' 96 Conference: Human Factors in Computing Systems*, ACM, New York (1996), 111–117.

Carriere, Jeremy and Kazman, Rick, Interacting with huge hierarchies: Beyond cone trees, *Proc. IEEE Information Visualization '95*, IEEE Computer Press, Los Alamitos, CA (1995), 74–81.

Chimera, Richard, Value bars: An information visualization and navigation tool for multiattribute listings, *Proc. CHI '92 Conference: Human Factors in Computing Systems*, ACM, New York (1992), 293–294.

Chimera, Richard and Shneiderman, Ben, Evaluating three user interfaces for browsing tables of contents, *ACM Transactions on Information Systems*, 12, 4 (October 1994), 383–406.

Chuah, Mei C., Roth, Steven F., Mattis, Joe, and Kolojejchcik, John, SDM: Malleable Information Graphics, *Proc. IEEE Information Visualization '95*, IEEE Computer Press, Los Alamitos, CA (1995), 66–73.

Cleveland, William, *Visualizing Data*, Hobart Press, Summit, NJ (1993).

Egan, Dennis E., Remde, Joel R., Gomez, Louis M., Landauer, Thomas K., Eberhardt, Jennifer, and Lochbum, Carol C., Formative design-evaluation of SuperBook, *ACM Transactions on Information Systems*, 7, 1 (January 1989), 30–57.

Egenhofer, Max and Richards, J., Exploratory access to geographic data based on the map-overlay metaphor, *Journal of Visual Languages and Computing*, 4, 2 (1993), 105–125.

Eick, Stephen G., Steffen, Joseph L., and Sumner, Jr., Eric E., SeeSoft: A tool for visualizing line-oriented software statistics, *IEEE Transactions on Software Engineering*, 18, 11 (1992) 957–968.

Eick, Stephen G. and Wills, Graham J., Navigating large networks with hierarchies, *Proc. IEEE Visualization '93 Conference* (1993), 204–210.

Fairchild, Kim M., Poltrock, Steven E., and Furnas, George W., SemNet: Three-dimensional representations of large knowledge bases. In Guindon, Raymonde (Editor), *Cognitive Science and its Applications for Human–Computer Interaction*, Lawrence Erlbaum, Hillsdale, NJ (1988), 201–233.

Fishkin, Ken and Stone, Maureen C., Enhanced dynamic queries via movable filters, *Proc. CHI' 95 Conference: Human Factors in Computing Systems*, ACM, New York (1995), 415–420.

Foltz, Peter W. and Dumais, Susan T., Personalized information delivery: An analysis of information filtering methods. *Communications of the ACM*, 35, 12 (1992), 51–60.

Furnas, George W., Generalized fisheye views, *Proc. CHI' 86 Conference: Human Factors in Computing Systems*, ACM, New York (1986), 16–23.

Goldstein, Jade and Roth, Steven F., Using aggregation and dynamic queries for exploring large data sets, *Proc. CHI' 95 Conference: Human Factors in Computing Systems*, ACM, New York (1995), 23–29.

Greene, S. L., Devlin, S. J., Cannata, P. E., and Gomez, L. M., No IFs, ANDs, or ORs: A study of database querying, *International Journal of Man–Machine Studies*, 32, (March 1990), 303–326.

Hendley, R. J., Drew, N. S., Wood, A. S., Narcissus: Visualizing information, *Proc. IEEE Information Visualization '95*, IEEE Computer Press, Los Alamitos, CA (1995), 90–96.

Humphrey, Susanne M. and Melloni, Biagio John, *Databases: A Primer for Retrieving Information by Computer*, Prentice-Hall, Englewood Cliffs, NJ (1986).

Inselberg, Alfred, The plane with parallel coordinates, *The Visual Computer*, 1, (1985), 69–91.

Jarke, M., and Vassiliou, Y., A framework for choosing a database query language, *ACM Computing Surveys*, 11, 3 (1986), 313–340.

Jerding, Dean F. and Stasko, John T., The information mural: A technique for displaying and navigating large information spaces, *Proc. IEEE Information Visualization '95*, IEEE Computer Press, Los Alamitos, CA (1995), 43–50.

Jog, Ninad and Shneiderman, Ben, Information visualization with smooth zooming on an starfield display *Proc. Visual Databases 3* (March 1995), 1–10.

Johnson, Brian, and Shneiderman, Ben, Tree-maps: A space-filling approach to the visualization of hierarchical information structures, *Proc. IEEE Visualization '91*, IEEE, Piscataway, NJ (1991), 284–291.

Keim, D. A. and Kriegal, H., VisDB: Database exploration using multidimensional visualization, *IEEE Computer Graphics and Applications* (September 1994), 40–49.

Kim, H. J., Korth, H. F., and Silberschatz, A., PICASSO: A graphical query language, *Software: Practice and Experience*, 18, 3 (1988), 169–203.

Koenemann, Juergen and Belkin, Nicholas, A case for interaction: A study of interactive information retrieval behavior and effectiveness, *Proc. CHI '96 Human Factors in Computing Systems*, ACM Press, New York (1996), 205–212.

Korfhage, Robert, To see or not to see: Is that the query?, *Communications of the ACM*, 34 (1991), 134–141.

Kumar, Harsha, Plaisant, Catherine, and Shneiderman, Ben, Browsing hierarchical data with multi-level dynamic queries and pruning, *International Journal of Human–Computer Studies*, 46, 1 (January 1997), 103–124.

Lamping, John, Rao, Ramana, and Pirolli, Peter, A focus + context technique based on hyperbolic geometry for visualizing large hierarchies, *Proc. of CHI '95*

Conference: Human Factors in Computing Systems, ACM, New York (1995), 401–408.

Laurini, R. and Thompson, D., *Fundamentals of Spatial Information Systems*, Academic Press, New York (1992).

Marchionini, Gary, *Information Seeking in Electronic Environments*, Cambridge University Press, UK (1995).

Marchionini, Gary and Shneiderman, Ben, Finding facts and browsing knowledge in hypertext systems, *IEEE Computer*, 21, 1 (January 1988), 70–80.

Mark, Leo, A graphical query language for the binary relationship model, *Information Systems*, 14, 3 (1989), 231–246.

McCormick, B., DeFanti, T, and Brown, R. (Editors), Visualization in scientific computing and computer graphics, *ACM SIGGRAPH*, 21, 6 (November 1987).

Michard, A., A new database query language for non-professional users: Design principles and ergonomic evaluation, *Behavioral and Information Technology*, 1, 3 (July–September 1982), 279–288.

Mitchell, Richard, Day, David, and Hirschman, Lynette, Fishing for information on the internet, *Proc. IEEE Information Visualization '95*, IEEE Computer Press, Los Alamitos, CA (1995), 105–111.

Mukherjea, Sougata, Foley, James D., and Hudson, Scott, Visualizing complex hypermedia networks through multiple hierarchical views, *Proc. of ACM CHI '95 Conference: Human Factors in Computing Systems*, ACM, New York (1995), 331–337 plus color plate.

North, Chris, Shneiderman, Ben, and Plaisant, Catherine, User controlled overviews of an image library: A case study of the Visible Human, *Proc. 1st ACM International Conference on Digital Libraries* (1996), 74–82.

Pirolli, Peter, Schank, Patricia, Hearst, Marti, and Diehl, Christine, Scatter/gather browsing communicates the topic structure of a very large text collection, *Proc. of ACM CHI' 96 Conference*, ACM, New York (1996), 213–220.

Plaisant, Catherine, Rose, Anne, Milash, Brett, Widoff, Seth, and Shneiderman, Ben, LifeLines: Visualizing personal histories, *Proc. of CHI' 96 Conference: Human Factors in Computing Systems*, ACM, New York (1996), 221–227, 518.

Reisner, Phyllis, Query languages. In Helander, Martin (Editor), *Handbook of Human–Computer Interaction*, North-Holland, Amsterdam, The Netherlands (1988), 257–280.

Rao, Ramana and Card, Stuart K., The Table Lens; Merging graphical and symbolic representations in an interactive focus + context visualization for tabular information, *Proc. CHI '94 Conference: Human Factors in Computing Systems*, ACM, New York (1994), 318–322.

Resnick, Paul, Iacovou, Neophytos, Suchak, Mitesh, Bergstrom, Peter and Riedl, John, GroupLens: An open architecture for collaborative filtering of netnews, *Proc. Conference on Computer Supported Cooperative Work '94*, ACM, New York (1994), 175–186.

Robertson, George G., Card, Stuart K., and Mackinlay, Jock D., Information visualization using 3-D interactive animation, *Communications of the ACM*, 36, 4 (April 1993), 56–71.

Robertson George G. and Mackinlay, Jock D., The document lens, *Proc. 1993 ACM User Interface Software and Technology*, ACM New York (1993), 101–108.

Roth, Steven F., Lucas, Peter, Senn, Jeffrey A., Gomberg, Cristina C., Burks, Michael B., Stroffolino, Philip J., Kolojejchick, John A. and Dunmire, Carolyn, Visage: A user interface environment for exploring information, *Proc. IEEE Information Visualization '96*, IEEE Computer Press, Los Alamitos, CA (1996), 3–12.

Salton, G., Automatic Text Processing: *The Transformation, Analysis, and Retrieval of Information by Computer*, Addison-Wesley, Reading, MA (1989).

Sarkar, Manojit and Brown, Marc H., Graphical fisheye views, *Communications of the ACM*, 37, 12 (July 1994), 73–84.

Schaffer, Doug, Zuo, Zhengping, Greenberg, Saul, Bartram, Lyn, Dill, John, Dubs, Shelli and Roseman, Mark, Navigating hierarchically clustered networks through fisheye and full-zoom methods, *ACM Transactions on Computer–Human Interaction*, 3, 2 (June 1996), 162–188.

Shneiderman, Ben, Tree visualization with tree-maps: A 2-D space-filling approach, *ACM Transactions on Graphics*, 11, 1 (January 1992), 92–99.

Shneiderman, Ben, Dynamic queries for visual information seeking, *IEEE Software*, 11, 6 (1994), 70–77.

Shneiderman, Ben, Brethauer, Dorothy, Plaisant, Catherine and Potter, Richard, Three evaluations of museum installations of a hypertext system, *Journal of the American Society for Information Science*, 40, 3 (May 1989), 172–182.

Shneiderman, Ben, Byrd, Donald, and Croft, Bruce, Clarifying search: A user-interface framework for text searches, *D-LIB Magazine of Digital Library Research* (January 1997), http://www.dlib.org/.

Spence, Robert and Apperley, Mark, Data base navigation: An office environment for the professional, *Behaviour & Information Technology*, 1, 1 (1982), 43–54.

Spoerri, Anslem, InfoCrystal: A visual tool for information retrieval and management, *Proc. ACM Conf. on Information and Knowledge Management* (1993), 150–157.

Tufte, Edward, *The Visual Display of Quantitative Information*, Graphics Press, Cheshire, CT (1983).

Tufte, Edward, *Envisioning Information*, Graphics Press, Cheshire, CT (1990).

Tweedie, Lisa, Spence, Robert, Dawkes, Huw, and Su, Hua, Externalising abstract mathematical models, *Proc. of CHI' 96 Conference: Human Factors in Computing Systems*, ACM, New York (1996), 406–412.

Weiland, William J. and Shneiderman, Ben, A graphical query interface based on aggregation/generalization hierarchies,, *Information Systems*, 18, 4 (1993), 215–232.

Welty, C., Correcting user errors in SQL, *International Journal of Man–Machine Studies*, 22 (1985), 463–477.

Williamson, Christopher, and Shneiderman, Ben, The Dynamic HomeFinder: Evaluating dynamic queries in a real-estate information exploration system, *Proc. ACM SIGIR '92 Conference*, ACM, New York (1992), 338–346. Reprinted in Shneiderman, B. (Editor), *Sparks of Innovation in Human–Computer Interaction*, Ablex Publishers, Norwood, NJ (1993), 295–307.

Wise, James A., Thomas, James, J., Pennock, Kelly, Lantrip, David, Pottier, Marc, Schur, Anne, and Crow, Vern, Visualizing the non-visual: Spatial analysis and interaction with information from text documents, *Proc. IEEE Information Visualization '95*, IEEE Computer Press, Los Alamitos, CA (1995), 51–58.

Wurman, Richard Saul, *Information Anxiety*, Doubleday, New York (1989).

Young, Degi and Shneiderman, Ben, A graphical filter/flow model for Boolean queries: An implementation and experiment, *Journal of the American Society for Information Science*, 44, 6 (July 1993), 327–339.

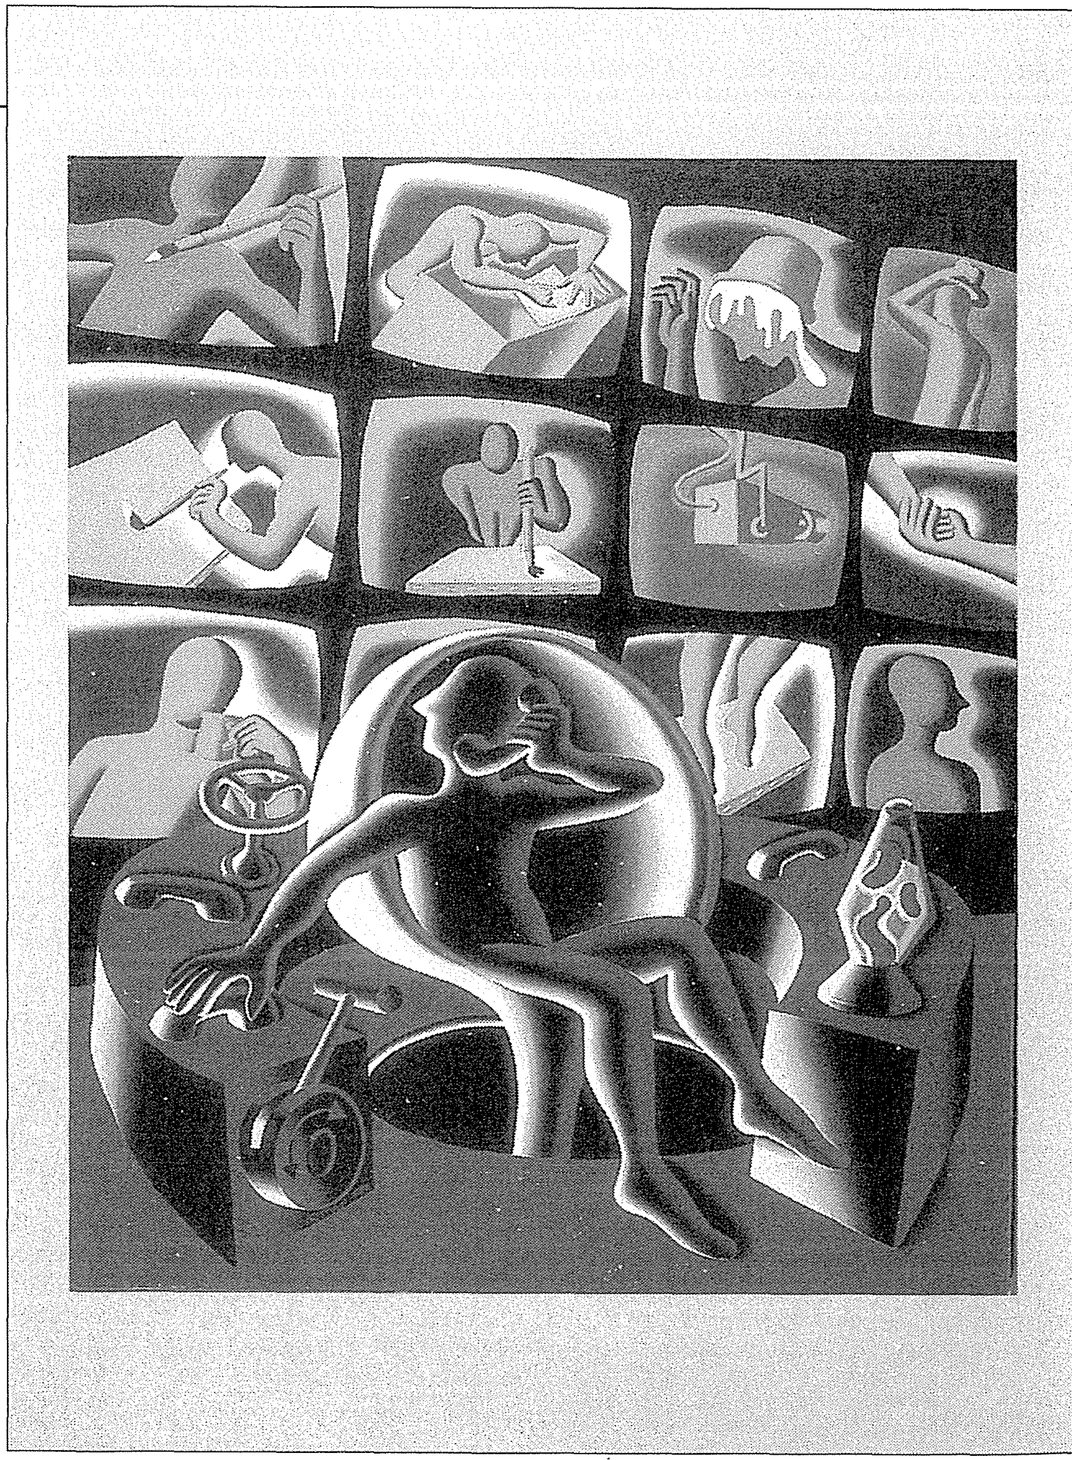

Mark Kostabi, *Quality Control*, 1990

CHAPTER

16

Hypermedia and the World Wide Web

Gradually I began to feel that we were growing something almost organic in a new kind of reality, in cyberspace, growing it out of information . . . a pulsing tree of data that I loved to climb around in, scanning for new growth.

Mickey Hart, *Drumming at the Edge of Magic: A Journey into the Spirit of Percussion,* **1990**

Look at every path closely and deliberately.
Try it as many times as you think necessary.
Then ask yourself, and yourself alone, one question . . .
Does this path have a heart?
If it does, the path is good; if it doesn't, it is of no use.

