INTERNATIONAL JOURNAL OF HUMAN–COMPUTER INTERACTION, *14*(1), 25–92

Content Preparation and Management for Web Design: Eliciting, Structuring, Searching, and Displaying Information

Session Organizers
Robert W. Proctor
Purdue University

Kim-Phuong L. Vu
Purdue University

Gavriel Salvendy
Purdue University and Tsinghua University

Participants (alphabetically)
Helmut Degen
Siemens AG

Xiaowen Fang
DePaul University

John M. Flach
Wright State University

Sherrie P. Gott
University of Texas Health Science Center at San Antonio

Douglas Herrmann
Indiana State University

Heidi Krömker
Siemens AG

Nancy J. Lightner
University of South Carolina

Kem Lubin
Siemens Corporate Research

Lawrence Najjar
Viant

Leah Reeves
University of Central Florida

Arnold Rudorfer
Siemens Corporate Research

Kay Stanney
University of Central Florida

Constantine Stephanidis
Institute of Computer Science, Foundation for Research and Technology-Hellas; University of Crete

Thomas Z. Strybel
California State University Long Beach

Misha Vaughan
Oracle Corporation

Huifang Wang
Cisco Systems

Harald Weber
Institute of Computer Science, Foundation for Research and Technology-Hellas

Yanxia Yang
Cadence Design Systems, Inc.

Wenli Zhu
Microsoft Corporation

The vast amount of information available through the Web has made it difficult to retrieve information relevant to a specific task. To help ensure that users' interactions with a system are successful, preparation of content and its presentation to users must take into account (a) what information needs to be extracted, (b) the way in which this information should be stored and organized, (c) the methods for retrieving the information, and (d) how the information should be displayed. The goal of this article is to discuss the generic problems facing content preparation and evaluate the current methods available to help remedy them, as well as identify areas in which more research is needed. The material presented in this article was a result of the collective efforts of the participants of a special "white paper" session that was part of the 9th International Conference on Human–Computer Interaction (HCI International 2001).

1. INTRODUCTION

The information age expanded significantly with the development of the modern Internet. The Internet makes it possible for a wide range of users to access any available information at any time and from anywhere. The information on the Web ranges from personal homepages to consumer information, catalogs of products, and large databases. With personal homepages, the amount of information contained is usually small, and the users know the information that they would like to make available to others. As a result, there is not much cost for personal homepages to be designed according to individual preferences. However, for sites that contain a large amount of data intended for access by a range of users, or even a targeted group of users, the way in which information is organized and displayed is vital. For some cases, such as a library catalog, the organization and structure of information, as well as the type of information that should be displayed, is straightforward: The information in the library catalog should be available by author, topic, call number, title, and so on. Although the usability of accessing the online catalog can be greatly affected by the way in which the Web pages are designed, there is little issue about what information needs to be conveyed.

Unlike the library catalog, many systems and Web sites contain material that cannot be readily classified into well-established, distinct, and limited topics for search. Furthermore, for these systems and Web sites, the way in which the content is prepared, organized, and retrieved is often of extreme importance. Consequently, the content designer may play an integral role in the success or failure of such systems. To

The current article is based, in its entirety, on information derived from participants invited to a special session at the 9th International Conference of Human–Computer Interaction International (HCI International 2001). The session was organized by Robert Proctor, Kim Vu, and Gavriel Salvendy. The organizers wrote the initial draft of the paper based on discussions from the white paper session. The draft was sent to all participants for their input regarding its contents, and comments and material that they provided were integrated into the manuscript. Subsequent drafts were also circulated for additional comments. Thus, the article integrates and reflects the various viewpoints of the participants.

Portions of this article incorporate information presented in the proceedings papers authored by the participants. The participants are listed in alphabetical order, and a brief description of each participant is included in the Appendix.

The "white paper session" was funded by Siemens. We thank Heidi Krömker for her support.

Requests for reprints should be sent to Robert W. Proctor, Department of Psychological Sciences, Purdue University, West Lafayette, IN 47907–1364. E-mail: proctor@psych.purdue.edu

help ensure that users' interactions with these systems are successful, the primary goals of the content designer and of content preparation in general should be to determine (a) what information needs to be extracted, (b) the way in which this information should be stored and organized, (c) the best methods for retrieving the information, and (d) how the information should be displayed (see Figure 1).

Today, the traditional issues associated with the efficiency of information retrieval have expanded enormously. With the advent of the Internet and Web-based storage systems, the information stored can be linked to different locations within the system (e.g., getting to the homepage from any Web page within the site), accessed by multiple users, or converted into a different platform (e.g., retrieving e-mail stored on a server using a personal computer or personal data assistant). Although the flexibility of the Internet allows information to be available and accessible in more than one way, there are drawbacks. For example, the amount of information available through the Web has made it more difficult to retrieve information relevant to a specific task, and the format in which the information is displayed in one platform may not be optimal for another platform.

1.1. How is Content Currently Organized and Presented?

Currently, there is no standardized way to organize and present information on the Web. Therefore, there is often a noticeable range in the quality of information orga-

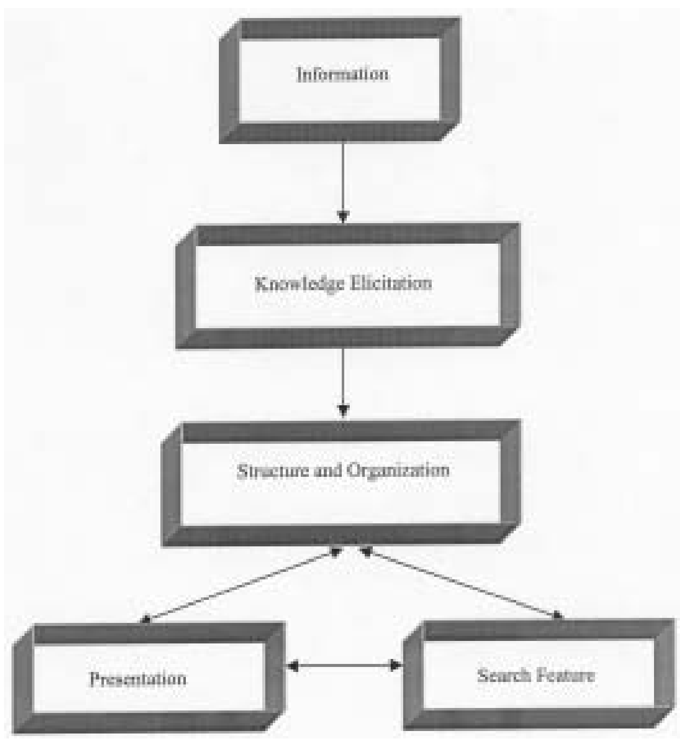

FIGURE 1 Schema for content preparation and management.