Carlos Castaneda, *The Teachings of Don Juan*

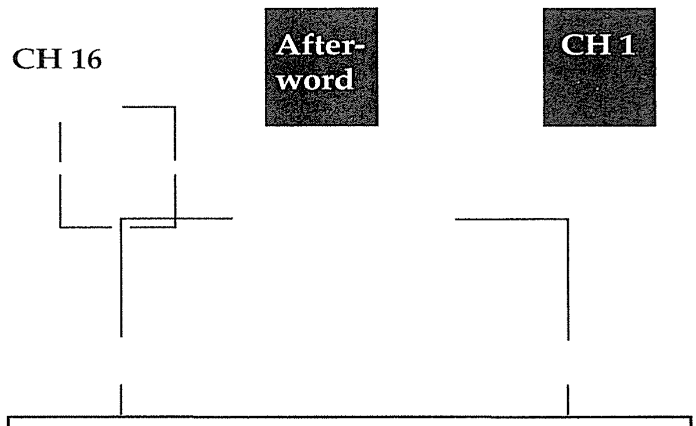

Chapter 16

16.1 Introduction

In July 1945, Vannevar Bush, President Franklin Roosevelt's Science Adviser, wrote a provocative article (Bush, 1945) offering his vision of science projects that might become feasible in the post-World War II period. He wisely identified the information-overload problem and sought to make cross-references within and across documents easy to create and traverse. His desktop information-exploration tool, *memex,* was based on microfilm and eye-tracking technology. Memex would enable readers to follow cross-references by merely staring at them:

> Wholly new forms of encyclopedias will appear, ready-made with a mesh of associative trails running through them, ready to be dropped into the memex

> and there amplified. The lawyer has at his touch the associated opinions and decisions of his whole experience, and of the experience of friends and authorities. . . . There is a new profession of trail blazers, those who find delight in establishing useful trails through the enormous mass of the common record. The inheritance from the master becomes, not only his addition to the world's record, but for his disciples the entire scaffolding by which they were erected. (Bush, 1945)

It has taken 50 years to create effective—although somewhat revised—models of Bush's vision. Now the technology is beginning to make possible a useful reading, browsing, linking, and annotating environment to support communal nonlinear writing and reading. The name *hypertext*, or *hypermedia*, has been applied to networks of nodes (also called articles, documents, files, cards, pages, frames, screens) containing information (in text, graphics, video, sound, and so on) that are connected by links (also called pointers, cross-references, citations). *Hypertext* is more commonly applied to text-only applications, whereas *hypermedia* is used to convey the inclusion of other media, especially sound and video. The World Wide Web extends the hypermedia to a vast network of computers in which millions of users can create and retrieve multimedia materials from around the world in seconds.

Ted Nelson coined the term *hypertext* in the 1960s as he was writing about his universal library and *docuverse*, with *stretch text* that expands when selected. Nelson's enthusiasm and imagination infected many people who shared his *computopian* hopes. Using less flamboyant terms, Douglas Engelbart created his Human Augmentation system at SRI during the 1960s, with hypertext point-and-click features, expanding outline processors, multiple windows, remote collaboration, and the mouse (Engelbart, 1984). In parallel, Andries van Dam developed early electronic books at Brown University using colorful dynamic graphics and three-dimensional animation (Yankelovich et al., 1985; van Dam, 1988).

By the mid-1980s, many research and commercial packages offered hypertext features to enable convenient jumps among articles (Conklin, 1987; Halasz, 1988; Shneiderman and Kearsley, 1989; Nielsen, 1995). Pioneering hypertext systems include NoteCards, developed at Xerox PARC; KMS, from Knowledge Systems, Inc.; Guide from OWL International; and Hyperties, originated at the University of Maryland (Shneiderman, 1989) and commercially developed by Cognetics Corporation of Princeton, New Jersey (Fig 7.5, Fig. 12.6, and Color Plate C1).

Hyperties was conceived as a publication tool, with which authors produce hypermedia for thousands of readers. It has separate tools for browsing and authoring documents. Hyperties was based on the metaphor of an electronic encyclopedia. Each document was called an article, and cross-references were implemented as highlighted text links and image maps. Using the metaphor of an encyclopedia comprising a collection of titled articles made acceptance easy and facilitated navigation. Built into Hyperties were author-generated and alphabetical tables of contents, including every article plus history lists supporting reversible actions. The Hyperties browser was one of the first software packages that needed no error messages, because the design prevented the user from making syntactic errors.

In the late 1980s, commercial hypertext applications began to appear. In 1987, Apple provided Bill Atkinson's HyperCard system free with every Macintosh. Although the brochures referred to Vannevar Bush's vision, Apple refrained from using the term *hypertext* in describing HyperCard (Fig. 16.1a–b). Building on the metaphor of cards arranged in stacks, Apple claimed in the online help that "you can use HyperCard to create your own applications for gathering, organizing, presenting, searching and customizing information."

The July 1988 *Communications of the ACM* contained eight papers from the first hypertext conference. Three electronic versions of this issue, built with KMS, HyperCard, and Hyperties (Shneiderman, 1988), were marketed by ACM to thousands of professionals. We used Hyperties the following year to create the first commercial electronic book, *Hypertext Hands-On!* (Shneiderman and Kearsley, 1989). Hewlett-Packard used Hyperties to distribute electronic documentation for its LaserJet 4 printers in 15 languages. That may have been the first world-wide distribution of hypertext prior to implementation of the World Wide Web.

Today, the World Wide Web uses hypertext to link tens of millions of documents together. The basic highlighted text link can be traced back to an innovation, developed in 1983, as part of *The Interactive Encyclopedia System (TIES)*, the research predecessor to Hyperties. The original concept was to eliminate menus by embedding highlighted link phrases directly in the text (Koved and Shneiderman, 1986). Earlier designs required typing codes, selecting from menu lists, or clicking on visually distracting markers in the text. The embedded-text-link idea was adopted by developers such as Tim Berners-Lee, who described "hot spots" in his 1989 proposal for the World Wide Web (Berners-Lee, 1994).

Other Hyperties features anticipated the World Wide Web. Charles Kreitzberg, Whitney Quesenbery, and programmers at Cognetics implemented image maps, animations, and a markup language called Hyperties Markup Language (HTML). It is quite similar to the HTML markup language used with web browsers; both drew on concepts in SGML, which continues to be an important markup language within the publishing community. Hyperties also had a Java-like scripting language that allowed processes to be attached to pages or to links.

Hypertext has become a mainstream interface paradigm with the emergence of the World Wide Web. In the web, the basic vision of Vannevar Bush, Ted Nelson and the researchers who picked up their challenge has become a reality. Some people might argue that the web does not capture the vision of the early hypertext pioneers, but the distinctions are modest, and innovations are added weekly.

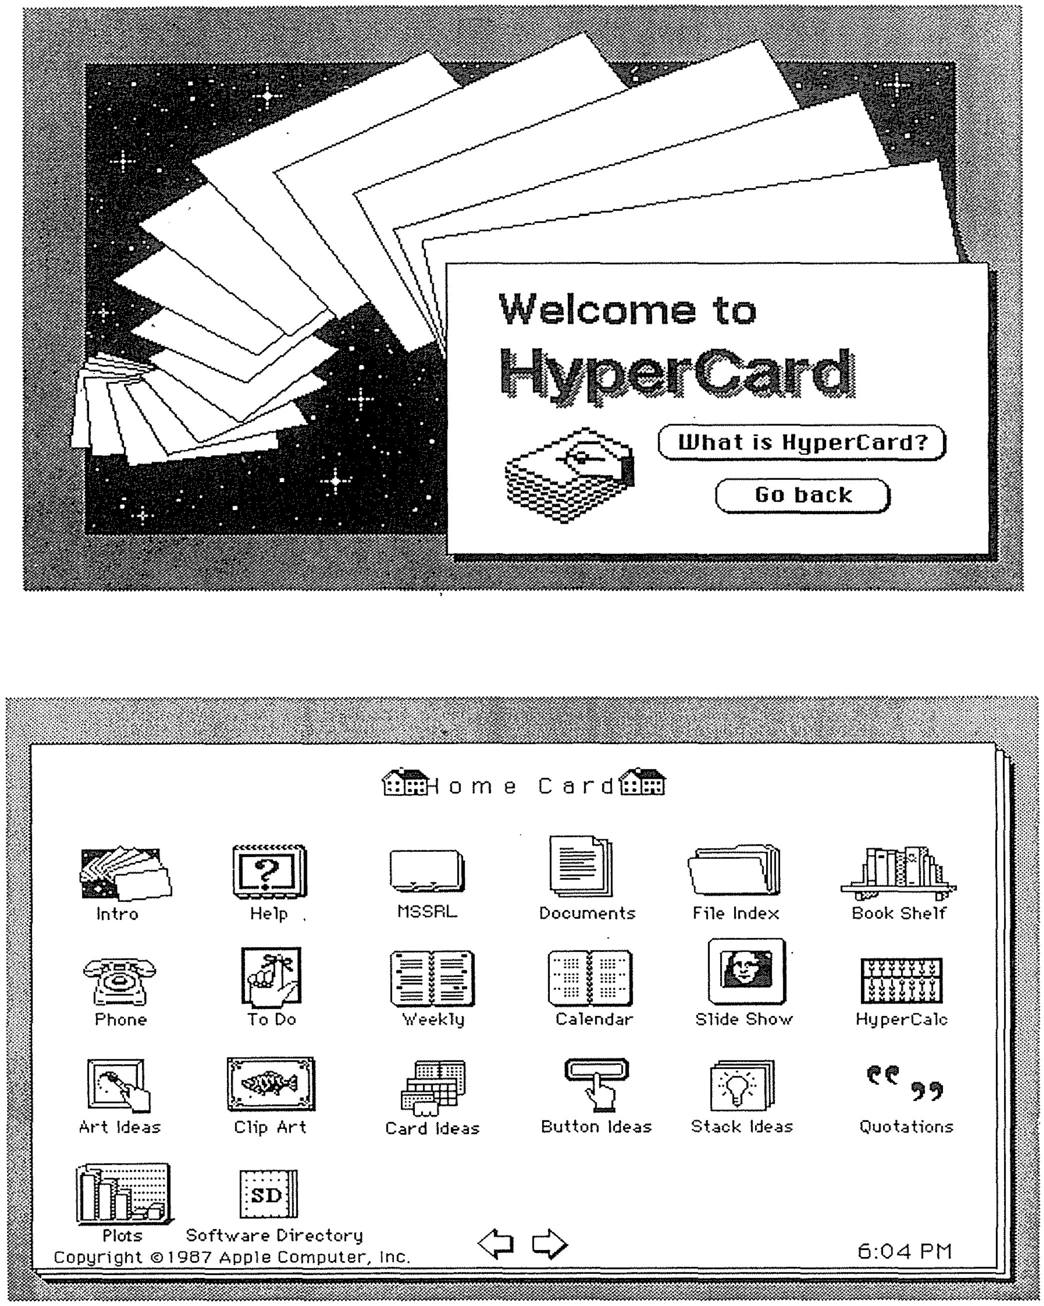

Figure 16.1

HyperCard displays from the original 1987 version. An iconic index of stacks is included. (Courtesy of Apple Computer, Cupertino, CA.)

With a simple mouse click, users can jump to related web pages, which may be delivered from millions of server computers located around the world. With graphics, maps, photos, sound, and the increasing degree of animation supplied by Java applets or VRML, the web is limited only by the bandwidth of the network and the imagination of designers.

16.2 Hypertext and Hypermedia

The intrigue of hypertext is that it extends traditional linear text with the opportunity for jumping to multiple related articles. Convenient backtracking, clickable indexes and tables of contents, string searching, bookmarks, and other navigation tools profoundly alter the reader's experience. For some purposes, hypertext can be a welcome improvement over linear paper documents, but there is a danger that jumping can also lead to hyperchaos. To reduce confusion, hypertext authors need to choose appropriate projects, to organize their articles suitably, and to adjust their writing style to make the best use of this new medium. The first step in creating effective hypertexts is to choose projects that adhere to the *Golden Rules of Hypertext* (Shneiderman, 1989):

1. There is a large body of information organized into numerous fragments.
2. The fragments relate to one another.
3. The user needs only a small fraction of the fragments at any time.

The dual dangers are that hypertext may be inappropriate for some projects and that the design of the hypertext may be poor (for example, too many links or a confusing structure). A traditional novel is written linearly, and the reader is expected to read the entire text from beginning to middle to end. Most poems, fairy tales, newspaper articles, and even the chapters of this book are written in a linear form. Of course, hypernovels, hyperpoems, hyper—fairy tales, hypernewspapers, and hyperbooks are possible; but they require creative rethinking of the traditional forms to satisfy the Golden Rules of Hypertext.

Poor design of hypertext is common: too many links, long chains of links to reach relevant material, or too many long dull articles (Rivlin et al., 1994). Inadequate tables of contents or overviews make it difficult for users to determine what is contained in the hypertext. Breaking a text into linked fragments does not ensure that the result will be effective or attractive. Just as turning a theater production into a movie requires learning new techniques of zooming, panning, closeups, and so on, creating successful hypertext requires learning to use the features of the new medium (Jones and Shneiderman, 1990).

Enthusiasts of hypertext systems often dwell on nonlinear reading, yet there is also a great sense of novelty and adventure in *writing* nonlinear hypertexts. Authoring tools should support at least the features in this tableau of actions and objects:

Actions	**Objects**
import	an article or node
edit	a link
export	collections of articles or nodes
print	webs of links
search	entire hypertext

In constructing the first commercial hyperbook, *Hypertext Hands-On!* (Shneiderman and Kearsley, 1989), we faced two key authoring issues: managing the articles and specifying the links. Hypertext systems should provide an index of all the articles that have been referenced or created, and should allow rapid specification of links. Marking a phrase or a region can usually be accomplished easily, but then it should also be easy to indicate the link destination. Furthermore, if the same phrase appears many times, it should be possible to resolve the link more easily the second time. Other features to consider in an authoring tool for stand-alone or the World Wide Web (for example, Claris HomePage or Microsoft Front Page) are these:

- *Range of editing functions* available (for example, copying, moving, insertion, deletion, global change within and across articles)
- *Availability of lists of links* (in and out), index terms, synonyms, and so on
- *Link verification* to check correctness of links
- *Range of display-formatting commands*, fonts, sizes, highlighting
- *Availability of search-and-replace functions* for making global changes across multiple articles
- *Control of color* (text, background); color can make the text look attractive, but it can also be distracting; since users have different preferences and tasks, it should be possible to reset color-usage parameters
- *Capability to switch easily* between author and browser modes to test ideas
- *Availability of graphics and video* facilities; embedded graphics editors and mechanisms for exploring video segments

- *Possibility of collaboration;* more than one person should be able to edit the hypertext at one time; different people should be able to author components, which are then merged
- *Data compression;* compression algorithms can reduce the size and facilitate distribution
- *Security control;* password control can restrict access to the hypertext or parts of it
- *Encryption;* encryption of sensitive nodes enhances security
- *Reliability;* bug-free performance with no loss of data
- *Possibility of integration* with other software or hardware
- *Import and export of standard interchange formats,* such as SGML

For at least the past 3000 years, authors and editors have explored ways to structure knowledge to suit the linear medium of the written word. When appropriate, authors have developed strategies for linking related fragments of text and graphics even in the linear format. Now, hypertext encourages nonlinear interconnecting links among articles.

The first challenge is to structure the knowledge such that an overview can be presented to the reader in an introductory article. The overall structure of articles must make sense to readers so that they can form a mental image of the topics covered. This image facilitates traversal, reduces disorientation, and lets users know what is and what is *not* in the hypertext.

Hypertext is conducive to the inclusion of appendices, glossaries, examples, background information, original sources, and bibliographic references. Interested readers can pursue the details; casual readers can ignore them.

Creating documents for a hypertext database introduces considerations beyond the usual concerns of good writing. No list can be complete, but this list, derived from our experience, may be useful:

- *Know the users and their tasks* Consult with users throughout the development process, and test your designs. You are not a good judge of your own design.
- *Ensure that meaningful structure comes first* Build the project around the structuring and presentation of information, rather than around the technology. Develop a *high concept* for the body of information that you are organizing.
- *Apply diverse skills* Make certain that the project team includes information specialists (trainers, psychologists, graphic artists), content specialists (users, marketers), and technologists (systems analysts, programmers).

- *Respect chunking* Organize information into chunks that deal with one topic, theme, or idea. Chunks may be 100 words or 1000 words—but when a chunk reaches 10,000 words, consider restructuring into multiple smaller chunks. Screens are still usually small and hard to read, so lengthy linear texts are not pleasant for users.
- *Show interrelationships* Write each article to contain links to other articles. Too few links bore readers; too many links overwhelm and distract. Author preferences range from putting in a maximum of one or two links per screen, to the more common range of two to 20 links per screen, to the extreme of dozens of links per screen. Although pages meant to be read thoroughly should have few links, index pages can contain hundreds of links.
- *Ensure simplicity in traversal* Design the link structure so that navigation is simple and consistent throughout the system.
- *Design each screen carefully* Design screens such they can be grasped easily. The focus of attention should be clear, headings should guide the reader, links should be useful guides that do not overwhelm the reader. Visual layout should be compact vertically, so as to minimize scrolling. Whitespace can be helpful, but blank space is wasteful.
- *Require low cognitive load* Minimize the burden on the user's short-term memory. Do not require the user to memorize terms or codes. The goal is to enable users to concentrate on the contents while the computer vanishes.

Key design questions are how to organize the hypertext, and how to convey that order to the reader. Authoring strategies for creating the introductory article include these:

- *Executive overview* Make the home page or introductory article an overview that summarizes the contents and contains links to all major concepts.
- *Top-down* Adopt a hierarchical approach in which the links in the home page are to major categories only.
- *Menu* Organize the home page as a detailed table of contents.
- *Search strategy* Make string search easily available as a possible first step.

A major concern of hypertext authors is the optimal length for articles. Research suggests that many short articles are preferable to a smaller number of long articles. An experiment at the University of Maryland using the Hyperties system compared two versions: 46 short articles (four to 83 lines) and five long articles (104 to 150 lines). Participants in the study were given

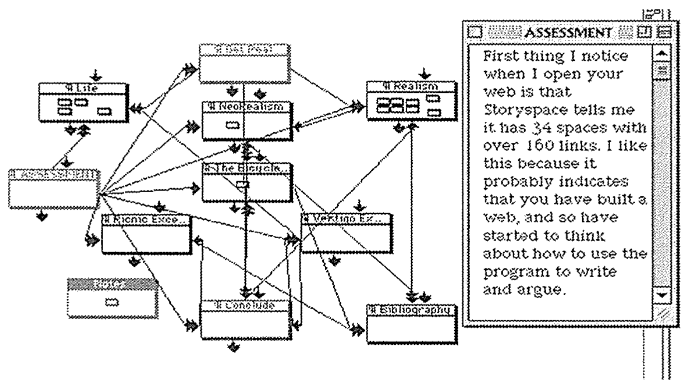

Figure 16.2

Eastgate Systems Storyscape, showing hypertext link structure and an assessment window with a comment. (Used with permission of Eastgate Systems, Watertown, MA.)

30 minutes to locate the answers to a series of questions by using the database. The 16 participants who worked with the short articles answered more questions correctly and took less time to answer the questions. The optimal article length will be affected by screen size, response time, nature of task, and experience of the users. With the longer time for retrieving articles on the web, the preferred length of home pages would be larger. Higher numbers of pointers per page can reduce the number of steps to reach a desired article.

Hypertext authoring has continued as a cottage industry among literary types, encouraged by Eastgate Systems, who also market the Storyscape system (Fig. 16.2). Broader application of hypertext has appeared in Microsoft's and other companies' online help systems, as well as in numerous CD-ROM reference works and encyclopedias, such as Encarta or Compton's (Color Plate C2). These are excellent systems, but the volume of material is dwarfed by the enormous and continuously expanding contents of the World Wide Web.

16.3 World Wide Web

The deluge of web pages has generated dystopian commentaries on the tragedy of the flood of information. It has also produced utopian visions of harnessing the same flood for constructive purposes. Within this ocean of information there are also lifeboat web pages offering design principles, but often the style parallels the early user-interface writings of the 1970s. The well-intentioned Noahs, who write from personal experience as web-site designers, often draw their wisdom from specific projects, so their advice is incomplete or lacks generalizability. Their experience is valuable, but the paucity of empirical data to validate or sharpen insight means that some guidelines are misleading. As scientific evidence accumulates, foundational cognitive and perceptual theories will structure the discussion and guide designers in novel situations.

It may take a decade until sufficient experience, experimentation, and hypothesis testing clarify design issues, so we should be grateful for the early and daring attempts to offer guidance. One of the better guides (Lynch, 1995) offers this advice:

> Proper World Wide Web site design is largely a matter of balancing the structure and relationship of menu or "home" pages and individual content pages or other linked graphics and documents. The goal is to build a hierarchy of menus and pages that feels natural and well-structured to the user, and doesn't interfere with their [sic] use of the Web site or mislead them.

This advice is helpful, but it does not tell designers what to do or how to evaluate the efficacy of what they have done. Lynch goes on to give constructive advice about not being too broad or too deep, finding the proper length of pages, using gridded layouts, and the challenge of "balancing the power of hypermedia Internet linkages against the new ability to imbed graphics and motion media within networked WWW pages." He sorts out the issues better than most, but still leaves designers with uncertainties.

Jakob Nielsen (1995) goes a step further by reporting on his case study of designing a web site for Sun Microsystems to showcase that company's products. His usability-testing approach revealed specific problems, and the web site discusses nine different versions of the home page. The subjective data reveal problems and highlight key principles—for example "Users consistently praised screens that provided overviews of large information spaces." Empirical testing should reveal what kinds of overviews are most effective and whether performance times, error rates, or retention are enhanced by certain overviews.

Refinement of the web is more than a technical challenge or commercial goal. As governments offer information plus services online and educational institutions increase their dependence on the web, effective designs will be essential. Universal access is an important economic and policy issue; it is also a fundamental design issue.

Until the empirical data and experience from practical cases arrive, we can use knowledge from other user-interface design domains, such as menu systems and hypertext (Isakowitz et al., 1995; Shneiderman and Kearsley, 1989).

Designers can apply the theoretical framework of the OAI model and experience from information-retrieval research (Marchionini, 1995). Improved guidelines appear regularly (IBM, 1997), so check the booksite for fresh pointers.

16.4 Genres and Goals for Designers

As they do in any medium, criteria for web-design quality vary with the genre and authors' goals. A dizzying diversity of web sites is emerging from the creative efforts of bold designers who merge old forms to create new information resources, communication media, business services, and entertainment experiences. Web sites can range from a one-page personal biography (Color Plate C3) to millions of pages in the Library of Congress's (Color Plate A5) American Memory project organized by the National Digital Library program (Color Plate C4). Common high-level goals include visual appeal, comprehensibility, utility, efficacy, and navigability, but finer discriminations come into play if we examine the categories of web sites.

A primary way of categorizing web sites is by the originator's identity: individual, group, university, corporation, nonprofit organization, or government agency. The originator's identity gives a quick indication of what the likely goals are and what contents to expect: corporations have products to sell, museums have archives to promote, and government agencies have services to offer.

A second way of categorizing web sites is by goals of the originators, as interpreted by the designers (Table 16.1). Such goals may be simple informa-

Table 16.1

Web-site goals tied to typical organizations

Goal	Organizations
Sell products	publishers, airlines, department stores
Advertise products	auto dealers, real estate agents, movie studios
Inform and announce	universities, museums, cities
Provide access	libraries, newspapers, scientific organizations
Offer services	governments, public utilities
Create discussions	public-interest groups, magazines
Nurture communities	political groups, professional associations

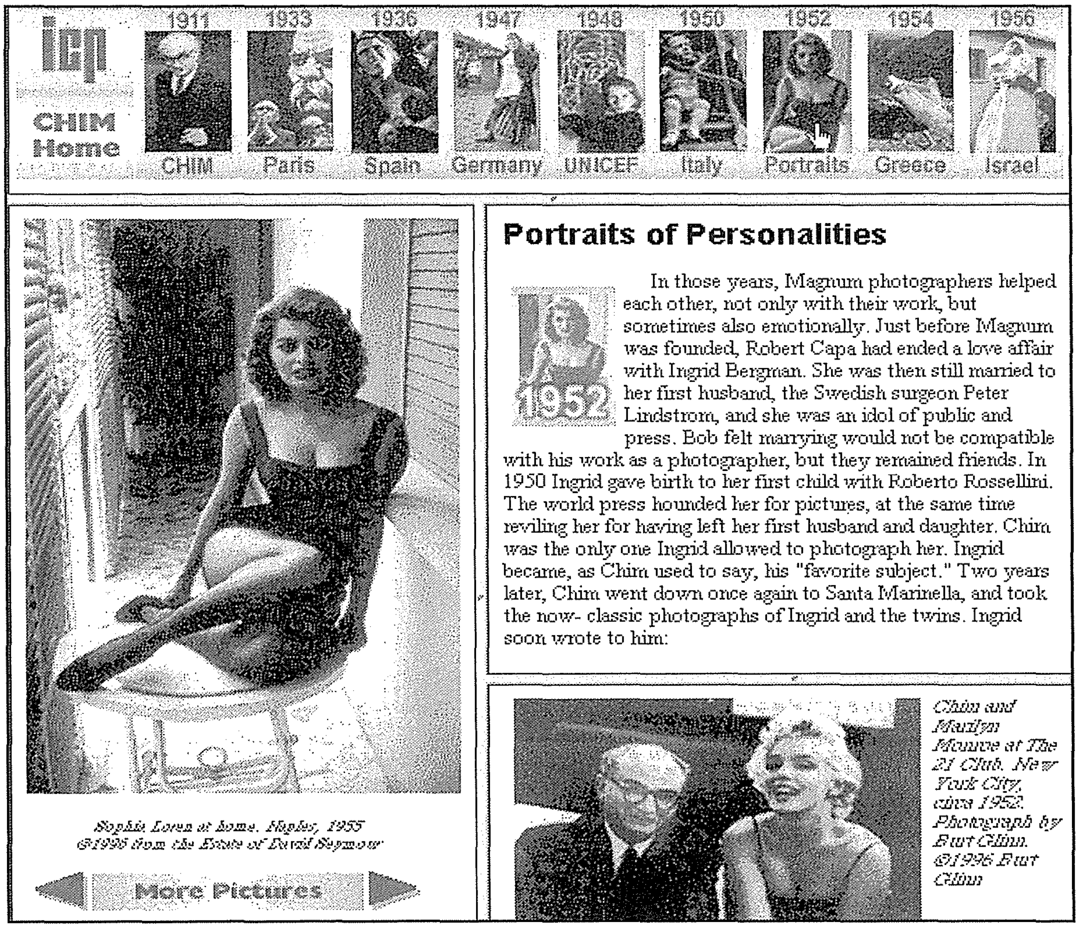

Figure 16.3

Life history of the photographer David Seymour ("Chim"), with a timeline showing eight segments of his work. Presented by the International Center of Photography in New York, NY (http://ww.icp.org/chim/chim2.html).

tion presentation in a self-publishing style, where quality is uncontrolled and structure may be chaotic. Information may be an index to other web sites, or it may be original material. Carefully polished individual life histories (Fig. 16.3) and impressive organizational annual reports are becoming common as expectations and designer experience increase. As commercial usage increases, elegant product catalogs, eye-catching advertisements, and lively newsletters will become the norm. Commercial and scientific publishers will join newspapers (Fig. 16.4) and magazines in providing access to

Figure 16.4

The New York Times compact vertical page layout fits the typical home computer screen (http://www.nytimes.com). (Reprinted by permission © 1997 The New York Times Electronic Media Company.)

information while exploring the opportunities for feedback to editors, discussions with authors, and reader interest groups. Digital libraries of many varieties are appearing (Color Plate C6), but full recognition of their distinct benefits and design features is emerging only slowly. Entertainment web sites are growing as fast as the audience gets online.

A third way of categorizing web sites is by the number of web pages or amount of information that is accessible (Table 16.2): one-page bios and project summaries are small, organization overviews for internal and external use are medium, and airline schedules and the telephone directories are large. Taxonomies of web sites from many perspectives are likely. The Yahoo home page, with its thematic categories, provides a starting point, and it changes as the web grows (Color Plate C5).

A fourth way of categorizing web sites is by measures of success. For individuals, the measure of success for an online resume may be getting a job or making a friend. For many corporate web sites, the publicity is measured in number of visits, which may be millions per day, independent of whether users benefit. For others, the value lies directly in promoting sales of other products, such as movies, books, events, or automobiles. Finally, for access providers who earn fees from hourly usage charges, success is measured by the thousands of hours of usage per week. Other measures include diversity of access as defined by what the number of users is; what their countries of origin are; or whether the users came from university, military, or commercial domains.

Table 16.2

Web-site genres, with approximate sizes and examples

Number of Web pages	Example genres	
1–10	Personal bio Project summary	Restaurant review Course outline
5–50	Scientific paper Conference program	Photo portfolio or exhibit Organization overview
50–500	Book or manual Corporate annual report	City guide or tour Product catalog or advertisement
500–5,000	Photo library Technical reports	Museum tour Music or film database
5,000–50,000	University guide	Newspaper or magazine
50,000–500,000	Telephone directory	Airline schedule
500,000–5,000,000	Congressional digest	Journal abstracts
>5,000,000	Library of Congress	NASA archives

16.5 Users and Their Tasks

As in any user-interface design process, we begin by asking: Who are the users and what are the tasks? Even when broad communities are anticipated, there are usually implicit assumptions about users being able to see and read English. Richer assumptions about the users' age group or educational background should be made explicit to guide designers. Just as automobile advertisements are directed to college-age males, young couples, or mature female professionals, web sites are more effective when directed to specific audience niches. Gender, age, economic status, ethnic origin, educational background, and language are primary audience attributes. Physical disabilities such as poor vision, hearing, or muscle control call for special designs.

Users' specific knowledge of science, history, medicine, or other disciplines will influence design. A web site for physicians treating lung cancer will differ in content, terminology, writing style, and depth from a web site on the same topic for patients. Communities of users might be museum visitors, students, teachers, researchers, or journalists. Their motives may range from fact finding to browsing, professional to casual, or serious to playful.

Knowledge of computers or web sites can also influence design, but more important is the distinction between first-time, intermittent, and frequent users of a web site. First-time users need an overview to understand what the range of services is, what is not available, and what buttons select which actions. Intermittent users need an orderly structure, familiar landmarks, reversibility, and safety during exploration. Frequent users demand shortcuts or macros to speed repeated tasks and extensive services to satisfy their varied needs (Kellogg and Richards, 1995).

Since many applications focus on educational services, appropriate designs should accommodate teachers and students from elementary through university levels. Adult learners and elderly explorers may also get special services or treatments.

Evidence from a survey of 15,000 web users conducted at the Georgia Institute of Technology (Pitkow and Kehoe, 1996) showed that the average age of respondents is 35, the mean household income is above $60,000, and 69 percent are male. A remarkable 82 percent are daily users, and are likely to have a professional connection to computing or education. These profiles have shifted from previous surveys and will probably continue moving toward a closer match with the population at large. Of course, survey response was voluntary, from the web community, so the sample is biased, but the results are still thought provoking. More carefully controlled marketing and user studies are beginning to emerge (Hoffman et al., 1996).

Identifying the users' tasks also guides designers in shaping a web site. Tasks can range from specific fact finding to more unstructured open-ended browsing of known databases, to exploration of the availability of information on a topic (Section 15.1).

The great gift of the web is its support for all these possibilities. Specific fact finding is the more traditional application of computer-based databases with query languages such as SQL, but the web has dramatically increased the capability for users to browse and explore. Equal challenges are to support users seeking specific facts and to help users with poorly formed information needs who are just browsing.

A *planning document* for a web site might indicate that the primary audience is North American high-school environmental-science teachers and their students, with secondary audiences consisting of other teachers and students, journalists, environmental activists, corporate lobbyists, policy analysts, and amateur scientists. The tasks might be identified as providing access to selected LANDSAT images of North America clustered by and annotated with agricultural, ecological, geological, and meteorological features. Primary access might be by a hierarchical thesaurus of keywords about the features (for example, `floods`, `hurricanes`, `volcanoes`) from the four topics. Secondary access might be geographical with indexes by state, county, and city, plus selection by pointing at a map. Tertiary access might be by specification of latitude and longitude.

16.6 Object–Action Interface Model for Web Site Design

The OAI model (Section 2.3) employs a hierarchical decomposition of objects and actions in the task and interface domains (see Fig. 2.2). It can be a helpful guide to web-site designers in decomposing a complex information problem and fashioning a comprehensible and effective web site.

The task of information seeking is complex, but it can be described by hierarchies of task objects and actions related to the information. Then, the designer can represent the task objects and actions with hierarchies of interface objects and actions. For example, a music library might be presented as a set of objects such as collections, which have shelves, and then songs. Users may perform actions such as entering a collection, searching the index to a shelf, and reading the score for a song. The interface for the music library could have hierarchies of menus or metaphorical graphic objects accompanied by graphic representations of the actions, such as a magnifying glass for a search. Briefly, the OAI model encourages designers of web sites to focus on four components in two areas:

1. *Task*
   - *Structured information objects* (for example, hierarchies, networks)
   - *Information actions* (for example, searching, linking)
2. *Interface*
   - *Metaphors for information objects* (for example, bookshelf, encyclopedia)
   - *Handles* (affordances) *for actions* (for example, querying, zooming)

The boundaries are not always clear, but this decomposition into components may be helpful in organizing and evaluating web sites. It was useful in comparing alternatives and analyzing the complex possibilities for the Library of Congress. We shall explore the OAI model and give examples of decompositions of object and actions.

16.6.1 Design of task objects and actions

Information seekers pursue objects relevant to their tasks and apply task-action steps to achieve their intention. Although many people would describe a book as a sequence of chapters and a library as a hierarchy organized by the Dewey Decimal System, books also have book jackets, tables of contents, indexes, and so on, and libraries have magazines, videotapes, special collections, manuscripts, and so on. It would be still harder to characterize the structure of university catalogs, corporate annual reports, photo archives, or newspapers, because they have still less standardized structures and more diverse access paths.

When you are planning a web site to present complex information structures, it helps to have a clear definition of the atomic task objects, and the aggregates that can be combined to build the universe. Atoms can be a birthdate, name, job title, biography, resumé, or technical report. With image data, an atomic object might be a color swatch, icon, corporate logo, portrait photo, or music video.

Information atoms can be combined in many ways to form aggregates, such as a page in a newspaper, a city guidebook, or an annotated musical score. Clear definitions help to coordinate among designers and inform users about the intended levels of abstraction within each project. Information aggregates are further combined into collections and libraries that form the universe of concern relevant to a given set of tasks.

Strategies for aggregating information are numerous. Here is a starting list of possibilities:

- *Short unstructured lists* City-guide highlights, organizational divisions, current projects (and this list)
- *Linear structures* Calendar of events, alphabetic list, human-body slice images from head to toe, orbital swath
- *Arrays or tables* Departure city/arrival city/date, latitude/longitude/time
- *Hierarchies, trees* Continent–country–city (for example, Africa, Nigeria, Lagos), or concepts (for example, sciences, physics, semiconductors, gallium arsenide)
- *Multitrees, faceted retrieval* Photos indexed by date, photographer, location, topic, film type
- *Networks* Journal citations, genealogies, World Wide Web

These aggregates can be used to describe structured information objects. An encyclopedia is usually seen as a linear alphabetical list of articles, with a linear index of terms pointing to pages. Articles may have a hierarchical structure of sections and subsections, and cross references among articles create a network.

Some information objects, such as a book table of contents, have a dual role, since people may read them to understand the topic itself or may browse them to gain access to a chapter. In the latter role, they represent the actions for navigation in a book.

The information actions enable users to follow paths through the information. Most information resources can be scanned linearly from start to finish, but their size often dictates the need for shortcuts to relevant information. Atomic information actions include these:

- Looking for Hemingway's name in an alphabetical list
- Scanning a list of scientific article titles
- Reading a paragraph
- Following a reference link

Aggregate information actions are composed of atomic actions:

- Browsing an almanac table of contents, jumping to a chapter on sports, and scanning for skiing topics
- Locating a scientific term in an alphabetic index and reading articles containing that term
- Using a keyword search in a catalog to obtain a list of candidate book titles
- Following cross reference from one legal precedent to another, repeatedly, until no new relevant precedents appear
- Scanning a music catalog to locate classical symphonies by eighteenth century French composers

These examples and the list in Section 15.1 create a diverse space of actions. Some are learned from youthful experiences with books or libraries, others are trained skills such as searching for legal precedents or scientific articles. These skills are independent of computer implementation; they are acquired through meaningful learning, are demonstrated with examples, and are durable in memory.

16.6.2 Design of interface objects and actions

Since many users and designers have experience with information objects and actions on paper and other traditional media, designing for computer implementation can be a challenge. Physical attributes such as the length of a book or size of a map, which vanish when the information is concealed behind a screen, need to be made apparent for successful use. So web-site designers have the burden of representing the desired attributes of traditional media, but also the opportunity of applying the dynamic power of the computer to support the desired information actions. Successful designers can offer users compelling features that go well beyond traditional media, such as multiple indexes, fast string search, bookmarks, history keeping, comparison, and extraction.

Metaphors for interface objects The metaphoric representation of traditional physical media is a natural starting point: electronic books may have covers, jackets, page turning, bookmarks, position indicators, and so on, and electronic libraries may show varied size and color of books on shelves (Pejtersen, 1989). These may be useful starting points, but greater benefits will emerge as web-site designers find new metaphors and handles for showing larger information spaces and powerful actions.

Information hierarchies are the most frequently represented metaphor, with at least these examples:

- File cabinet with folders and documents
- Book with chapters
- Encyclopedia with articles
- Television with channels
- Shopping mall with stores
- Museum with exhibits

Richer environments include a library with doors, help desk, rooms, collections, and shelves, and the City of Knowledge with gates, streets, buildings, and landmarks. Of course, the information superhighway is often presented as a metaphor, but rarely is it developed as a visual search environment. The metaphor needs to be useful in presenting high-level concepts, appropriate for expressing middle-level objects, and effective in suggesting pixel-level details.