nization and presentation. Cymfony (2001) indicated that 70% of business data (e-mails, faxes, memos, Web sites, etc.) consists of unstructured text and that this percentage is increasing each year. Due to the inefficient structure of the information, users may waste a lot of time and energy searching for specific items. For instance, Cymfony also estimated that 53% of initial data searches do not yield the desired information. Furthermore, it has been shown that users looking for a desired item on e-commerce Web sites cannot find it 36% of the time (Nielsen & Tahir, 2001). Given the high rate of unsuccessful searches, improving the hit rate may increase sales significantly. The bottom line is simple—e-commerce sites cannot sell products if users cannot find the products within their Web sites. The implication is that companies that are able to organize and structure information in a way that promotes efficient and effective retrieval will save time and money, promote customer satisfaction and continued business, and get an advantage over their competitors.

One way to improve a system design and guarantee more successful retrieval of information is through usability engineering. Usability engineering has evolved to ensure that human–interactive systems are intuitive, effective, and subjectively acceptable to users (Nielsen, 1993). Although the ideal approach to usability engineering is to consider these objectives from the inception of the system development process, usability may be evaluated in the later development stages or may focus on existing systems. Usability testing is the evaluation of the ease with which humans can use a system (Nielsen, 1997b). It is probably the most well-known method for assessing the adequacy of systems in terms of their compatibility with the characteristics of the potential users and effectiveness in helping users achieve their goals.

Through usability testing, substantial knowledge has been gained about aspects of human–computer interaction (HCI) such as how to design interfaces to accommodate the capabilities of humans and to optimize human performance. The benefits of usability testing include but are not limited to the following: reduced system redesign costs, increased system productivity, enhanced user satisfaction, decreased user training, and decreased technical support (Nielsen, 1993). Improving usability can also improve sales. When the usability of their sites was improved, IBM found a 400% increase in sales (Tedeschi, 1999), and Digital Equipment Corporation reported an 80% increase in revenue (Wixon & Jones, 1992). One study (Creative Good, 1999–2001) estimated that improving the customer experience increases the number of buyers by 40% and increases order size by 10%. The initial and continuous preparation and design of Web page information is critical to the success of Web sites.

1.2. General Trends in Content Preparation and Management

Numerous companies have begun to note this need for content preparation, and a new emphasis on providing content preparation services is emerging. Companies providing services in this and related areas include the following: Interwoven, Documentum, Verity, and Vignette. Brief descriptions about their content preparation and management services, taken from the individual Web sites of these companies, are provided to illustrate the development of this service:

• Interwoven: "Interwoven is the leader in Content Infrastructure systems and services for the enterprise Web. Interwoven's flagship product, TeamSite, manages the development and deployment of business-critical Web sites. TeamSite currently powers such renowned eBusiness sites as Cisco Systems, General Electric, Xerox, and the U.S. Department of Education … . Interwoven provides the infrastructure that supports electronic commerce, customer relationship management, supply-chain management and knowledge management" (http://www.interwoven.com/company/).

• Documentum: "Documentum 4i™ eBusiness Platform is the industry standard for managing and distributing large volumes of content within and beyond the enterprise. Based on long-standing expertise for managing electronic content, Documentum 4i™ provides an open, scalable, and completely reliable platform for building and deploying e-business solutions, enabling collaborative portals, meeting regulatory requirements, and powering global Web sites" (http://www.documentum.com/products/content-management_products.html).

• Verity: "Through the Verity Webtop—the convenient, single point of access to information enterprise-wide—users and administrators can personalize the sources, layout and delivery of content … . Portals powered by Verity feature the most precise and accurate search, navigate and view capabilities available—the backbone of any effective portal solution … . Enhancing this advanced search is intuitive navigation that lets users browse through familiar directories or limit searches to specific categories" (http://www.verity.com/pdf/MK0360_CorporateBro.pdf).

• Vignette: "Vignette's software enables organizations to

- Manage and deliver relevant, timely content to communicate effectively using any electronic customer touch-point.
- Integrate e-business applications to collaborate within and among enterprises to meet real-time customer demand.
- Analyze the customer's online experience to comprehend and adapt to their changing needs."

(http://www.vignette.com/CDA/Site/0,2097,1–1-1329–2056,00.html?vgn_rms_id=8_10).

As the aforementioned descriptions convey, content preparation and management will be of extreme importance to any organization, especially for those relying on Web access.

1.3. Purpose

Although there has been an increased interest in the area of content preparation and management, along with initial efforts to develop tools to combat problems associated with poor organization and retrieval of information, the currently available mechanisms are not fully developed or adequate. Current issues include, among others, the following: Web sites are not designed in a consistent manner (e.g., navigation), many Web sites do not utilize established guidelines to improve usability, search engines yield results that are irrelevant, information posted on the Web may be inaccurate or unavailable, and the download time for many Web pages is longer than what the user will tolerate.

The purpose of this "white paper" is to discuss the generic problems facing content preparation and to evaluate the current methods available for minimizing or eliminating these problems. Although the area of content preparation and management is broad and applicable to any organization, we place most emphasis on content preparation for e-business and Web design to narrow the focus. E-business is defined as business-to-consumer and business-to-business commerce using electronic media for exchange. E-business also includes information exchange within organizations (called intranets) and between organizations (extranets). Four areas are considered relating to content preparation and management: knowledge elicitation, organization and structure of information, retrieval of information (search engines), and presentation of information. Within these areas, an attempt to answer the following three questions is made:

1. What are current issues pertaining to content preparation and management?
2. What are the methods that are currently being used?
3. On what topics or areas should future research focus?