Design of computer-based metaphors extends to support tools for the information seeker. Some systems provide maps of information spaces as an overview to allow users to grasp the relative size of components and to discover what is not in the database. History stacks, bookmarks, help desks, and guides offering tours are common support tools in information environments. Communications tools can be included to allow users to send extracts, to ask for assistance from experts, or to report findings to colleagues.

Handles for interface actions The web-site representation of actions is often conveyed by action handles: the labels, icons, buttons, or image regions that indicate where users should click to invoke an action. Navigation action handles can be a turned page corner to indicate next-page operation, a highlighted term for a link, and a magnifying glass to zoom in or open an outline. Other action handles might be a pencil to indicate annotation, a funnel to show sorting, a coal car to indicate data mining, or filters to show progressive query refinement. Sometimes, the action handle is merely a pull-down–menu item or a dialog box offering rich possibilities. The ensemble of handles should allow users to decompose their action plan conveniently into a series of clicks and keystrokes.

16.6.3 Case study with the Library of Congress

The OAI model is still in need of refinement plus validation, but it may already be a useful guide for website designers and evaluators. It offers a way to decompose the many concerns that arise and provides a framework for structured design processes and eventually software tools. It is not a predictive model, but a guide to designers about how to break a large problem into many smaller ones and an aid in recognizing appropriate features to include in a website. In my experience, designers are most likely to focus on the task or interface objects, and the OAI model has been helpful in bringing out the issues of permissible task actions and visible representations of interface actions.

In the early 1990s, the U.S. Library of Congress staff developed a touchscreen catalog interface to replace the difficult-to-learn command-line interface. In this project, the design was relatively simple; the task objects were the set of catalog items that contained fields about each item. The task actions were to search the catalog (by author, title, subject, and catalog number), browse the result list, and view detailed catalog items. The interface objects were a search form (with instructions and a single data entry field), result lists, brief catalog items, and detailed catalog items. The interface actions were represented by buttons to select the type of search, to scroll the result lists, and to expand a brief catalog entry into a detailed catalog entry. Additional actions, also represented by buttons, were to start a new search, get help, print, and exit. Even in this simple case, explicit attention to these four domains helped to simplify the design.

In the more ambitious case of the Library of Congress website, many potential task objects and actions were identified; more than 150 items were proposed for inclusion on the homepage. The policy and many design decisions were made by a participative process involving the Librarian of Congress, an 18-person Policy Committee, four graphic designers, and staff from many divisions. The resulting design (Color Plate A5) for the hierarchy of task objects is rich, including the catalog, exhibits, copyright information, Global Legal Information, the THOMAS database of bills before Congress, and the vast American Memory resources, but it does not include the books. The exclusion of books is a surprise to many users, but copyright is usually held by the publishers and there is no plan to make the full text of the books available. Conveying the absence of expected objects or actions is also a design challenge.

For brevity we focus on the American Memory component. It will contain 200 collections whose items may be searchable documents, scanned page images, and digitized photographs, videos, sound, or other media. A collection also has a record that contains its title, dates of coverage, ownership, keywords, etc. Each item may have a name, number, keywords, description, etc. The task actions are rich and controversial. They begin with the actions to browse a list of the collection titles, search within a collection, and retrieve an item for viewing. However, searching across all collections is difficult to support and is not currently available. Early analysis revealed that collection records might not have dates or geographic references, thereby limiting the ways that the collection list could be ordered and presented. Similarly, at the next level down, the item records may not contain the information to allow searching by date or photographer name, and restricting search to specific fields is not always feasible.

Continuing within the American Memory component, the interface objects and actions were presented explicitly on the homepage (Color Plate C4). Since many users seek specific types of objects, the primary ones were listed explicitly and made selectable: Prints & Photos, Documents, Motion Pictures, and Sound Recordings. The interface actions were stated simply and are selectable: Search, Browse, and Learn (about using the collections for educational purposes). Within each of these objects and actions, there were further decompositions based on what was possible and what a detailed needs analysis had revealed as important.

At the lowest level of interface objects were the images and descriptive text fields. At the lowest level of interface actions were the navigation, home page, and feedback buttons.

The modest nature of the OAI model means that it can lead to varying outcomes, but it would be unreasonable to assume that there is one best organization or decomposition of a website. In dealing with complex resources and services, it offers designers a way to think about solving their problems.

16.6.4 Detailed design issues

Many web-site design issues are not yet resolved properly. The four-phase framework (see Section 15.2) can provide guidance to web search-engine designers to improve the currently confusing situation. Other issues include query previews to reduce the zero-hit problem while facilitating browsing of large information spaces, and session management to support multiple step plans while providing user assistance.

Query Previews For large collections, especially when they are searched across the network, search actions can be split into two phases: a rapid rough search that previews only the number of items in the result set, followed by a query-refinement phase that allows users to narrow their search and to retrieve the result set (Doan et al., 1996).

For example, in a search for a restaurant (Color Plate C7), the query-preview screen gives users limited choices, with buttons to indicate the type of food (for example, Chinese, French, Indian), double-boxed range sliders to specify average price and time, and maybe a map to specify regions. As users make selections among these attributes, the query-preview bar at the bottom of the screen is updated immediately to indicate the number of items in the result set. Users can quickly discover that there are no cheap French restaurants in downtown New York, or that there are several Caribbean restaurants open after midnight. When the result set is too large, users can restrict their criteria; when the result set is too small, they can alter their plans.

Query previews require database maintainers to provide an updated table of contents that users can download from the server. Then, users can perform rapid searches on their client machines. The table of contents contains the number of items satisfying combinations of attributes, but the size of the table is only the product of the cardinality of the attributes, which is likely to be much smaller than the number of items in the database. With 12 kinds of restaurants, eight regions, and three kinds of charge cards, a simple table of contents would contain only 288 entries. Storing the table of contents burdens users who may have to keep tables of contents (1000 to 100,000 bytes) for each database that they search. Of course, the size of the table of contents can be cut down dramatically if there are simply fewer attributes or fewer values per attribute. The burden of storing tables of contents seems moderate when weighed against the benefits, especially if users search a database repeatedly. The table of contents is only as big as a typical image in a web site, and it can be downloaded for use automatically when Java applets are used.

Query previews are implemented for a complex search on NASA environmental databases. Users of the old system must understand the numerous and complex attributes of the database, which is distributed across eight archival centers. Many searches result in zero hits because users are uncertain about what data are available, and broad searches take many minutes yet yield huge and unwieldy result displays. The query preview uses only three parameters: locations (clustered into 15 geographic regions), 171 scientific parameters (soil type, ocean temperature, ozone, and so on), and dates (clustered into 10 1-year groups) (Fig. 16.5). There is thus a total of $15 \times 171 \times 10 = 25650$ data values in the table of contents. In the prototype, users could quickly discover that the archive held no ozone measurements in Antarctica before 1985. Once a reasonable-sized result set is identified, users can download the details about these data sets for the query-refinement phase.

Session Management Each search or jump is one action to accomplish a task within a session. When sessions cover many complex tasks, it is helpful

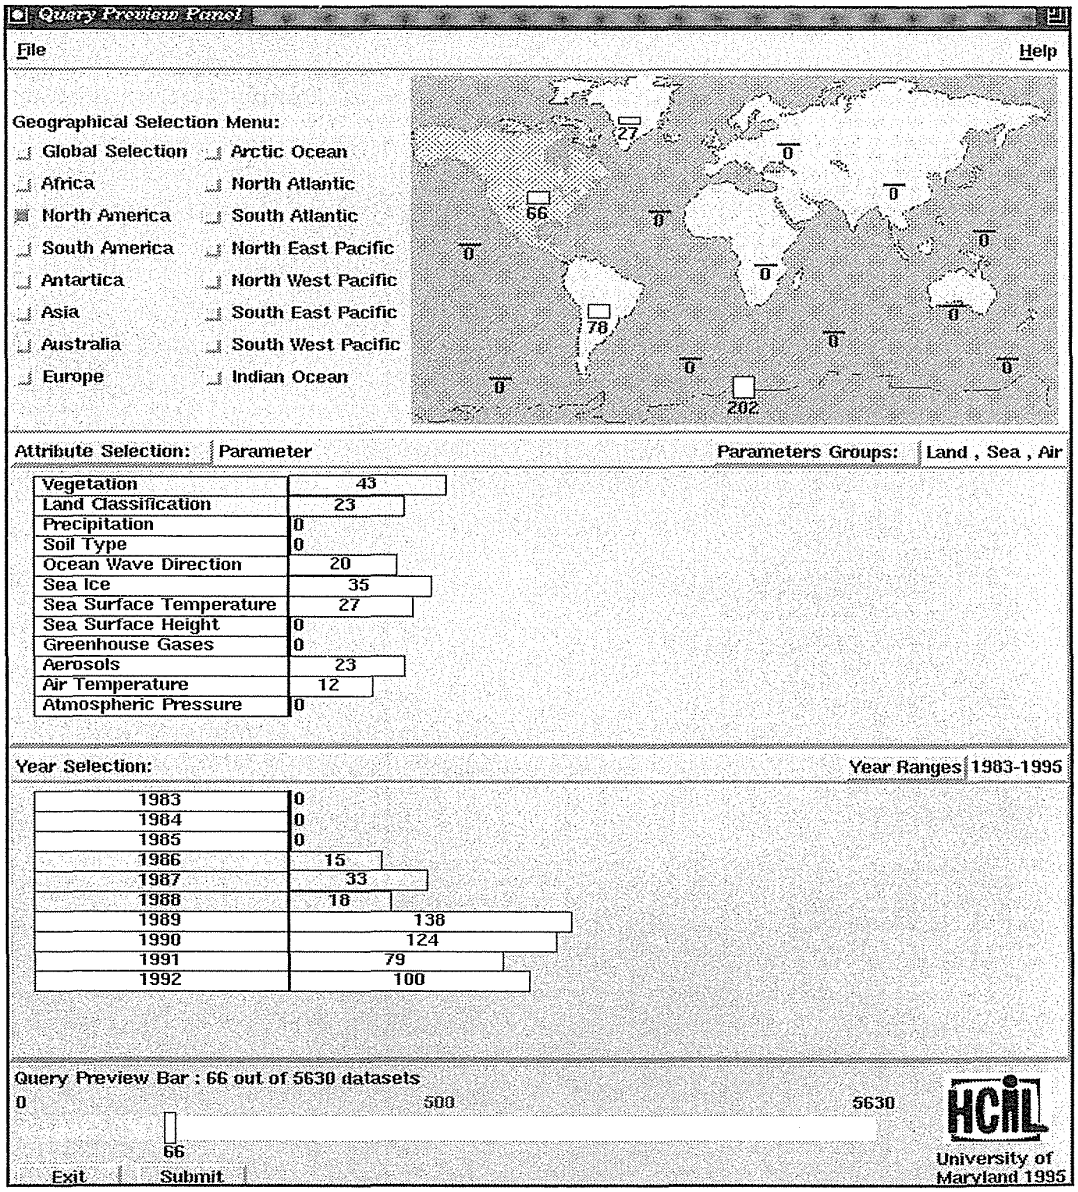

Figure 16.5

NASA query preview. The system applies this technique to a complex search for professional scientists. The set of more than 20 parameters is distilled to three, thus helping to speed the search and to reduce wasted effort. Users select values for the parameters and immediately see the size of the result bar on the bottom, thus avoiding zero-hit and mega-hit queries. (Doan et al., 1996) to have a history of the session, so that users can review their progress, return to key actions, learn from their failures, and save or edit or reapply previous actions (possibly on related databases). Other useful services allow users to annotate the objects that they have retrieved, and to extract text or images for use in their own projects. Communication with the owners or designers or maintainers of a web site and with other users should also be supported. Web-site users should be encouraged to comment about errors or suggest improvements.

No matter how well designed a web site is, there will always be a need for online and human assistance. The designer's goal should be to reduce the need for assistance, but with ever-greater numbers of users with ever-greater expectations, careful planning and trained personnel are needed. Online tutorials, descriptions of interface objects and actions, files with frequently asked questions (FAQs), electronic-mail help desks, and telephone service (possibly for a fee) should be made available.

16.6.5 Web-Page design

According to the OAI model, web-page designers should begin by identifying tasks in terms of information objects and actions. Then, designers can present interface metaphors for information objects accompanied by handles for actions. Success also requires wise choices during detailed page design to show objects (for example, menus, search results, fonts, colors) and to invoke actions (for example, button press, selection from a list). These visible design elements are often the most discussed aspects of design, and are the ones most directly implemented by HTML or Java coding. Initial subjective satisfaction is strongly determined by these surface features; therefore, they deserve intense attention (Horton et al., 1996).

Compactness and branching factors The most discussed issues are page length and number of links (branching factor). An extremely long page with no links is appealing only if users are expected to read the entire text sequentially. They rarely are, so some form of home or index page to point to fragments is necessary. Meaningful structures that guide users to the fragments that they want is the goal, but excessive fragmentation disrupts people who wish to read or print the full text. As the document and web site grow, the number of layers of index pages can grow as well, and that poses a severe danger. A higher branching factor is almost always preferred for index pages, especially if it can save an extra layer that users must traverse. The extra layers are more disorienting than are longer index pages. In a redesign for the Library of Congress home page (http://www.loc.gov) (Color Plate A5), the seven links to general themes were replaced with a compact display with 31 links to specific services. The Yahoo home page has almost 100 links in a compact two-column presentation.

Within a page, compact vertical design to reduce scrolling is recommended. Although some white space can help to organize a display, often web pages contain harmful dead space that lengthens the page without benefit to users. A typical mistake is to have a single left-justified column of links that leaves the right side of the display blank, thus forcing extra scrolling and preventing users from gaining an overview. A second common mistake is to use excessive horizontal rules or blank lines to separate items.

Sequencing, clustering, and emphasis Within a page—especially the highly visible home page of an organization—designers must consider carefully the sequencing, clustering, and emphasis for objects. Users expect the first item in a page to be an important one and are likely to select it. Clustering related items shows meaningful relationships. More important items can be emphasized with large fonts, color highlights, and surrounding boxes. The Library of Congress home page emphasizes the American Memory collections by placing them first and giving them a large fraction of the space. Public services such as the catalog and THOMAS (for searching legislation) are clustered in the center, and library services are clustered on the right side.

Support for universal access Designers must accommodate small and large displays, monochrome and color, slow and fast transmission, and various browsers that may not support desired features. The pressure for lowest-common-denominator design is often outweighed by the desires to assume larger displays, to use more detailed and more numerous graphics, to support Java applets, and to employ newer browser features. Fortunately, balanced approaches that enable users to indicate their environment and preferences are possible. Since many key design decisions involve task issues, several versions of the interface can be developed for relatively small incremental costs.

Providing text-only versions for users with small displays and low-bandwidth access is likely to be strongly recommended for many years to come. Users using low-cost devices, users in developing countries with poor communication infrastructure, users wanting low-bandwidth wireless access, users with small personal display devices, and handicapped users constitute a large proportion of the potential users.

The great disparity in transmission speeds (low-speed modems at 1200 baud to direct-connection lines at 4 megabits per second) has compelled many designers to build two versions of a web site: text only and graphical. Another solution is to display the textual components first, and to fill in the graphics as time permits. The use of image thumbnails that can be optionally expanded is another appropriate accommodation. Opinions are split over the use of low-resolution graphics that become enhanced as time permits.

Accommodating diverse users should be a strong concern for most designers, since it enlarges the market for commercial applications and provides democratic access to government services. Web sites should be tested an gray-scale displays, low-bandwidth transmission lines, and small displays. In addition, access by way of telephone or voice input–output devices will serve handicapped users and enlarge access. Access to web sites might also come from wristwatch projection displays, wallet-sized pocket PCs, or personal video devices mounted on eyeglasses.

Pointing devices could be the traditional mouse, trackball, trackpoint, touchpad, touchscreen, or eyegaze-detection devices. Cursor-control arrow keys could be used to support jumping among highlighted items.

Good graphical design Many personal web sites are developed by individuals who learn a modest amount of HTML or use a graphical interface to generate HTML. Simple web sites can be created successfully with these tools, but doing innovative and effective page layout for a large web site requires as much care and skill as does laying out a newspaper, magazine, or book. Page layout is a well-developed topic for graphic designers, whose expertise is vital for innovative and effective designs (Cotton and Oliver, 1993; McAdams, 1996; Weinman, 1996). Grid layouts and consistent structure help to guide the reader. Distinctive headings and graphics signal boundaries and provide familiar landmarks during navigation on the first visit and return visits. Indexes and shortcuts give frequent users paths for rapid traversal.

An hour spent browsing web sites will reveal diverse graphic-design philosophies. The poster designers use centered titles, large graphics, ample white space, and a small number of visually striking buttons, sometimes from the *Wired* magazine school of garish colors and extreme imagery. Book designers use left-justified titles, a few small graphics, dense text that goes on for many pages, and numerous text links. Newspaper designers start several stories on the home page, each with its own heading, font, column, inset photos, and continuations. GIF-heads are eager to show their scanned photos, art, scientific images, or logo, and place these graphics on the page with little care for captions, layout, or the burden on users who have low-bandwidth access. Hypertext fanatics chop up documents into paragraph-sized chunks or smaller, and put as many links as possible per sentence. Traditionalists simply put a lengthy text in a single file and expect users to scroll happily and linearly.

Each design philosophy is meant to appeal to certain users and to support certain tasks. Abuses that hinder task completion are likely to fade over time. Page layouts are likely to favor convenient online browsing, but special layouts may be needed to produce effective printed versions.

Traditional graphic-design rules often apply in the web environment. Large fonts or boldface type typically indicate major headings, and medium fonts can signal sub headings. Text is best left as black on white or gray, enabling links to be highlighted by color or underlining. A graphic logo is typical for an organization home page, and, as users move down in a hierarchy, moderate- or small-sized logos can indicate location in the hierarchy. Four sizes of logos are probably as many as most users can grasp quickly. For different branches of a large hierarchy, variant logos, and color coding of banners or backgrounds can be effective, with a limit of six to eight variations. These recommendations emerge from graphic-design books for paper documents; new opportunities will appear for electronic documents.

Navigation support In a paper book, the reader's progress is easily seen. Since this is not available online, innovative substitutes are appearing. The simple approach of indicating `page 171 of 283` can be effective, but various analog progress indicators, such as scroll bars and page bars, are emerging. More elaborate indicators, such as a tree or network diagram, sometimes called a *site map*, help to orient users in larger sites (see Fig. 7.7). Dynamic indicators that respond to mouse cursor placement by opening up a hierarchy or popping up detailed information in a small window are still novel. Animated indicators that reveal underlying structures or offer more details are likely to emerge, along with auditory feedback, three-dimensional displays, and rich information visualizations.

Although scroll bars are the primary navigation tool because they provide a simple and standard mechanism, a paging strategy using page bars (a scroll bar with discrete jumps) is cognitively less demanding because users have a clearer sense of position in a document. Designers can make use of tops and bottoms of pages to provide navigational cues (headers, footers, page numbers). Users become familiar with a document by remembering a photo or figure at the top of a certain page. Unfortunately, this strategy is undermined by the wide variation in screen sizes, so designers have to commit to a specific size, such as 640 × 480 pixels, and then users of larger or smaller screens have to accommodate to the standard. When designers can assume a larger screen and resolution-indepen-

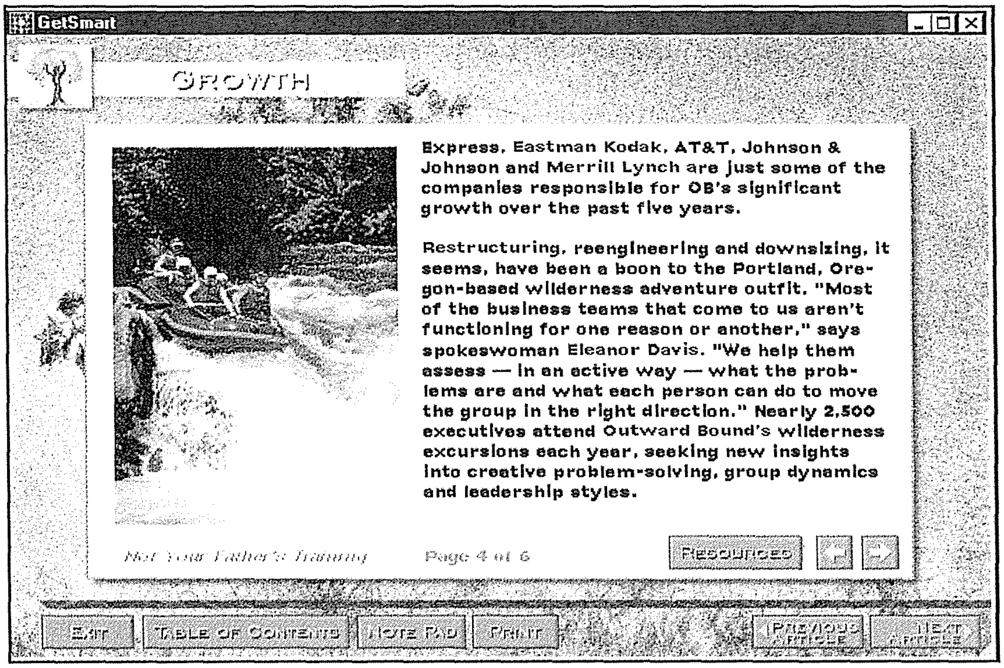

Plate C1: An electronic magazine Get Smart, built with Hyperties, that allows integration of images, video, and sound with embedded links to related topics. (Used with permission of Cognetics Corp., Princeton Junction, NJ.)

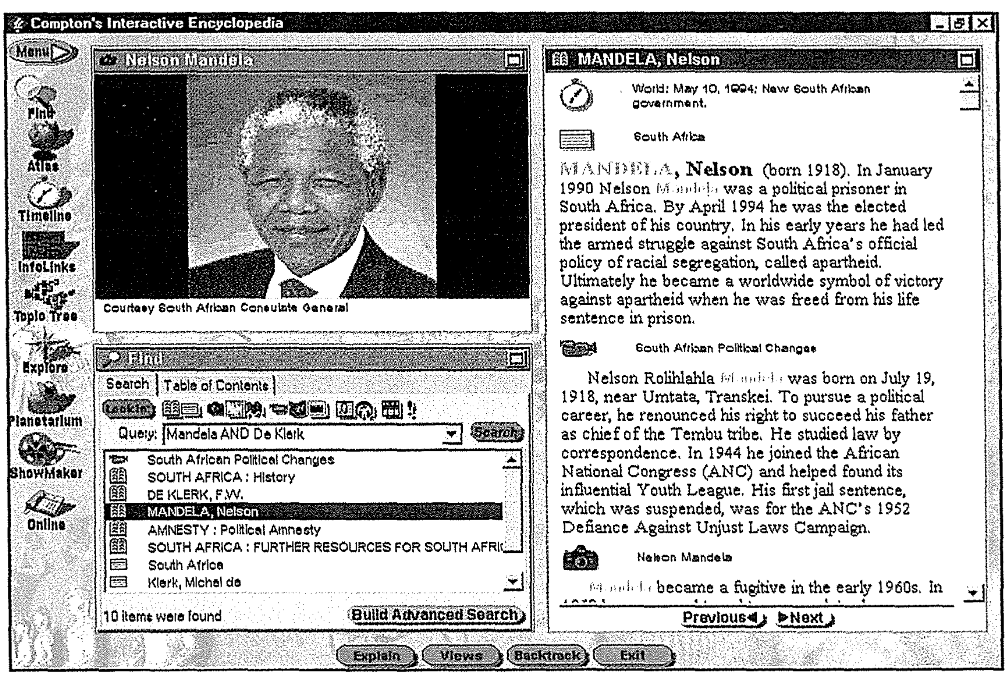

Plate C2: Compton's Encyclopedia uses multimedia and multiple windows to enrich the presentation and support browsing.

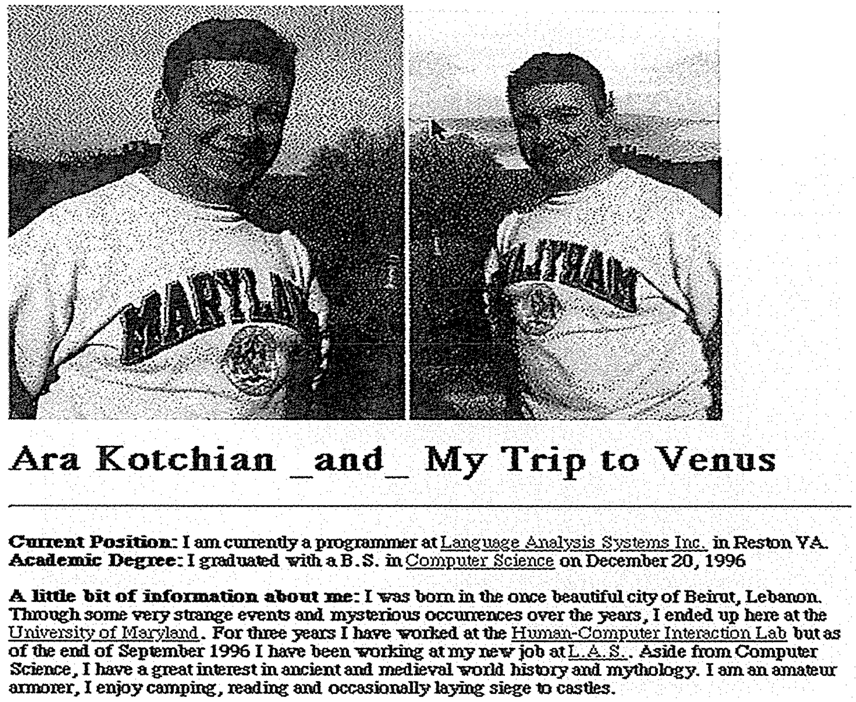

Plate C3: One-page personal biography of Ara Kotchian, a student at the University of Maryland. (Used with permission.) (http://www.cs.umd.edu/projects/hcil/People/ara/index.html).

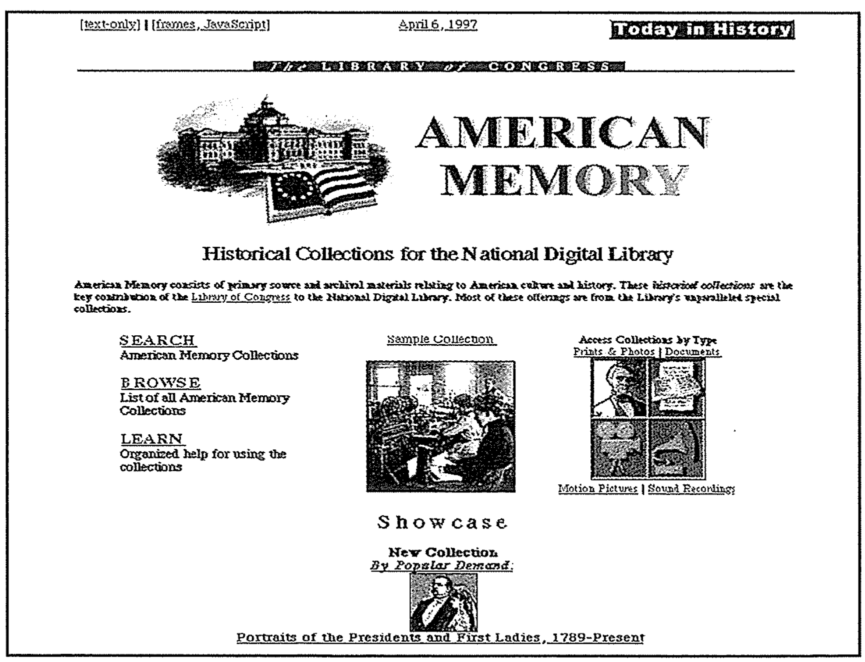

Plate C4: American Memory home page from the Library of Congress, which will offer more than 5 million images, texts, videos, and so on by the year 2000. (http://lcweb2.loc.gov/ammem).

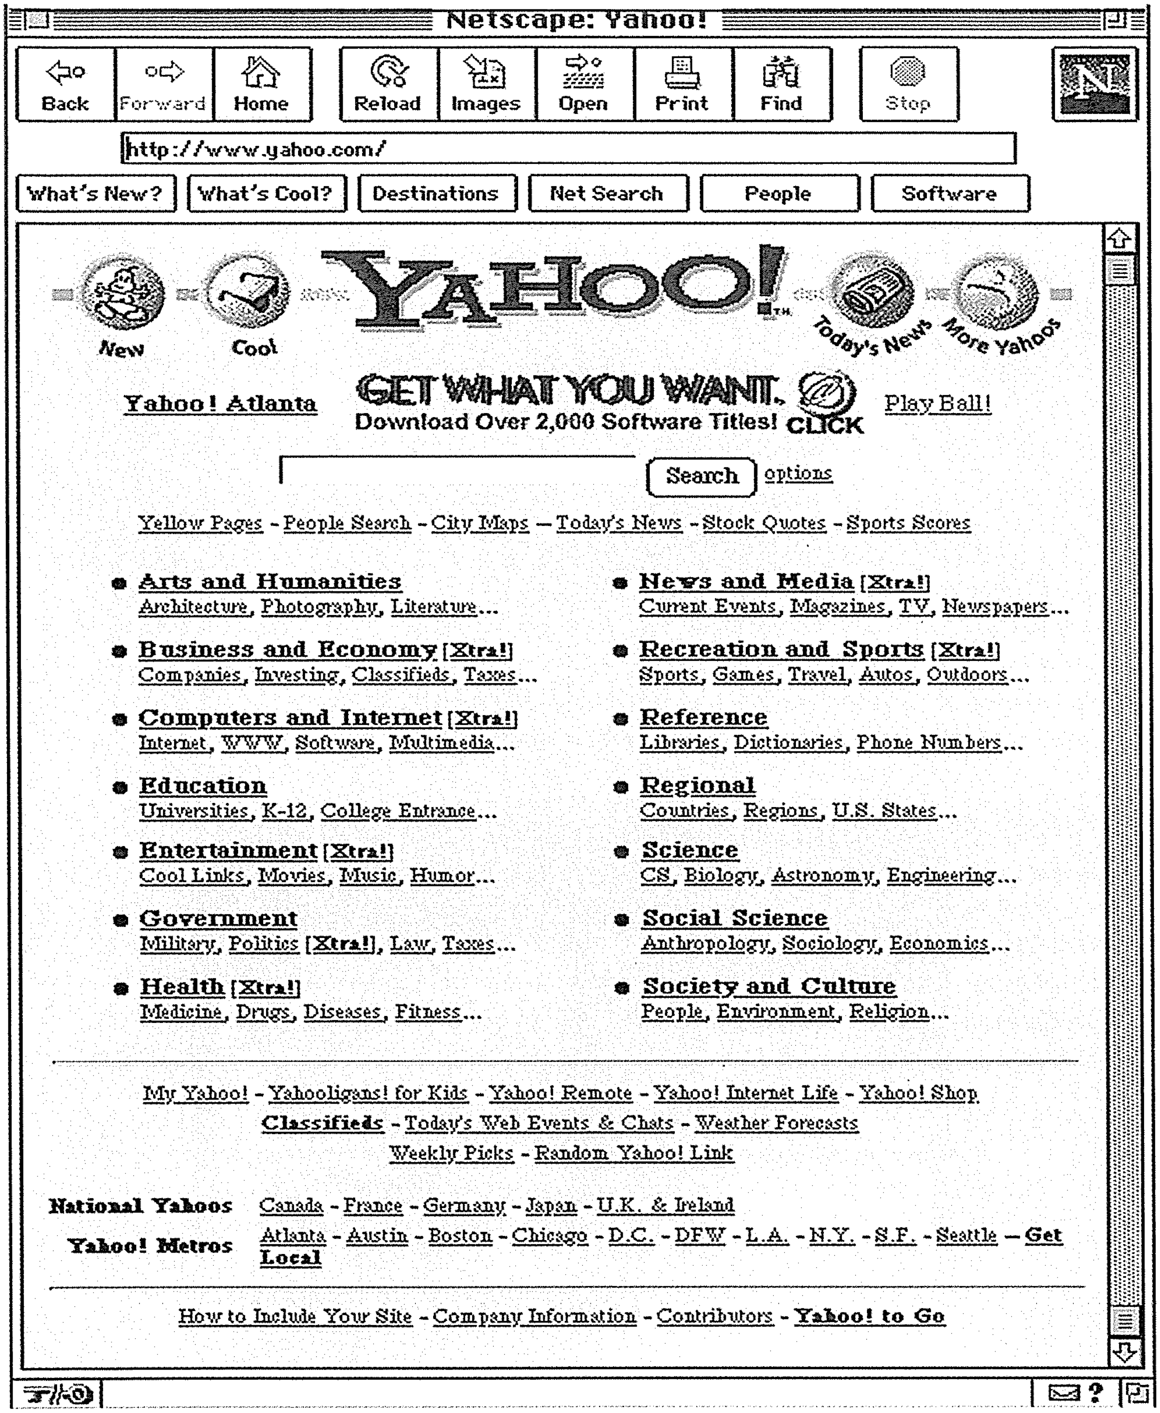

Plate C5: Yahoo index page showing a 14-item thematic categorization with 51 second-level links, and more than 30 other links. (http://www.yahoo.com) (Used with permission of Reuters, Inc.)

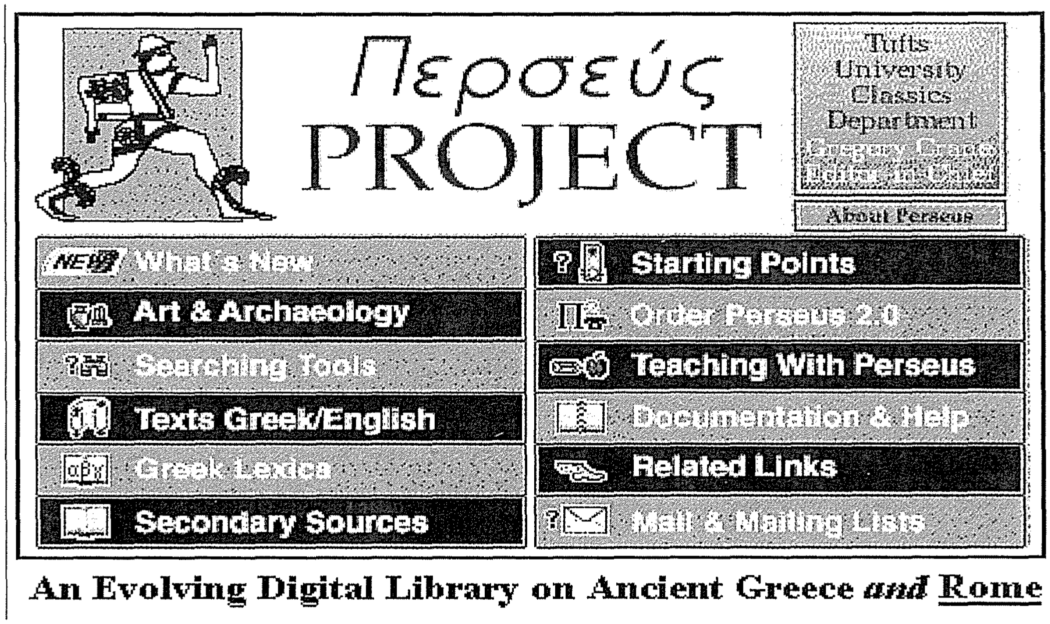

Plate C6: Perseus digital library, which contains ancient Greek texts in original and English forms with maps, photos, architectural plans, vases, coins, and so on for students and researchers. (http://www.perseus.tufts.edu).

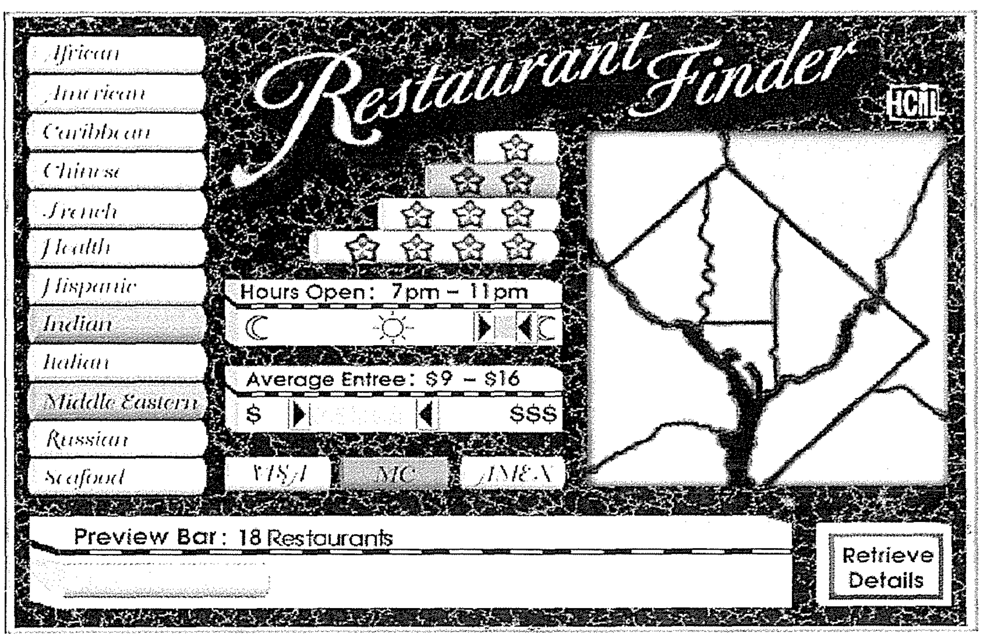

Plate C7: Restaurant finder, a system that demonstrates query preview. Users can quickly adjust the parameters and see the effect on the size of the preview bar at the bottom. Zero-hit or mega-hit results are visible immediately, so users can always be sure that their search will produce an appropriate number of results. (Graphic design by Teresa Cronnell.) (Doan et al., 1997.)

dent layouts as standard, then the shift toward page orientation will be more common.