2. KNOWLEDGE ELICITATION

One reason why finding relevant information on the Web or in large databases is difficult is that the systems are not programmed to "know" what information should be extracted for the various possible uses and how to extract this information. One solution to this problem is to consult experts in the field to determine what information is needed and how it should be structured and displayed. Knowledge elicitation involves identifying the knowledge structures and processes that competent individuals bring to task performance in a domain (Lehto, Boose, Sharit, & Salvendy, 1992).

Initially, knowledge elicitation techniques were developed in the context of expert systems, which are software systems designed to solve problems in particular domains much like experts (e.g., Sell, 1985). Consequently, the major focus of knowledge elicitation techniques has been on extracting knowledge from experts in a given domain, with the expectation that the expert's knowledge is maximally organized and tied to its conditions of use. According to Shadbolt and Burton (1995), two questions dominate in knowledge elicitation from experts: How do we get experts to tell us, or else show us, what they know that enables them to be experts at what they do?, and How do we determine what constitutes the expert's problem-solving competence? Researchers working in the field of knowledge elicitation can attest to the difficulties involved in getting answers to these questions.

However, advances in the cognitive sciences have produced a variety of techniques that can help reveal the expert's underlying cognitive processes and knowledge structures—the components of performance that are most useful to know in developing intelligent, interactive training systems and human–computer interfaces (Gott & Morgan, 2000; Shraggen, Chipman, & Shalin, 2000). These techniques go beyond simply asking experts to verbalize the sequence of steps they follow to complete a task. Advanced techniques elicit, for example, why the step is executed at a given point in time (as opposed to all the other steps that could be taken), as well as the internalized conceptual models and strategic knowledge that experts use to solve a problem by sequencing a series of steps in a particular way.

When the goal is to convey information to a wide range of users who vary in their expertise, knowledge is needed not only from experts but also from the full range of target users. Therefore, emphasis has recently been placed on eliciting knowledge from target users, who are mostly nonexperts, to determine what information should be conveyed to them and how it should be conveyed. In this case, the knowledge elicitation task changes from trying to understand how experts solve problems to trying to understand the knowledge of targeted users for the purpose of conveying the information they desire. Developers need to consider the knowledge structures that users bring for interacting with the system because this knowledge will likely predict their actions.

2.1. What Are Current Issues Regarding Knowledge Elicitation Techniques?

This section reviews the most salient issues regarding techniques for eliciting knowledge from experts, novices, targeted user groups, and metadata. Furthermore, problems with eliciting, documenting, and organizing knowledge are discussed. The importance of ensuring that elicited knowledge is valid is also addressed.

Eliciting Knowledge From Experts

In premodern times, an apprentice elicited knowledge from an expert by observing the expert perform tasks for many years. The more overt and observable the elements of the performance, the more the task performance lent itself to traditional forms of apprenticeship training and assessment. For example, a carpenter's apprentice could generally learn a great deal by following a conventional apprentice practice of observing the master, executing a task with support and critique from the master, and then continuing to practice with the master's support, eventually working independently. Similarly, assessment could be accomplished by evaluating observable behaviors and products of the behaviors using standards provided by experts at the craft.

Learning through apprenticeship (as a premodern form of eliciting an expert's knowledge) has historically worked well, but it takes years to train one apprentice, and this method seems to apply mainly to skills that are overtly physical, not primarily cognitive. Therefore, apprenticeship as a type of knowledge elicitation method (used by the apprentice) is not easily applied to the modern workplace, where the desired skill to be elicited is cognitive and less physical. Internalized cognitive structures and processes have replaced the external behavioral elements of work tasks from an earlier time. Because so much of "the action" is now internalized and unobservable, learning through observation has been significantly hampered. No longer is it effective to focus on overt behaviors and observable end products. Rather, practical learning experiences and assessments for modern work environments must be targeted at the internalized cognitive concepts and processes that lie behind expertise (see Gott, 2001).

Because modern expertise is dominated by internalized knowledge, a major problem for knowledge elicitation is to get experts to explain what they know and how they use their knowledge to solve domain problems. One reason why eliciting knowledge from experts is such a difficult task is that much of the expert's performance has become automatic. Once automated, expert knowledge often becomes tacit, or implicit, making it difficult to articulate. In other words, experts may know what to do next but cannot explain why or how they came to that conclusion. One way to overcome this limitation of expert knowledge is to develop authentic problem-solving scenarios that experts pose to each other (where the solution is not known in advance by the expert solver, only the poser). In the process, the expert posing the problem questions the solving expert about articulated solution steps. Questions would include, What other steps could be taken at this juncture?, and Which of the possible steps is best, and why?

It has been shown that knowledge elicitation should be conducted using multiple experts (usually two to four is sufficient), rather than a single expert, to avoid individual idiosyncrasies that would not be applicable to most experts in the domain (Boose, 1986; Dauer, 1990; DiPiazza, 1990). Furthermore, although a group of individuals may be experts in a domain, they generally do not execute tasks in exactly the same way. Therefore, there is also the problem of integrating various expert performances to represent the broadest expert solution space. When multiple experts are used, it is often constructive to examine the input from a dissenting expert (Boose & Bradshaw, 1988).

Eliciting Knowledge From Nonexperts

Knowledge elicitation from nonexperts has been used typically to identify normal learning progression in a given domain and to find the common barriers to expertise (Gott, 2001). By comparing the performance of experts and nonexperts on the same task, important expert–novice skill differences can be detected. In a given domain, the thinking that experts apply to tasks differs in fundamental ways from the thinking of performers at less-advanced stages of development. Experts typically deploy diverse types of knowledge and skill, which interact in the problem-solving (or critical thinking) process.