16.6.6 Testing and maintenance of web sites

Usability testing is recommended for all user-interface projects; in addition, there are special needs for web-site testing. As always, the questions of who the users are and what the tasks are guide designers. The testing should be done with representatives of each of the primary user communities and of as many of the secondary communities as time and money allow (Nielsen, 1995). Users of various ages, genders, and ethnic backgrounds, as well as international users, may be included. The task-frequency list developed during needs assessment provides guidance for construction of test tasks.

Users should be tested in realistic settings that resemble the office or home environments. The number of users and length of testing will depend on the project's importance. Various screen sizes and transmission speeds should be tested. Voice access should be tested for handicapped and other users. Since browsers have such varied features, testing should be done with several of them.

After early in-house testing with limited numbers of users, a more extensive in-house test might be conducted. Then, intensive field testing can begin, before a public announcement. A phased roll-out process will protect against disaster, improve quality, and ensure the highest satisfaction by the largest number of users.

The web-site developer's work is never done. The more successful a site is, the more opportunities there are for revision and improvement. Software logs should capture the frequency of use for each page, or at least for each component of a database. Such statistics can reveal patterns that provide guidance for improving a web site. If some components are never accessed, then they can be removed, or the references to these components can be improved to draw attention. Logging software can also reveal patterns of use over a month, week, or day, and can indicate the paths users take to arrive at and to traverse a web site.

In addition to using automated logging, web-site maintainers can solicit feedback from users by electronic mail or survey questionnaires embedded in the web site. Knowledge of user demographics and motivations may be helpful in refining a web site. To get in-depth understanding, web-site maintainers may interview users individually by telephone or personally, or conduct focus-group discussions among users.

User expectations and organizational policy guide the rate of change of the web-site contents and interface. Some web sites are stable, and users depend on the permanent availability of the contents; libraries, government archives, and online journals are examples. Other sites are volatile and are expected to change hourly, daily, weekly, or seasonally; weather information, newspapers, magazines, and train schedules are examples.

16.7 Practitioner's Summary

Linked information in text, graphics, image, sound, animation, and video formats can be used in commercial projects that satisfy the Golden Rules of Hypertext. Effective hypermedia products follow basic principles of user-interface design, but place a greater emphasis on content organization and presentation.

Careful web-site design makes the difference between a must-see, top-10 site and a worst–web page award. Specifying the users and setting the goals come first, followed by design of information objects and actions. Next, designers can create the interface metaphors (bookshelf, encyclopedia, shopping mall) and the handles for actions (scrolling, linking, zooming). Finally, the web-page design can be created in multiple visual formats and international versions, with access provided for handicapped or otherwise special readers. Every design project, including web-site development, should be subjected to usability testing and to other validation methods. Monitoring of use should guide revision.

16.8 Researcher's Agenda

Hypertext, hypermedia, multimedia, and the World Wide Web are still in the Ford Model T stage of development. Strategies for blending text, sound, images, and video are in need of refinement, and effective rhetorics for hypermedia are only now being created. Who will be the first to write the Great American Hypernovel or Hypermystery? Many results from other user-interface topics—such as menu selection, direct manipulation, and screen design—can be applied to web-site design. On the other hand, the novel communities of users, innovative databases, ambitious services, emphasis on linking and navigation, and intensive use of graphics present fresh challenges and rich opportunities to researchers to validate hypotheses in this environment. Theories of information structuring are emerging, as are standards for representing traversal actions. The creative frenzy on the web is likely to present new opportunities for design research for many years to come.

Controlled experimental studies are effective for narrow issues, whereas field studies, data logging, and online surveys are attractive alternative research methods in the wide-open web. Focus groups, critical-incident studies, and interviews may be effective for hypothesis formation. Other opportunities include sociological studies about impact of web use on home or office life, and political studies of web-use influence on democratic processes. Broader concerns—such as copyright violation, invasion of privacy, pornography, or criminal activity—merit attention as the impact of the World Wide Web increases. We can influence the direction and societal impact of technology only if we have the scientific foundation to understand the issues.

World Wide Web Resources WWW

It should not be a surprise that World Wide Web contains numerous documents about the World Wide Web, including style guides for authors, numerous navigation tools, research and survey reports, and ample discussion.

http://www.aw.com/DTUI

References

Berners-Lee, Tim, Cailliau, Robert, Luotonen, Ari, and Nielsen, Henrik F., and Secret, Arthur, The World Wide Web, *Communications of the ACM*, 37, 8 (1994), 76–82.

Bush, Vannevar, As we may think, *Atlantic Monthly*, 76, 1 (July 1945), 101–108.

Chimera, R. and Shneiderman, B., Evaluating three user interfaces for browsing tables of contents, *ACM Transactions on Information Systems*, 12, 4 (October 1994), 383–406.

Conklin, Jeff, Hypertext: A survey and introduction, *IEEE Computer*, 20, 9 (September 1987), 17–41.

Cotton, Bob and Oliver, Richard, *Understanding Hypermedia: From Multimedia to Virtual Reality*, Phaidon Press, London, U.K. (1993).

Doan, Khoa, Plaisant, Catherine, and Shneiderman, Ben, Query previews for networked information services, *Proc. Advances in Digital Libraries Conference*, IEEE Computer Society, Los Alamitos, CA (May 1996), 120–129.

Engelbart, Douglas, Authorship provisions in AUGMENT, *Proc. IEEE CompCon Conference* (1984), 465–472.

Flynn, Laurie, Making searches easier in the web's sea of data, *The New York Times* (October 2, 1995).

Halasz, Frank, Reflections on NoteCards: Seven issues for the next generation of hypermedia systems, *Communications of the ACM*, 31, 7 (July 1988), 836–852.

Hoffman, Donna L., Kalsbeek, William D., and Novak, Thomas P., Internet and web usage in the U.S., *Communications of the ACM*, 39, 12 (December 1996), 36–46.

Horton, William, Taylor, Lee, Ignacio, Arthur, and Hoft, Nancy L., *The Web Page Design Cookbook*, John Wiley and Sons, New York (1996).

IBM, World Wide Web Design Guidelines, http://www.ibm.com/ibm/hci/guidelines/web/web_design.html (April 1997).

Isakowitz, Tomas, Stohr, Edward A., and Balasubramanian, P., RMM: A methodology for hypermedia design, *Communications of the ACM*, 38, 8 (August 1995), 34–44.

Jones, Trish and Shneiderman, Ben, Evaluating usability for a training-oriented hypertext: Can hyper-activity be good?, *Electronic Publishing 3*, 4 (November 1990), 207–225.

Kellogg, Wendy A. and Richards, John T., The human factors of information on the internet, In Nielsen, Jakob (Editor), *Advances in Human–Computer Interaction*, Volume 5, Ablex, Norwood, NJ (1995), 1–36.

Koved, Larry and Shneiderman, Ben, Embedded menus: Selecting items in context, *Communications of the ACM*, 29, 4 (April 1986), 312–318.

Lemay, Laura, *Teach Yourself Web Publishing with HTML in a Week*, Sams Publishing, Indianapolis, IN (1995).

Lynch, Patrick J., Yale, *University C/AIM WWW Style Guide*, http://info.med.yale.edu/caim/StyleManual_Top.HTML (September 5, 1995).

Marchionini, Gary, *Information Seeking in Electronic Environments*, Cambridge University Press, U.K. (1995).

Marchionini, Gary and Shneiderman, Ben, Finding facts and browsing knowledge in hypertext systems, *IEEE Computer*, 21, 1 (January 1988), 70–80.

McAdams, Melinda, Information design and the new media, *ACM interactions*, II.4, (October 1995), 38–46.

Nielsen, Jakob, *Multimedia and Hypermedia*, Academic Press, San Diego, CA (1995).

Nielsen, Jakob, A home-page overhaul using other web sites, *IEEE Software*, 12, 3 (May 1995), 75–78.

Nielsen, Jakob, Using paper prototypes in home-page design, *IEEE Software*, 12, 4 (July 1995), 88–97.

Nielsen, Jakob, Sun studies of WWW design, http://www.sun.com/sun-on-net/www.sun.com/uidesign.

Pejtersen, A. M., A library system for information retrieval based on a cognitive task analysis and supported by an icon-based interface, *Proc. ACM SIGIR Conference* (1989), 40–47.

Pitkow, Jim and Kehoe, Colleen, GVU's Fourth WWW User Survey, www-survey@cc.gatech.edu (1996).

Plaisant, Catherine, Guide to Opportunities in Volunteer Archaeology: Case study of the use of a hypertext system in a museum exhibit. In Berk, Emily and Devlin, Joseph, (Editors), *Hypertext/Hypermedia Handbook*, McGraw-Hill, New York (1991), 498–505.

Rivlin, Ehud, Rotafogo, Rodrigo, and Shneiderman, Ben, Navigating in hyperspace: Designs for a structure-based toolbox, *Communications of the ACM*, 37, 2 (February 1994), 87–96.

Shneiderman, Ben (Editor), Hypertext on Hypertext, Hyperties disk with 1Mbyte data and graphics incorporating July 1988 *CACM*, ACM Press, New York, NY (July 1988).

Shneiderman, Ben, Reflections on authoring, editing, and managing hypertext. In Barrett, E. (Editor), *The Society of Text*, MIT Press, Cambridge, MA (1989), 115–131.

Shneiderman, Ben, Brethauer, Dorothy, Plaisant, Catherine and Potter, Richard, Three evaluations of museum installations of a hypertext system, *Journal of the American Society for Information Science*, 40, 3 (May 1989), 172–182.

Shneiderman, Ben and Kearsley, Greg, *Hypertext Hands-On! An Introduction to a New Way of Organizing and Accessing Information*, Addison-Wesley, Reading, MA (1989).

van Dam, Andries, Hypertext 87: Keynote Address, *Communications of the ACM*, 31, 7 (July 1988), 887–895.

Weinman, Lynda, *Designing Web Graphics*, New Riders Publishing, Indianapolis, IN (1996).

Yankelovich, Nicole, Meyrowitz, Norm, and van Dam, Andries, Reading and writing the electronic book, *IEEE Computer*, 18, 10 (October 1985), 15–30.

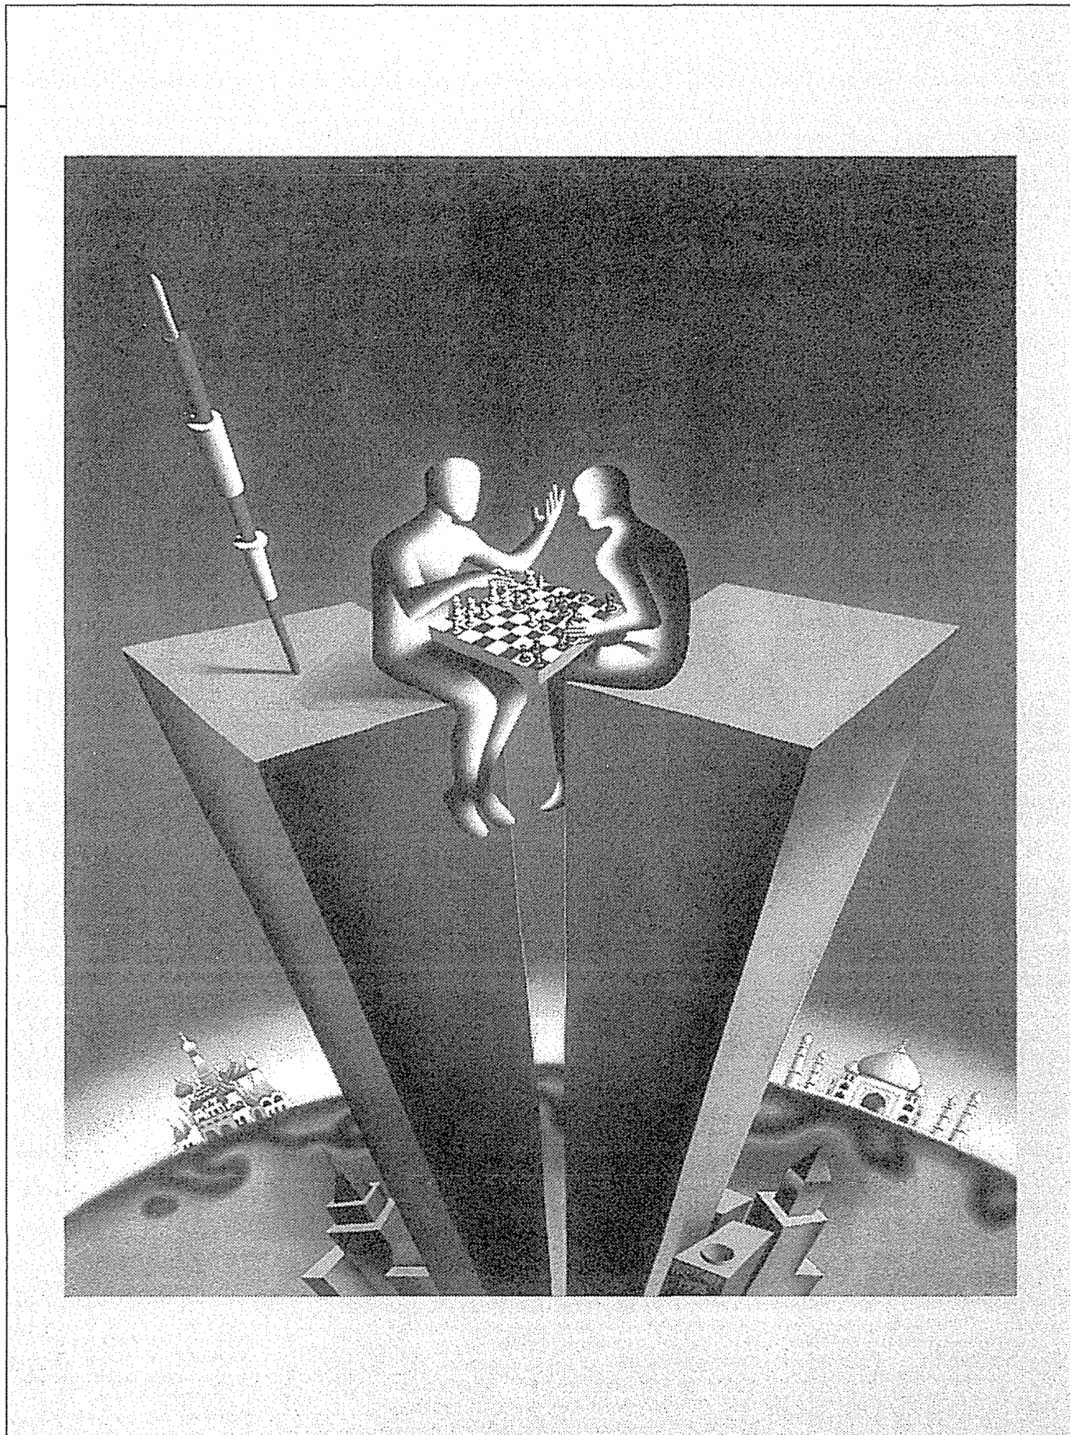

Mark Kostabi, *Grandmasters*, 1995

AFTERWORD

Societal and Individual Impact of User Interfaces

The machine itself makes no demands and holds out no promises: it is the human spirit that makes demands and keeps promises. In order to reconquer the machine and subdue it to human purposes, one must first understand it and assimilate it. So far we have embraced the machine without fully understanding it.

Lewis Mumford, *Technics and Civilization,* **1934**

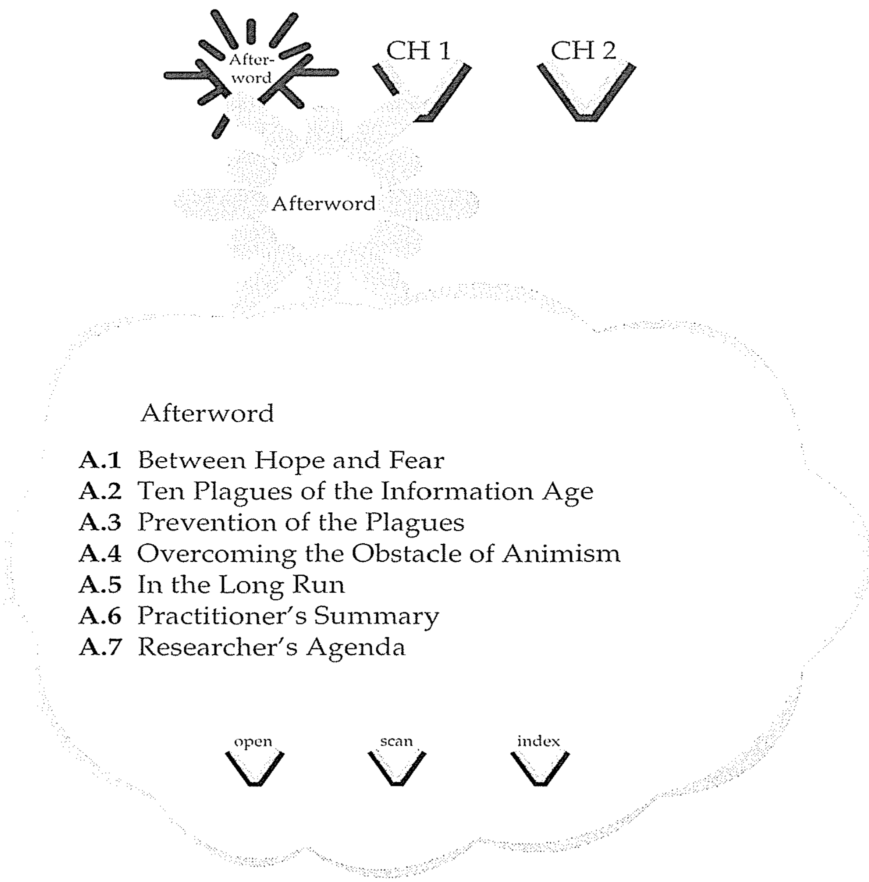

A.1 Between Hope and Fear

Hope is a vital human emotion, stimulated by the desire to make life better and infused with the belief that change is possible. Hope has a strong rational component that shapes plans and reasons about possible outcomes, but hope depends on passion for forward movement.

Deeply held hopes can invigorate other people to join in purposeful action. Martin Luther King's "I Have a Dream" speech is inspirational because of its image of racial harmony. Similarly, John F. Kennedy's vision of a human setting foot on the moon helped to bring about that event. His passion provided the propelling force for the rational scientific work that followed.