Research into the nature of expertise has consistently revealed differences between expert and novice knowledge structures. For example, mental models of novices are typically organized around the surface-level features of objects or scenarios in a field of practice, whereas experts have "deep structure" models (Chi, Glaser, & Farr, 1988). Due to their limited knowledge structures or mental models, the reasoning of novices is restricted to the superficial dimensions of a problem, whereas experts can reason about a problem at deeper, functional levels of understanding. Further, on diagnostic tasks, where hypothesis testing is demanded, studies have shown that experts distinguish themselves from novices by engaging in forward reasoning, for example, reasoning from what is known about a medical condition to some unknown diagnosis (Groen & Patel, 1988). Novices, by comparison, often engage in backward reasoning, where they fix on a diagnosis early on and then conduct tests (or test their hypotheses) to confirm their earlier diagnoses. When they encounter any disconfirming evidence, they either ignore it or misinterpret it as support for their initial diagnosis.

Eliciting Knowledge From Targeted End Users

Although most research on knowledge elicitation with nonexperts has been directed at finding differences in the way they solve problems as compared to experts, knowledge elicitation with targeted end users can also yield valuable information. For instance, the owners of most e-commerce Web sites want to present information in ways that will attract multiple user types. Eliciting information on how users interact with these Web sites can give designers insight into presenting information that is consistent with the users' expectations. Consequently, when designing based on the user experience, the process commences with exploratory rather than definitional activities, where designers are focused on discovery and invention (Bear, Teasley, & Carroll, 2001). Through contextual inquiries, competitive analyses, and prior-art research, designers focus on the actual needs and intents of target users. They identify users' goals, desires, and behavioral patterns, usage contexts, and competitors, as well as technology barriers (Crow & Yang, 2001). Designers can make decisions based on observations of the real behaviors of target users (Armitage, 2001).

Stanney, Maxey, and Salvendy (1997) identified key factors that should be elicited during exploratory activities. These factors include the following:

- Target users' capabilities and limitations (both cognitive and physical).
- User requirements (e.g., situated roles and responsibilities).
- Organizational factors (e.g., interaction with customer base, social issues).
- Task requirements (e.g., cooperative task activities; critical inputs and outputs to a service system).
- Equipment and system specification (e.g., artifacts that support task activities).
- Environmental factors (e.g., informal practices) that the service environment or product supports.

Through the familiarity gained by this involvement, designers can format and lay out systems so the content presentation is appropriate for target user groups. Unfortunately, eliciting this information can be a time-consuming and costly process. For example, MONKEYmedia reports spending one third of its project development time and budget on design analysis (Bear et al., 2001).

Eliciting Knowledge From Metadata

Knowledge can also be elicited from existing databases to identify and analyze different types of trends (Welbank, 1990). Metadata are basically data about existing data. For example, a library catalog consists of information (metadata) about books, journals, and authors (data). Metadata can be used in a wide range of applications such as e-commerce, intelligent software agents, content rating, and cataloging. In e-commerce, metadata can be used to encode information for locating products, companies and consumers, prices for products, different payment options, delivery options, and so forth. (Williams & Kotnour, 1993).

Defining Expertise

A number of problems exist in eliciting knowledge from experts. Finding and identifying a world-class expert in a domain is difficult and often expensive. Furthermore, determining who is an appropriate expert is not a simple matter. For example, Flach (2000) noted the following on the basis of studies of laparoscopic surgery:

> Who is the expert? Who is the authority with respect to the amount of anesthesia an older patient might tolerate—the surgeon or anesthesiologist? Who is in a better position to predict the consequences of an extended hospital stay—the surgeon, nurse, or general practitioner? In aviation, who has the better knowledge of the performance envelope of a particular aircraft—the pilot or aeronautical engineer? Most complex domains involve multiple experts each with different perspectives on the domain. Who should be the target for the knowledge elicitation? (pp. 93–94)

The previous examples describe how a person may be an expert in one subset of a domain but not in another subset. An individual may be an expert at a domain level, a task level, a strategic level, or an inference level. Moreover, different experts may have different knowledge structures for the same domain. For example, Gillan, Breedin, and Cooke (1992) elicited knowledge from two groups of experts in HCI—software development experts and human factors experts. The knowledge networks generated from card-sorting tasks for the two expert groups were quite different (see Figure 2). The human factors experts organized their concepts exclusively by usability, and the software design experts organizing theirs by dimensions related to technology and implementation, as well as user characteristics.

Costs Associated With Knowledge Elicitation Techniques

The current problems with knowledge elicitation techniques outlined earlier demonstrate the difficulty in obtaining knowledge from targeted groups (e.g., experts or end users; Boose & Gaines, 1988; Gaines & Boose, 1988; Williams & Kotnour, 1993). Knowledge elicitation methods are usually time consuming, making it necessary for companies to determine a cost–benefit ratio. However, companies do not always know when they should stop the process (or keep it going) to accomplish their goals. Furthermore, because knowledge elicitation techniques must be performed accurately to yield valid results, companies that do not have qualified knowledge elicitation engineers or do not know what methods should be used for a given scenario, may have to rely on consultants at an additional cost. Because many companies cannot afford the high costs associated with this process,

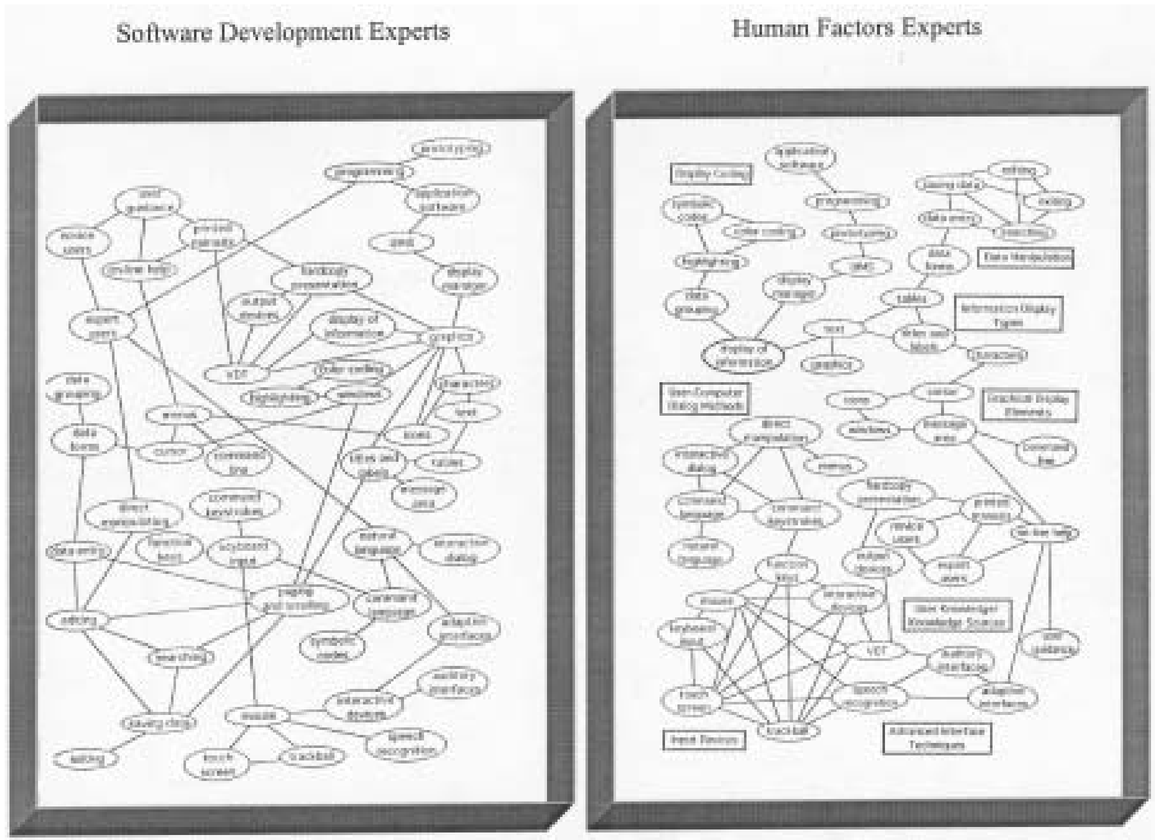

FIGURE 2 Pathfinder networks of human–computer interaction concepts for software development experts and human factors experts. Based on Gillan, Breedin, and Cooke, *International Journal of Man–Machine Studies.* Copyright (1992) by Harcourt, Inc. Adapted with permission.

this process may be eliminated entirely or companies just rely on knowledge elicitation software tools.

2.2. What Are the Existing Knowledge Elicitation Methods Being Used?

Methods for Eliciting Knowledge From Experts

Interviews. The most well-known knowledge elicitation method is an interview (Shadbolt & Burton, 1995). There are basically two types of interviews: structured and unstructured. With the structured interview, the interviewer plans and directs the session, and the expert generates scenarios or explanations relating to the questions asked by the interviewer. The benefit of structured interviews is that the content is easier to analyze because it was given in a specific context. In a typical structured interview, the interviewer usually asks the expert to describe a task that is performed by the expert and the methods and strategies that are used to accomplish the task. In some cases, the interviewer uses specific sets of probes to elicit more details or to provide the interviewer with a better understanding of the description.

Unstructured interviews have little predetermined direction that is specified by the interviewer. Rather, the interviewer provides the expert with the general goal of the session and lets the expert express his or her knowledge in any manner that the expert desires. Although the content derived from this type of interview may be harder to analyze, the unstructured interview does not exert stringent constraints on the interview, allowing the expert to express knowledge in a manner that is comfortable and consistent with the expert's mental model. Structured and unstructured interviews are the two extremes of interview types, and the interviewer can choose to conduct a session that is intermediate to these extremes.

Verbal protocol analysis. Protocol analysis refers to the ways in which performance of experts solving problems is analyzed by having them report their thought processes in solving the problem (Ericsson & Simon, 1984). There are two different types of protocol analyses. The first is online protocol analysis in which the expert "thinks aloud" while performing the task (describes what he or she is doing and why), or another expert provides a description of what the observed expert is doing. Offline protocol analysis allows the expert (or other experts) to view the recording of an expert solving the task. This method permits the expert actually performing the task to reflect on the actions the expert performed previously for the purpose of providing a report on what the expert was thinking while performing the task. If other experts view the video, there can be a group discussion regarding the task.

Protocol analysis requires that the knowledge elicitation engineer be familiar with the domain for two primary reasons: (a) for the engineer to understand which behaviors are important and should be emphasized, and (b) to select the tasks or problems for the analysis. Although protocol analyses can yield critical information, the verbal descriptions obtained from the expert performing the task or other experts observing the task can be difficult to analyze. Furthermore, tasks may be performed in superficially similar manners, yielding similar protocols, but be the result of different processes that are not captured by the protocol analysis (J. Karat, 1997; Nielsen, 1993; Stanney, Smith, Carayon, & Salvendy, 2001).

Group task analysis. In a group task analysis, several experts are brought together to describe what they do and to create a flow chart that represents the steps they take to perform a specific task (Dayton, McFarland, & Kramer, 1998; Lafrehiere, Dayton, & Muller, 1999). A group task analysis is typically facilitated by a nonexpert. The nonexpert's role is to record the steps that the experts take when working on the task. The nonexpert clarifies each step and identifies steps that the experts unconsciously perform automatically. In addition, the nonexpert can obtain preference information along the way. The benefits of group task analysis include being able to capture dynamic processes involved in performing a task and being able to document thought processes for the task. The disadvantages are that this approach needs to be conducted with another knowledge elicitation method, there is no research validating this method, and it is developed through trial and error.