Often, hope must overcome resistance—the *fear* that action will fail or leave us worse off. Fear can be a terrifying barrier to change, but also can create energy for action. Confronting fears and summoning the courage to press forward requires self-confidence and a determination to succeed. Because of these challenges, people and civilizations are often remembered for their deep hopes, or, in the words of Ezra Pound, "One measure of a civilization, either of an age or of a single individual, is what that age or person really wishes to do."

Computing professionals can reflect proudly on 50 years of accomplishments, and it is appropriate for us to consider our deep hopes for the next 50 years. Computing has grown into a worldwide infrastructure that touches every country and soon may touch every individual on the planet. But what are our deep hopes for the next 50 years? If our hopes inspire action, our profession will be appreciated for contributing to a better society (Brooks, 1996; Shneiderman, 1995, 1990).

Through the half-century of our profession, visionaries have inspired constructive development. In the 1940s, Vannevar Bush envisioned memex, a desk with microfilm libraries to extend memory by accessing vast resources of patents, scientific papers, or legal citations (Bush, 1945). J.C.R. Licklider carried the digital-library idea into the world of electronic computers and recognized the potential for teleconferencing to bring people closer together (Licklider, 1965). Douglas Engelbart envisioned computers as symbol manipulators that could augment human intellect (Engelbart, 1968). He created an ambitious workstation with a mouse, a chorded keyboard, an outliner, and links across documents that he demonstrated at the Fall Joint Computer Conference of 1968. Later visionaries brought us personal computers, networks, electronic mail, GUIs, and more. These innovations helped to launch the modern computer industry; finding the next breakthrough is our next challenge.

An obvious vision of hope is by *technology extrapolation*, which posits that advances in technology are in themselves beneficial to society. This approach leads to dreams of gigahertz processors producing rapid user-controlled three-dimensional animations on gigapixel displays. Technology extrapolation also suggests terabyte hard disks and web spaces with petabytes of information at our fingertips. Progress is relatively easy to recognize if we follow technology extrapolation, but a more challenging path is to consider what technologies we want to change ourselves and our civilization.

A more elaborate form of technology extrapolation is to dream of intelligent agents, speech interaction, or information at our fingertips (Gates, 1995; Negroponte, 1995). These goals are technology-oriented, but they are not directly linked to clear societal benefits, such as world peace, improved health care, or civil rights. Linking grand goals to realistic scenarios for accomplishing them takes impassioned imagination combined with scientific rigor.

Let's start with imagination. The fisherman who rubbed Aladdin's lamp evoked a genie who offered three wishes. If rubbing your keyboard could produce a modern digital genie, what wishes would you choose to shape the

> **Web inspiration:** In trying to fathom the link between the emotional quality of hope and the rational world of technology, I ventured onto the World Wide Web. A quick search revealed an encouraging pattern: more than one million entries for `hope`, and only ⅓ million for `fear`. Some probing yielded a web site from a recent international conference in Japan on "The Future of Hope" (http://iij.asahi.com/paper/hope/english/index.html). Along with many speakers, Elie Wiesel called for recognition of past and current suffering. The closing *Hiroshima Declaration* stressed continued reduction in nuclear weapons and support for human and civil rights. It mentioned the potential for "revolutionary technologies" that "offer ever more opportunities to bring people and leaders together in dialogue and thus resolve their conflicts."

future? After much reflection, my choices are for universal access to computing technology, universal medical records, and universal educational support.

A.1.1 Universal access to computing technology

My first hope is for universal access, in which progress is defined by the percentage of the population with convenient low-cost access to specific World Wide Web services, such as electronic mail, distance education, or community networks (Anderson et al., 1995). Providing electricity, hardware, and communications is just the beginning. Applications and services will have to be reengineered to meet the differing needs of the many still forgotten users. We must think about how electronic mail can be reshaped to accommodate unskilled writers and readers, while helping to improve their skills. How can job training and hunting be organized to serve those people with currently poor employment skills and transient lifestyles? How can services such as voting registration, motor-vehicle registration, or crime reporting be improved if universal network access is assumed?

Perhaps we can begin by redesigning interfaces to simplify common tasks. We can provide novel training and help methods so that using a computer is a satisfying opportunity, rather than a frustrating challenge. Evolutionary learning with level-structured interfaces would allow first-time users to succeed with common tasks and would provide a growth path to reveal more complex features. With millions of new users, improved strategies for filtering electronic mail, searching directories, finding information, and getting online assistance will be needed. Low-cost manufacturing is a central requirement to achieve universal access for low-income Americans or the many still lower-income citizens in less technologically developed nations.

Facile tailoring of interfaces for diverse populations could be accomplished with control panels that allow users to specify their national language, units of measurement, skill level, and more. Portability to nonstandard hardware, accommodation of varying screen sizes or modem speeds, and design for handicapped or elderly users should be common practices.

Support for increased plasticity of information and services is technologically possible, but attention to this area has been limited. Convenient semantic tagging of items would enable software designers to reformat presentations, to remove unnecessary information selectively, or to integrate related materials dynamically to adjust to users' needs. Comprehensible software tools to support platform-independent authoring will enable many more people to contribute to the growing worldwide information infrastructure, as well as to their local resources.

Universal access is a policy issue: Common practices and a guiding vision are helpful to making it happen. Regulatory policies for telephones, television, and highways have been successful in creating near-universal access to these technologies, but computing economics, designs, and services apparently need revision to reach a broader audience. Fears of inappropriate intervention in free markets are legitimate, but commercial producers are likely to be major beneficiaries of universal-access policies. How might decision makers encourage industry to support universal access so as to create an expanding market that also benefits producers of commercial products and services?

In communities where adequate housing, sanitation, and food are still problems, telephones or computing are not primary needs, but the technology can still be helpful as part of an overall development plan. Community-networking technologies are being tried in well-off locales such as Taos, New Mexico, Seattle, Washington, and Blacksburg, Virginia, but adapting these designs to mountainous Nepal, urban Rio de Janeiro, or rural Botswana will take creative engineering, in addition to financial resources.

A.1.2 Universal medical records

My second hope is for improved medical record keeping. Resistance to changing the current paper-based approaches limits the availability of medical information for clinical decision making, quality control, and research. It is a paradox that airline reservations are available around the world, crossing hostile political boundaries and spanning networks of competing companies, but your medical records are inaccessible even when they might help to save your life.

A physician at any emergency room in the world should be able to review your history and see your most recent electrocardiogram or chest X-ray image within 15 seconds of your arrival, either via network or on a personal datacard. The display should appear in the local language, using familiar units of measurement, with easy access to details and convenient links for electronic consultation with physicians who know you personally.

Progress on standardizing clinical records, speeding data entry for patient histories, and designing effective overviews for viewing patient records (Color Plate B3) (Plaisant et al., 1996) could be dramatically accelerated. With a one-screen overview of patient histories, physicians could quickly spot previous surgeries or chronic diseases that might affect current decisions. Privacy protection, patient rights, and cost containment are serious concerns, but the potential for improved health care and reduced costs is compelling.

Further benefits of online medical records include assistance in formulating treatment plans and clinical research. Imagine if your physician were able to review the past year's outcomes for potential treatment plans in a sample of 10,000 patients who had your diagnosis? Imagine if scientists were able to study case histories retrospectively to support research on treatment plans and associated clinical outcomes?

With careful attention to personal privacy and costs, online medical records can become the basis for improved accountability of individual physicians and health-management organizations, as well as of improved medical understanding. Physicians may resist such visibility of their decisions, but objective comparisons with peer performance seem preferable to the current complexity and cost of malpractice litigation.

A.1.3 Universal educational support

Education is the hope of civilization. Computing is already dramatically altering education, but it is not enough to teach children about surfing the net—we must also teach them about making waves (Shneiderman, 1993). Finding information is useful only if students have a meaningful goal and a chance to influence their world.

My approach combines education with social benefit and authentic experiences to teach students how to participate in workgroups, political systems, and communities. Powerful information technologies enable students to collaborate effectively in constructing meaningful results that benefit someone outside the classroom. These action-oriented and authentic service projects done in teams produce a high level of motivation among students and give them the satisfaction of helping other people while learning.

A favorite student project involved a team that was interested in computing for the elderly. They read the literature, made a plan, brought computers to a nearby nursing home, and trained the elderly residents for several weeks. Then, their final report was written to the director of the home, with a well-reasoned plan for what might be done and pointers to helpful organizations. Another project set up a database system for a charitable organization that continues to manage volunteer and donor lists with more than 20,000 names.

Student projects could be educationally oriented, such as writing an online textbook for their course or developing an *Encyclopedia of X*, where *X* is a variable in their discipline. Creating services for other groups is compelling to students and is in harmony with the efforts in many states, such as Maryland, which requires 60 hours of community service for high-school graduation.

This *relate–create–donate* approach enlivens the educational process, pushes students to learn the relevant fundamentals, and encourages them to strive for practical goals (Denning, 1992). I'm encouraged by reports from other educators who have replicated and adapted this strategy from elementary schools (fifth graders creating a multimedia course on the animals of Africa for third graders) to graduate business schools (M.B.A. candidates set up web pages for two dozen campus and community groups).

Current technologies provide support for relate–create–donate styles of education, but four phases of creative work could be improved with advanced technology:

1. Reliable retrieval of existing knowledge relevant to team projects
2. Creative activities with brainstorming tools, simulation modeling, design exploration, and authoring tools
3. Consultation with peers and experts using convenient group-support tools
4. Dissemination of results through community-information tools

Imagine online science festivals in which student projects could build on one another over the years. New student teams could view previous projects, conduct research, and develop creative contributions, while consulting with other teams who are working on related problems or with professional scientists. The results could be reviewed by award panels, disseminated to interested people, and posted for future students.

Resistance to team projects is natural from faculty who have never had the experience themselves, but many are learning to guide computer-mediated team projects. The shift from "sage on the stage" to "guide on the side" is a challenge, and finding appropriate team projects plus management strategies takes experience. Those educators who succeed are enthusiastic about the power of collaboration and the thrill of intense experiences.

Taking responsibility for the future is a substantial challenge. It is my sincere belief that we, as computing professionals, should accept the challenge to look beyond the technology and to create a vision that inspires action. If we do so, we may be well remembered by history.

There are so many important problems to solve that there is room for everyone to contribute: jobs can be more rewarding, communities can be safer, and lives can be happier. Each one of us can make a difference.

Universal access to computing technology, universal medical records, and universal educational support are my ambitious hopes. There are surely other hopes and visions that can steer computing toward higher societal benefits, while providing unlimited challenges for researchers, entrepreneurs, and practitioners. For those people who feel inspired and wish to contribute, the time to begin is *now* and the leader to look to is *you*.

A.2 Ten Plagues of the Information Age

> The real question before us lies here: do these instruments further life and enhance its values, or not?
>
> **Mumford, *Technics and Civilization,* 1934**

It would be naive to assume that widespread use of computers brings only benefits. There are legitimate reasons to worry that increased dissemination of computers might lead to a variety of oppressions—personal, organizational, political, or social. People who fear computers have good reason for their concerns. Computer-system designers have an opportunity and a responsibility to be alert to the dangers and to make thoughtful decisions about reducing the dangers they apprehend (Huff and Finholt, 1994). Here, then, is a personal list of potential and real dangers from use of computer systems:

1. *Anxiety*

   Many people avoid the computer or use it with great anxiety; they suffer from *computer shock, terminal terror,* or *network neurosis*. Their anxieties include fear of breaking the machine, worry over losing control to the computer, trepidation about appearing foolish or incompetent ("computers make you feel so dumb"), or more general concern about facing something new. These anxieties are real, should be acknowledged rather than dismissed, and can often be overcome with positive experiences. Can we build improved user interfaces and systems that will reduce or eliminate the current high level of anxiety experienced by many users?

2. *Alienation*

   As people spend more time using computers, they may become less connected to other people (Sheridan, 1980). Computer users as a group are more introverted than are other people, and increased time with the computer may increase their isolation. One psychologist (Brod, 1984) fears that computer users come to expect rapid performance, yes–no or true–false responses, and a high degree of control not only from their machines but also from their friends, spouses, and children. The dedicated video-game player who rarely communicates with another person is an extreme case, but what happens to the emotional relationships of a person who spends two hours per day dealing with electronic mail, rather than chatting with colleagues or family members (Kraut et al., 1996)? Can we build user interfaces that encourage more constructive human social interaction?

3. *Information-poor minority*

   Although some utopian visionaries believe that computers will eliminate the distinctions between rich and poor or will right social injustices, often computers are just another way in which the disadvantaged are disadvantaged (Friedman and Nissenbaum, 1996). Those people who are without computer skills may have a new reason for not succeeding in school or not getting a job. Already, great disparity exists in the distribution of educational computers. The high-income school districts are considerably more likely to have computer facilities than are the poorer school districts. Access to information resources is also disproportionately in the hands of the wealthy and established social communities. Can we build systems that empower low-skilled workers to perform at the level of experts? Can we arrange training and education for every able member of society?

4. *Impotence of the individual*

   Large organizations can become impersonal because the cost of handling special cases is great. Individuals who are frustrated in trying to receive personal treatment and attention may vent their anger at the organization, the personnel they encounter, or the technology that limits rather than enables. People who have tried to find out the current status of their social-security accounts or tried to have banks explain accounting discrepancies are aware of the problems, especially if they have language or hearing deficits, or other physical or cognitive handicaps. Interactive computer systems can be used to increase the influence of individuals or to provide special treatment, but this application requires alert committed designers and sympathetic managers. How can we design so that individuals will feel more empowered and self-actualized?

5. *Bewildering complexity and speed*

   The tax, welfare, and insurance regulations developed by computer-based bureaucracies are so complex and fast changing that it is extremely difficult for individuals to keep up and to make informed choices. Even knowledgeable computer users are often overwhelmed by the torrent of new software packages, each with hundreds of features and options. The presence of computers and other technologies can mislead managers into believing that they can deal with the complexities that they are creating. Rapid computer systems become valued, speed dominates, and more features are seen as preferable. This situation is apparent in nuclear-reactor control rooms, where hundreds of brightly lit annunciators overwhelm operators when indicating failures. Simplicity is a simple—but too often ignored—principle. Stern adherence to basic principles of design may be the only path to a safer, saner, simpler, and slower world where human concerns predominate.

6. *Organizational fragility*

   As organizations come to depend on more complex technology, they can become fragile. When breakdowns occur, they can propagate rapidly and can halt the work of many people. With computer-based airline ticketing, telephone switching, or department-store sales, computer failures can mean immediate shutdowns of service. A more subtle example is that computer-based inventory control may eliminate or dramatically reduce stock on hand, after which disruptions spread rapidly. For example, a strike in a ball-bearing plant can force the closure of a distant automobile assembly line within a few days. Computers can cause concentration of expertise, and then a small number of people can disrupt a large organization. Can developers anticipate the dangers and produce robust designs?

7. *Invasion of privacy*

   The widely reported threat of invasion of privacy is worrisome because the concentration of information and the existence of powerful retrieval systems make it possible to violate the privacy of many people easily and rapidly. Of course, well-designed computer systems have the potential of becoming more secure than paper systems if managers are dedicated to privacy protection. Airline, telephone, bank, medical, legal, and employment records can reveal much about an individual if confidentiality is compromised. Can managers seek policies and systems that increase rather than reduce the protection of privacy in a computer-based organization?

8. *Unemployment and displacement*

   As automation spreads, productivity and overall employment may increase, but some jobs may become less valued or eliminated. Retraining can help some employees, but others will have difficulty changing lifetime patterns of work. Displacement may happen to low-paid clerks or highly paid typesetters whose work is automated, as well as to the bank vice-president whose mortgage-loan decisions are now made by an expert system. Can employers develop labor policies that ensure retraining and guarantee jobs?

9. *Lack of professional responsibility*

   Faceless organizations may respond impersonally to, and deny responsibility for, problems. The complexity of technology and organizations provides ample opportunities for employees to pass the blame on to others or to the computer: "Sorry, the computer won't let us loan you the library book without your machine-readable card." Will users of medical diagnostic or defense-related systems be able to escape responsibility for decisions? Will computer printouts become more trusted than a person's word or a professional's judgment? Complex and confusing systems enable users and designers to blame the machine, but with improved designs, responsibility and credit will be given, and will be accepted by the users and designers.

10. *Deteriorating image of people*

    With the presence of *intelligent terminals, smart machines,* and *expert systems,* it seems that the machines have indeed *taken over* human abilities. These misleading phrases not only generate anxiety about computers, but also may undermine the image that we have of people and their abilities. Some behavioral psychologists suggest that we are little more than machines; some artificial-intelligence workers believe that the automation of many human abilities is within reach. The rich diversity of human skills, the generative or creative nature of daily life, the emotional or passionate side of human endeavor, and the idiosyncratic imagination of each child seem lost or undervalued (Rosenbrock, 1982). Rather than be impressed by smart machines, accept the misguided pursuit of the Turing test, or focus on computational skills in people, I believe that we should recognize that designs that empower users will increase users' appreciation of the richness and diversity of unique human abilities.

Undoubtedly, more plagues and problems exist. Each situation is a small warning for the designer. Each design is an opportunity to apply computers in positive, constructive ways that avoid these dangers.

A.3 Prevention of the Plagues

People who are so fascinated by the computer's lifelike feats—it plays chess! it writes poetry!—that they would turn it into the voice of omniscience, betray how little understanding they have of either themselves, their mechanical-electrical agents or the potentialities of life.

Lewis Mumford, *The Myth of the Machine*, 1970

There is no sure vaccine for preventing the 10 plagues that we discussed. Even well-intentioned designers can inadvertently spread them, but alert, dedicated designers whose consciousness is raised can reduce the dangers. The strategies for preventing the plagues and reducing their effects include the following:

- *Human-centered design* Concentrate attention on the users and on the tasks that they must accomplish. Make users the center of attention and build feelings of competence, mastery, clarity, and predictability. Construct well-organized menu trees, provide meaningful structure in command languages, present specific and constructive instructions and messages, develop uncluttered displays, offer informative feedback, enable easy error handling, ensure appropriate display rates and response time, and produce comprehensible learning materials.
- *Organizational support* Beyond the software design, the organization must also support the user. Explore strategies for participatory design and elicit frequent evaluation and feedback from users. Techniques include personal interviews, focus groups, online surveys, paper questionnaires, and online consultants or suggestion boxes.
- *Job design* European labor unions have been active in setting rules for computer users to prevent the exhaustion, stress, or burnout caused by an *electronic sweatshop*. Rules might be set to limit hours of use, to guarantee rest periods, to facilitate job rotation, and to support education. Similarly, negotiated measures of productivity or error rates can help to reward exemplary workers and to guide training. Monitoring or meter-

ing of work must be done cautiously, but both managers and employees can be beneficiaries of a thoughtful plan.