Narratives and scenarios. Narratives, or scenarios, refer to the use of stories to elicit information from people. They are "stories about people and their activities" (Carroll, 1999, p. 123). Scenarios have characteristic elements such as a setting (situation state), agents or actors, goals or objectives, sequence of actions and events, and outcome. Carroll argued that scenario-based design can "promote use-oriented reflection-in-action, suspend commitment but support progress, enhance work orientation, support many needs and purposes and cumulate/integrate design work and results" (p. 123).

Three reasons why narrative structure may provide unique insights are as follows. First, the narrative structure may be particularly well suited to abductive forms of reasoning. Peirce (1878a, 1878b) introduced the term *abduction* to refer to the process or reasoning involved when one looks for a pattern in a phenomenon to suggest a hypothesis. Abduction can be thought of as the construction of a plausible "story" to account for a set of observations, particularly when these observations are surprising. If people reason in abductive terms, then elicitation techniques that invite operators to relate in narrative modes (i.e., to tell stories) may provide insights into the reasoning process. Second, the narrative structure can capture elements of context that can help to "situate" the cognitive processes. The person providing the narrative describes the reasoning and acts in which they engaged in a specific situation. Third, in the process of telling the stories, experts often discover dimensions of problem solving that had been implicit. For example, implicit assumptions may need to be addressed to produce a coherent story (e.g., Crandall & Calderwood, 1989).

One company, Info.Design, uses narratives as their primary way to collect information. Their Web page states the following:

> We use narrative as a route to improve understanding. As human beings, we look to stories as a way to organize, interpret, and create meaning from our experiences. Stories enable us to create a deep understanding of the complex relationships that exist within our business communities.

Info.Design is committed to improving understanding. We see narrative as a route to understanding. We've researched reasons why narrative has traditionally been dismissed by contemporary business, and we've explored ways forward-thinking businesses are re-introducing narrative in the workplace" (http://www.infodn.com/tell.shtml).

Critical incident reports. Critical incident analysis is a type of narrative method in which operators tell stories about specific events that tested their expertise (Klein, 1999). The hope is that these stories will provide insights into the processes involved in generating and evaluating plausible hypotheses about the nature of an unusual phenomenon or data set, as in an emergency situation. This type of reasoning is particularly important when a user is solving an unexpected problem. Use of critical incident reports can help designers to better understand how experts reason within a domain and how they search the problem space to find the appropriate responses to unexpected events.

Methods for Eliciting Knowledge From User Groups

Questionnaires. When designers want to elicit information from a large group of users, questionnaires or surveys are often used. The benefits of questionnaires include being able to obtain information from different user populations and getting information that is relatively easy to code. However, the return rate on questionnaires is typically low. In addition, the types of questions asked affect the validity of the questionnaire. Vu, Hanley, Strybel, and Proctor (2000) found that users' judgments are relatively reliable when the users are familiar with the task. However, Nielsen (2001a) indicated that caution should be taken when users are asked to predict what they would do in a certain situation because their actions and verbal reports may not correspond. Furthermore, Bailey (1993) showed that users may indicate that they prefer a design that would lead to poorer performance.

Focus groups. Focus groups usually consist of 6 to 10 users who are brought together in a session to discuss different issues regarding the features of a system. The group is directed by a moderator, who is in charge of keeping the group on track and getting all users involved. The benefits of focus groups include getting information for different aspects of the system and allowing users to react to ideas presented by other users. Focus groups have the same disadvantages as questionnaires and other knowledge elicitation methods: What users say may not truly reflect what they do. Users may not be able to articulate all the steps or knowledge due to automation and the tendency to summarize. Also, a single talkative individual can dominate or influence the discussion. This method works well for high-level goals, such as generating a list of functions or features for a product. It does not work well for discovering specific usability problems in a product, such as detailed navigational issues based on a specific user task.

Wants and needs analysis. In a wants and needs analysis, a facilitator asks users to brainstorm about the content, or information, they want and need to have included in an ideal system (Byrne & Barlow, 1993; Delbecq & Van den Ven, 1971;

Graetz, Barlow, Proulx, & Pape, 1997). Once users have determined the type of information they would like the system to convey, the facilitator then asks the users to indicate why each piece of information should be included. The results are then incorporated into the design process.

Observation and contextual inquiry. Observation and contextual inquiries are methods for evaluating how users interact with a product by observing the users in their natural surroundings (Beyer & Holtzblatt, 1998). That is, understanding the context in which a product is used is important. The evaluator can observe users' behavior directly and make notes on points of interest or use a video- or audiotape to record users' interactions and perform analyses later. The results of these techniques can be both qualitative and quantitative in nature (Hackos & Redish, 1998).

Ethnographic studies. Ethnographic studies seek to understand the users' culture and their work environments, and to identify the artifacts used to accomplish their goals (Takahashi, 1998). Ethnographic studies derive from anthropology, where ancient or remote cultures are studied. "Ethnography is a form of anthropological practice. It is both a methodology and a perspective … . At its core, ethnography relies on 'participant observation'" (Bell, 2001, p. 1). The researchers typically start the research with the purpose of understanding the potential users and their daily lives by "hanging out" with them (e.g., the *shadowing technique*). The questions regarding specific products or evaluations are formed based on the initial understanding of the users' daily lives, rather than derived ahead of time. In addition, ethnographic studies typically take longer to conduct, with the goal of the researcher being to understand the users so well that the researcher in a sense becomes a target user. The researcher can then act accordingly even in the event of new issues not previously encountered in the study period. The results of ethnographic research are typically told via stories, pictures, audios, and videos that capture the target users and their environments. The ultimate goal is to bring the target users to life so that the product teams have a sense of who and what the potential users are, develop empathy for them, and ultimately develop the products to fit their needs (Fulton-Suri, 1999). This method works well for the purposes of discovering new strategic product lines, but it takes longer and sometimes it is difficult to translate the broad and diversified data into product designs.

User diary. A user diary can also be used to elicit knowledge. The basic idea behind a diary is to observe, record, and evaluate actions of the user over a period of time (Rieman, 1993, 1996). It is based on a real-time tracking system. Usually, the researcher gives the user a diary in which to record events related to the task, as well as thoughts and insights regarding those events. The researcher can provide a camera so the user can take pictures that would complement the user's notes. Another way of obtaining a diary is to have the user wear a wireless video camera to record performance while engaging in the task. Although user diaries can provide detailed information, the process must not be invasive or difficult to implement, or users will not keep detailed records. One additional, negative factor is the tendency of users to delay entering the information in the diary until a "convenient time." This delay removes some of the insights that users make as they perform the task.

Concept sorting. Concept sorting is used when the goal is to determine the relation between a fixed set of concepts. Usually, the potential user is given a set of cards with a concept word written on each card. The cards are shuffled to randomize the order of concept words. The user is asked to sort the cards into groups and then to provide a representative name for each group of cards. The number of groups that the user is asked to generate may be fixed or left to the user's discretion. This task is repeated many times by many users. The goal is to determine multiple user views on how the concepts should be organized. Because the cards must be labeled with the appropriate concepts prior to the card-sorting session, the knowledge elicitation engineer must be familiar with the domain or have the categories generated from people who are familiar with the domain.

Log files. The aforementioned techniques are derived from asking different user groups to provide information on what they know or how they would behave. With log files, users are not explicitly asked to provide answers or descriptions for certain tasks. Instead, data are recorded from the users' behavior. The actions of the users are recorded as the users interact with the system, and these actions are analyzed to determine trends and predict behavior. The benefits of log files are that a large amount of data can be obtained from a variety of users, and the data collection process does not interfere with how each user would normally interact with the system. The drawback of log files is that irrelevant or incorrect data may be logged or important behaviors may not be logged. Furthermore, the data do not reflect the cognitive processes in which the users were engaged when performing the logged actions.

Validating Knowledge

One way to rectify the problem of different experts using various solution strategies is to use several experts of different types. Multiple experts are more likely to have expertise across the entire domain. Agreement in the knowledge extracted from the experts may be used to validate the knowledge. When experts are tested as a group, any incorrect information is likely to be recognized and corrected. Another way to solve this problem is by looking at the user tasks and choosing appropriate knowledge networks that will most likely facilitate the user tasks. The resultant knowledge networks may include aspects from both the software design experts and the human factors experts.

The fact that much of the knowledge of experts is tacit poses a particular problem for techniques that rely primarily on verbal reports. Cooke (1994) suggested that two methods can be used to ensure that the elicited knowledge is accurate and complete. The first is to use multiple knowledge elicitation methods. The general idea is that any shortcoming of one method will be overcome by the use of other methods. The second technique is to incorporate the elicited knowledge into a prototype expert system and then to have experts examine and interact with the expert system to evaluate whether the knowledge is accurately represented. A third solution, suggested in the white paper session, is to observe experts performing in their work domains and not simply rely on their verbal reports of what they do. This may require an additional evaluator, but the added insights gained from an extra observer may outweigh the additional cost. This technique may also allow an observer to stop the expert and ask the expert questions about a task or technique that may be important but may not have seemed so to the expert.

An Example of Using Knowledge Elicitation Techniques

A project by Vaughan, Candland, and Wichansky (2001) provides an example of how various knowledge elicitation techniques can be used to elicit information that is to be conveyed on a Web-based application. Vaughan et al. were asked to design an application that is accessible through the Web and used by two groups of users (professional buyers and accounts payable). The goal of a professional buyer is to negotiate and make purchases for a company. The goal of an accounts payable person is to pay a company's bills. Therefore, although they may use the same information to complete certain tasks, they perform distinctly different tasks with that information.

Three knowledge elicitation methods were used to identify the types of information and transactions each user group wanted in the Web application, the organization of this information, and which features should be presented at different levels of the user interface design. The team conducted a wants and needs analysis to determine what types of information the users wanted in the Web application, a card-sorting task to determine information organization, and a questionnaire to capture the frequency with which each group performed certain tasks to determine which features should be presented at which levels of the user interface.

For the wants and needs analysis, participants were asked to brainstorm about the content or information that they would like to have included in a system to help them manage their orders with vendors. Participants were asked to determine the five most important pieces of information and to give a reason why that information was important. This resulted in a prioritized list of product content categories and features for each user group.

For the card-sorting task, each participant was given a set of cards on which the name of one conceptual object, derived from the wants and needs analysis, was printed. Participants worked independently to sort the cards into groups that were meaningful to that person's work experience. The participants were allowed to add additional conceptual objects to the group. Cluster analysis was conducted on the card-sort data for each user group. Results showed that both user groups had similar and different items identified in each group of cards. For example, both user groups indicated that *user profile, logon,* and *preferences* should be organized together, but professional buyers also indicated that this material should be organized with *advertising* and *news* information.

A questionnaire listing a set of activities that the software was intended to support was given to the user groups. The user groups were asked to indicate the frequency with which they performed each activity. Items that were identified as being of high priority were those identified by participants as being used on a daily or weekly basis. These items served as the basis for the layout of the interface. For example, information that was important to both groups, such as shipping information, had shortcuts on the home tab to quickly take users to the needed information.

2.3. Summary and Future Research

Because users have certain mental models or scripts regarding the processes that are involved in a particular task, being able to elicit knowledge from experts and targeted users is an important part of designing systems. The mental models for different user groups differ, as do those for expert and nonexpert users. Therefore, eliciting information regarding the targeted users allows developers to predict the knowledge structures that the users bring for interacting with the system, and to design products in a manner that is consistent with or predicts their actions. Knowledge elicitation is not a simple task, and although there are many methods available for knowledge elicitation, a single best method cannot be identified because there are too many variables involved. Furthermore, the same techniques are not applied in the same way for every case, and different methods or different combinations of methods are used depending on the goals of the developers. Table 1 summarizes the major strengths and weaknesses of each knowledge elicitation method described in this section to help determine which methods will suit particular needs.