- *Education* The complexity of modern life and computer systems makes education critical. Schools and colleges, as well as employers, all play a role in training. Special attention should be paid to continuing education, on-the-job training, and teacher education.
- *Feedback and rewards* User groups can be more than passive observers. They can ensure that system failures are reported, that design improvements are conveyed to managers and designers, and that manuals and online aids are reviewed. Similarly, excellence should be acknowledged by awards within organizations and through public presentations. Professional societies in computing might promote awards, similar to the awards of the American Institute of Architects, the Pulitzer Prize Committee, or the Academy of Motion Picture Producers.
- *Public consciousness raising* Informed consumers of personal computers and users of commercial systems can benefit the entire community. Professional societies, such as the ACM and the IEEE, and user groups can play a key role through public relations, consumer education, and professional standards of ethics.
- *Legislation* Much progress has been made with legislation concerning privacy, right of access to information, and computer crime, but more work remains. Cautious steps toward regulation, work rules, and standardization can be highly beneficial. Dangers of restrictive legislation do exist, but thoughtful legal protection will stimulate development and prevent abuses.
- *Advanced research* Individuals, organizations, and governments can support research to develop novel ideas, to minimize the dangers, and to spread the advantages of interactive systems. Theories of user cognitive behavior, individual differences, acquisition of skills, visual perception, and organizational change would be helpful in guiding designers and implementers.

A.4 Overcoming the Obstacle of Animism

Unlike machines, human minds can create ideas. We need ideas to guide us to progress, as well as tools to implement them. . . . Computers don't contain "brains" any more than stereos contain musical instruments. . . . Machines only manipulate numbers; people connect them to meaning.

Penzias, 1989

The emergence of computers is one of the fundamental historical changes. Such upheavals are neither all good nor all bad, but rather are an amalgam of many individual decisions about how a technology is applied. Each designer plays a role in shaping the direction. The computer revolution has passed its infancy, but there is still tremendous opportunity for change.

The metaphors, images, and names chosen for systems play a key role in the designers' and the users' perceptions. It is not surprising that many computer-system designers still mimic human or animal forms. The first attempts at flight imitated birds, and the first designs for microphones followed the shape of the human ear. Such primitive visions may be useful starting points, but success comes most rapidly to people who move beyond these fantasies and apply scientific analyses. Except for amusement, the goal is never to mimic the human form, but rather is to provide effective service to the users in accomplishing their tasks.

Lewis Mumford, in his classic book, *Technics and Civilization* (1934), characterized the problem of "dissociation of the animate and the mechanical" as the "obstacle of animism." He described Leonardo da Vinci's attempt to reproduce the motion of birds' wings, then Ader's batlike airplane (as late as 1897), and Branca's steam engine in the form of a human head and torso. Mumford wrote: "The most ineffective kind of machine is the realistic mechanical imitation of a man or another animal . . . for thousands of years animism has stood in the way of . . . development."

Choosing human or animal forms as the inspiration for some projects is understandable, but significant advances will come more quickly if we recognize the goals that serve human needs and the inherent attributes of the technology that is employed. Hand-held calculators do not follow human forms, but serve effectively for doing arithmetic. Designers of championship chess-playing programs no longer imitate human strategies. Vision-systems researchers realized the advantages of radar or sonar range finders and retreated from using humanlike stereo depth-perception cues.

Robots provide an informative case study. Beyond stone idols and voodoo dolls, we can trace modern robots back to the devices built by Pierre Jacquet-Droz, a Swiss watchmaker, from 1768 to 1774. The first child-sized mechanical robot, called the Scribe, could be programmed to write any message up to 40 characters long. It had commands to change lines, to skip a space, or to dip the quill in the inkwell. The second, called the Draughtsman, had a repertoire of four pencil sketches: a boy, a dog, Louis XV of France, and a pair of portraits. The third robot, the Musician, performed five songs on a working pipe organ and could operate for 1.5 hours on one winding. These robots made their creators famous and wealthy, since they were in great demand at the courts of the kings and in public showings. Eventually, however, printing presses became more effective than the Scribe and the Draughtsman, and tape players and phonographs were superior to the Musician.

Robots of the 1950s included electronic components and a metallic skin, but their designs were also strongly influenced by the human form. Robot arms were of the same dimension as human arms and the hands had five fingers. Designers of modern robots have finally overcome the obstacle of animism and now construct arms whose dimensions are appropriate for the steel and plastic technology and for the tasks. Two fingers are more common than five on robot hands, and the hands can often rotate more than 270 degrees. Where appropriate, fingers have been replaced by rubber suction cups with vacuum pumps to pick up parts.

In spite of these improvements, the metaphor and terminology of human form can still mislead the designers and users of robots. Programmers of one industrial robot were so disturbed by the labels "upper arm" and "lower arm" on the control panel that they scratched out the words. They thought that the anthropomorphic terms misled their intuitions about how to program the robot (McDaniel and Gong, 1982). The terms *programmable manipulators* and the broader *flexible manufacturing systems* are less exciting, but describe more accurately the newer generation of robotic systems.

The banking machine offers a simple example of the evolution from anthropomorphic imagery to a service orientation. Early systems had such names as Tillie the Teller or Harvey Wallbanker and were programmed with such phrases as "How can I help you?" These deceptive images rapidly gave way to a focus on the computer technology, with such names as the Electronic Teller, CompuCash, Cashmatic, or CompuBank. Over time, the emphasis has moved toward the service provided to the user: CashFlow, Money Exchange, 24-Hour Money Machine, All-Night Banker, and Money Mover.

The computer revolution will be judged not by the complexity or power of technology, but rather by the service to human needs. By focusing on users, researchers and designers will generate powerful yet simple systems that permit users to accomplish their tasks. These tools will enable short learning times, rapid performance, and low error rates. Putting users' needs first will lead to more appropriate choices of system features, giving users a greater sense of mastery and control, and the satisfaction of achievement. At the same time, users will feel increased responsibility for their actions and may be more motivated to learn about the tasks and the interactive system.

Sharpening the boundaries between people and computers will lead to a clearer recognition of computer powers and human reasoning (Weizenbaum, 1976; Winograd and Flores, 1986). Rapid progress will occur when designers accept that human–human communication is a poor model for human–computer interaction. People are different from computers, and human operation of computers is vastly different from human relationships. Vital factors that distinguish human behavior include the diversity of skills and background across individuals; the creativity, imagination, and inventiveness incorporated in daily actions; the emotional involvement in every act; the desire for social contact; and the power of intention.

Ignoring these primitive but enduring aspects of humanity leads to inappropriate technology and to hollow experiences. Embracing these aspects can bring about powerful tools, joy in learning, the capacity to realize goals, a sense of accomplishment, and increased social interaction.

Although designers may be attracted to the goal of making impressive and autonomous machines that perform tasks as well as humans do, realizing this goal will not provide what most users want. I believe that users want to have sense of their own accomplishment, rather than to admire a smart robot, intelligent agent, or expert system. Users want to be empowered by technology to apply their knowledge and experience to make judgments that lead to improved job performance and greater personal satisfaction. Sometimes, predefined objective criteria can be applied to a task, but often human values must be applied and flexibility in decision making is a necessity.

Some examples may help us to clarify this issue. Doctors do not want a machine that does medical diagnosis; rather, they want a machine that enables them to make a more accurate, reliable diagnosis; to obtain relevant references to scientific papers or clinical trials; to gather consultative support rapidly; and to record that support accurately. Similarly, air-traffic or manufacturing controllers do not want a machine that automatically does their job; rather they want one that increases their productivity, reduces their error rates, and enables them to handle special cases or emergencies effectively. I believe that an increase in personal responsibility will result in improved service.

A.5 In the Long Run

How do we use the power of technology
without adapting to it so completely
that we ourselves behave like machines, lost in the levers and cogs,
lonesome for the love of life,
hungry for the thrill of directly
experiencing the vivid intensity of the
ever-changing moment?

Al Gore, *Earth in the Balance,* 1992

Successful interactive systems will bring ample rewards to the designers, but widespread use of effective tools is only the means to reach higher goals. A computer system is more than a technological artifact: Interactive systems, especially when linked by computer networks, create human social systems.

As Marshall McLuhan pointed out, "the medium is the message," and therefore each interactive system is a message from the designer to the user. That message has often been a harsh one, with the underlying implication that the designer does not care about the user. Nasty error messages are obvious manifestations; complex menus, cluttered screens, and confusing dialog boxes are also sentences in the harsh message.

Most designers want to send a more kind and caring message. Designers, implementers, and researchers are learning to send warmer greetings to the users with effective and well-tested systems. The message of quality is compelling to the recipients and can instill good feelings, appreciation for the designer, and the desire to excel in one's own work. The capacity for excellent systems to instill compassion and connection was noted by Sterling (1974) at the end of his guidelines for information systems: "In the long run what may be important is the *texture* of a system. By texture we mean the *quality* the system has to evoke in users and participants a feeling that the system increases the kinship among people."

At first, it may seem remarkable that computer systems can instill a kinship among people, but every technology has the potential to engage people in cooperative efforts. Each designer can play a role—not only that of fighting for the users, but also that of nurturing, serving, and caring for them.

A.6 Practitioner's Summary

High-level goals might include world peace, excellent health care, adequate nutrition, accessible education, communication, freedom of expression, support for creative exploration, safety, and socially constructive entertainment. Computer technology can help us to attain these high-level goals if we clearly state measurable objectives, obtain participation of professionals, and design effective human–computer interfaces. Design considerations include adequate attention to individual differences among users; support of social and organizational structures; design for reliability and safety; provision of access by the elderly, handicapped, or illiterate; and appropriate user-controlled adaptation.

A.7 Researcher's Agenda

The goals of universal access, advanced applications for life services, and tools to support innovation contain enough ambitious research projects for a generation. Medical information, education, and community networks are the most appealing candidates for early research, because the impact of changes could be so large. If we are to provide novel services to diverse users, we need effective theories and rigorous empirical research to achieve ease of learning, rapid performance, low error rates, and good retention over time, while preserving high subjective satisfaction.

World Wide Web Resources

WWW

Organizations dealing with ethics, social impact, and public policy are doing their best to make computing and information services as helpful as possible. Ways for you to become an activist are also included.

http://www.aw.com/DTUI

References

Anderson, Robert H., Bikson, Tora K., Law, Sally Ann, and Mitchell, Bridger M., *Universal Access to Email: Feasibility and Societal Implications*, RAND, Santa Monica, CA (1995). Also at http://www.rand.org.

Brod, Craig, *Technostress: The Human Cost of the Computer Revolution*, Addison-Wesley, Reading, MA (1984).

Brooks, Frederick, Jr., The computer scientist as toolsmith II, *Communications of the ACM*, 39, 3 (March 1996), 61–68.

Bush, Vannevar, As we may think, *Atlantic Monthly* (July 1945). Also at http://www2.theAtlantic.com/atlantic/atlweb/flashbks/computer/tech.htm.

Denning, Peter J., Educating a new engineer, *Communications of the ACM*, 35, 12 (December 1992), 83–97.

Engelbart, Douglas C. and English, William K., A research center for augmenting human intellect, *AFIPS Proc. Fall Joint Computer Conference*, 33 (1968), 395–410.

Friedman, Batya and Nissenbaum, Helen, Bias in Computer Systems, *ACM Transactions on Information Systems*, 14, 3 (July 1966), 330–347.

Gates, Bill, *The Road Ahead*, Viking Penguin, New York (1995).

Huff, Chuck W. and Finholt, Thomas (Editors) *Social Issues in Computing: Putting Computing in Its Place*, McGraw-Hill. New York (1994).

Kraut, Robert, Scherlis, William, Mukhopadhyay, Tridas, Manning, Jane, and Kiesler, Sara, The HomeNet field trial of residential internet services, *Communications of the ACM*, 39, 12 (December 1996), 55–63.

Licklider, J. C. R., *Libraries of the Future*, MIT Press, Cambridge, MA (1965).

McDaniel, Ellen, and Gong, Gwendolyn, The language of robotics: Use and abuse of personification, *IEEE Transactions on Professional Communications*, PC-25, 4 (December 1982), 178–181.

Mumford, Lewis, *Technics and Civilization*, Harcourt Brace and World, New York (1934).

Negroponte, Nicholas, *Being Digital*, Hodder and Stoughton, London, UK (1995).

Norman, Don, *The Psychology of Everyday Things*, Basic Books, New York (1988).

Penzias, Arno, *Ideas and Information*, Simon and Schuster, New York (1989).

Plaisant, Catherine, Rose, Anne, Milash, Brett, Widoff, Seth, and Shneiderman, Ben, LifeLines: Visualizing personal histories, *Proc. of ACM CHI '96 Conference: Human Factors in Computing Systems*, ACM, New York (1996), 221–227, 518.

Rosenbrock, H. H., Robots and people, *Measurement and Control*, 15, (March 1982), 105–112.

Sheridan, Thomas B., Computer control and human alienation, *Technology Review*, 83, 1 (October 1980), 51–73.

Shneiderman, B., Durango Declaration, *Communications of the ACM*, 38, 10 (1995), 13.

Shneiderman, B., Engagement and construction: Educational strategies for the post-TV era, *Journal of Computers in Higher Education*, 2, 4 (Spring 1993), 106–116.

Shneiderman, B., Human values and the future of technology: A declaration of empowerment, Keynote address, ACM SIGCAS Conference on Computers and the Quality of Life CQL '90, *SIGCAS Computers and Society*, 20, 3 (October 1990), 1–6.

Sterling, T. D., Guidelines for humanizing computerized information systems: A report from Stanley House, *Communications of the ACM*, 17, 11 (November 1974), 609–613.

Weizenbaum, Joseph, *Computer Power and Human Reason*, W. H. Freeman, San Francisco, CA (1976).

Winograd, Terry and Flores, Fernando, *Understanding Computers and Cognition: A New Foundation for Design*, Ablex, Norwood, NJ (1986).

Name Index

C

D

I

J

K

T

Y

Z

Subject Index

F

P

T

U

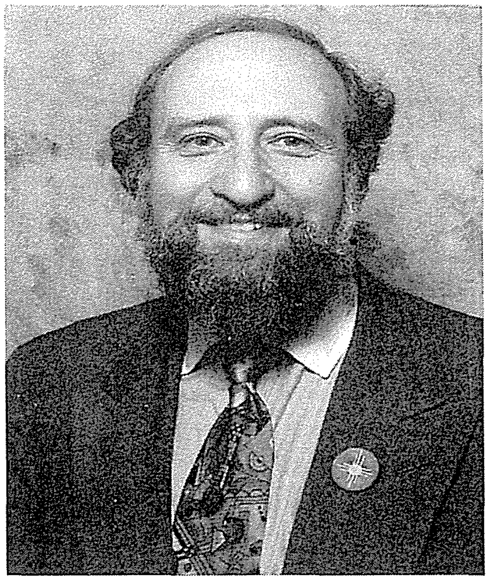

BEN SHNEIDERMAN

Ben Shneiderman is Professor in the Department of Computer Science, Head of the Human–Computer Interaction Laboratory, and Member of the Institutes for Advanced Computer Studies and for Systems Research, all at the University of Maryland at College Park. He taught previously at the State University of New York and at Indiana University. He regularly gives conference keynote speeches, tutorials, and short courses on user-interface design, information visualization, and educational technology. He also organizes an annual satellite–television presentation—*User Interface Strategies*—that has been seen by thousands of professionals since 1987. He has consulted and lectured for many organizations, including Apple, AT&T, Citicorp, GE, Honeywell, IBM, Intel, Library of Congress, NASA, and university research groups.

Recognized worldwide as a leader in human-computer interaction, Ben authored the first edition of *Designing the User Interface* in 1987, the second edition in 1992, and now this throroughly updated third edition. His early works include another influential book, *Software Psychology: Human Factors in Computer and Information Systems* (1980), and a seminal paper (1981) in which he coined the term *direct manipulation* to describe graphical user interface design principles. Later, with Greg Kearsley, Ben coauthored the first commercial electronic book, *Hypertext Hands-On!* (Addison-Wesley, 1989). This book, comprising both a print version and a full hypertext version on diskettes, pioneered the highlighted embedded link. He developed this concept in the Hyperties hypermedia system.

A prolific writer and editor, Ben has coauthored two other textbooks, edited three technical books, and published nearly 200 technical papers and book chapters. His edited book, *Sparks of Innovation in Human-Computer Interaction* (Ablex, 1993), collects 25 papers from 10 years of research at the University of Maryland. He has been on the Editorial Advisory Boards of nine journals, including two important ACM publications, *Transactions on Computer–Human Interaction* and *Interactions*.

Ben received his B.S. from City College of New York in 1968, and his Ph.D. from State University of New York at Stony Brook in 1973. He received an Honorary Doctorate of Science from the University of Guelph, Ontario, Canada, in 1996, and was elected as a Fellow of the Association for Computing (ACM) in 1997.

You can learn more about Ben, the work at his lab, and resources for this book from the following World Wide Web sites:

http://www.cs.umd.edu/~ben
http://www.cs.umd.edu/projects/hcil
http://www.aw.com/DTUI

Software Engineering / User Interface

Recognizing this book, ACM's Special Interest Group on Documentation (SIGDOC) presented Ben Shneiderman with the Joseph Rigo Award (1996). SIGDOC praised the book as one "that took the jargon and mystery out of the field of human–computer interaction," and attributed the book's success to "its readability and emphasis on practice as well as on research."

Designing the User Interface

Strategies for Effective Human–Computer Interaction
Third Edition

Ben Shneiderman, *University of Maryland*

In revising this popular book, Ben Shneiderman again provides a complete, current, and authoritative introduction to user–interface design. The user interface is the part of every computer system that determines how people control and operate that system. When the interface is well designed, it is comprehensible, predictable, and controllable; users feel competent, satisfied, and responsible for their actions. Shneiderman discusses the principles and practices needed to design such effective interaction.

Based on 20 years experience, Shneiderman offers readers practical techniques and guidelines for interface design. He also takes great care to discuss underlying issues and to support conclusions with empirical results. Interface designers, software engineers, and product managers will find this book an invaluable resource for creating systems that facilitate rapid learning and performance, yield low error rates, and generate high user satisfaction.

Coverage includes the human factors of interactive software (with a new discussion of diverse user communities), tested methods to develop and assess interfaces, interaction styles such as direct manipulation for graphical user interfaces, and design considerations such as effective messages, consistent screen design, and appropriate color.

Highlights of the Third Edition:

- New chapters on the World Wide Web, information visualization, and computer-supported cooperative work
- Expanded and earlier coverage of development methodologies, evaluation techniques, and user–interface–building tools
- Thought-provoking discussion of speech input-output, natural-language interaction, anthropomorphic design, virtual environments, and agents
- An associated Web site that provides additional illustrations and examples, updates, and links to useful resources: **http://www.aw.com/DTUI**

Ben Shneiderman is a professor of computer science at the University of Maryland, where he heads the Human–Computer Interaction Laboratory. A pioneer in user–interface design, he is known throughout the world for his research and writing. You can learn more about him, and about this book, from the DTUI Web site.

ADDISON-WESLEY

Addison-Wesley is an imprint
of Addison Wesley Longman, Inc.

0000

ISBN 0-201-69497-2