To reduce the problems associated with knowledge elicitation techniques, more research is needed on the conceptual content and structure of different domains. It would be helpful if both researchers and practitioners engaged in comprehensive, descriptive research regarding users' concepts and conceptual structures for some key domains looking to move to the Web or e-commerce. Healthcare is an excellent example of this. Companies are talking about making an individual's medical records portable (e.g., e-records). This can be accomplished from a technical standpoint. However, little is known about issues such as the following: What is the key information doctors need to see in a patient's record? How would patients want to use and view their own medical information? To what information should patients be able to control access by their employers or insurance companies? Before designing mobile or Web-based versions of such technology, research also needs to be conducted on understanding how people currently structure and use this information.

Although identifying both content and methods appropriate for different types of knowledge domains and user profiles is important, research should also seek to address the aforementioned issues from a methodological standpoint. That is, research is needed to identify the methods that are best suited for eliciting different types of knowledge because some knowledge elicitation techniques seem to work well for limited problems and domains. For example, questionnaires are a good method to identify users' preferences, but they are not a good method for understanding users' knowledge representations. Research should also be conducted to identify the methods best suited for eliciting information from different user profiles. The same methods may not be used to elicit knowl-

Table 1: Knowledge Elicitation Methods for Recommended Groups and Their Major Advantages and Disadvantages

Method	*Group*	*Major Advantages*	*Major Disadvantages*
Interviews	Experts and user groups	Most well-known method for eliciting knowledge Qualitative data	Time consuming Expensive
Verbal protocol analysis	Experts	Qualitative data Document thought process related to performance	Time consuming Hard to analyze
Group task analysis	Experts	Obtain different viewpoints Document thought process related to performance	Developed through trial and error No research validating this method
Narratives, scenarios, and critical incident reports	Experts and user groups	Provide insight to reasoning processes and implicit knowledge Good for ill-defined problems	Reliance on self-reports
Questionnaires	User groups	Quantitative data Easy to code	Low return rate Response may not correspond with actual behavior
Focus groups	User groups	Allows exchange of ideas Good for generating lists of functions and features for products	An individual may dominate the discussion Not good for discovering specific problems
Wants and needs analysis	User groups and experts	Exchange of ideas Determine areas of focus	What users say they want and need may not be realistic
Observation and contextual inquires	User groups	Studied in natural environment Qualitative and quantitative data	Time consuming Dependent on detailed notes of the observer
Ethnographic studies	User groups	Studied in natural environment Good for discovering new products	Time consuming Hard to generalize results to other product designs
User diary	User groups	Real-time tracking Qualitative data	Can be invasive or difficult to implement Delay in entries by users
Concept sorting	User groups and experts	Determine relations among components Helps structure information	Grouping may not be optimal Resulting structure may be too elaborate
Log files	User groups	Uses actual recorded behaviors Can collect data from a range of users	Irrelevant or wrong information may be recorded Data do not reflect cognitive processes

edge from experts and novices or from a busy CEO and an administrative assistant. Furthermore, research to determine whether knowledge elicitation techniques are generalizable across domains or relatively domain specific will also be valuable.

Most knowledge elicitation techniques have been developed in the context of eliciting knowledge from experts. However, when the goal is to design for target users who range in expertise, knowledge must be elicited from a large range of users. Therefore, future research should focus on evaluating and further developing methods that allow easy and rapid elicitation of knowledge from many users. For example, the methods for evaluating narratives should be developed more because narratives provide people with a method of building plausible explanations as guides for controlling events in their work domains. Another potential method that should be developed further is the use of log files to elicit knowledge from a large data set. With log files, users' behaviors can be analyzed to determine the frequency of certain actions and patterns of behaviors.

A major task for future research is to determine standardized ways of eliciting knowledge for specific goals. As a first step toward accomplishing this task, it would be valuable to survey in detail the knowledge elicitation studies conducted in different fields to determine which methods are most effective for different tasks and goals. The overall goal should be to determine which combinations of knowledge elicitation methods yield the most complete sets of information for specific applications.

3. ORGANIZING AND STRUCTURING INFORMATION

Organizing and structuring information is central to the success of content preparation. The notion of designing a system that organizes materials in a logical and efficient way is not new. Traditionally, paper documents are stored and filed in file cabinets. This method of storing and filing paper documents takes up a lot of space, is time consuming, and makes retrieving documents difficult. With the invention of the modern computer and graphical user interfaces, user interface designers used the "file cabinet" metaphor. This system assigns a single file name to materials to represent a distinct unit of information. The user organizes and classifies this file in a hierarchical structure of directories and subdirectories (see Nielsen, 1996). For a user to access a particular piece of information, the user must recall where that information is located in this structure or perform a search for it. This method is not very problematic if the information is organized and accessed by a single user. However, because users may not remember or may misremember where an item was stored, the time spent searching for a particular piece of information could be great.

The term *information architecture* is used to refer to the practice of organizing and structuring information to create a highly usable navigation structure. To create a successful structure, an information architect must know (a) the business context, (b) the nature of the content, and (c) the information needs of the users. Each business has different goals, activities, and organizational structures that need to be reflected in the structure and organization of the information. The content may differ in terms of the types of documents, the degree and nature of the structure of the information, and the format in which that information is stored. The major types of users can vary in terms of computer expertise, domain expertise, information needs, and more. The information architect must design the information structure to meet the needs of the business and the intended users. In sum, the organization and structure of information should reflect the context, content, and users.

3.1. What Are Current Issues Associated With Structuring and Organizing Information?

Organizing for the Context of the Web

The problem of making information available in multiple places without having multiple files has been reduced by modern technology. For example, hyperlinks in the Web and shortcuts in typical graphical user interfaces allow information to be accessed in multiple locations, but stored in a single location. However, because so much information is available on the Web, the task of organizing the information is difficult, and designers must ensure that the structure of the site can accommodate storage of this material. Moreover, because the information is constantly changing, organizing the material is a continuous task. Therefore, although there is a need to make the same information available in different places, the structure and organization of this information is also important.

Content Structure Often is not Based on a Theory of Meaning

Many approaches to task analysis focus on the procedures for performing tasks. These approaches typically compare the actual behavior with prescribed or normative procedures. In this context, an error is defined as an action that deviates from the prescribed procedures. However, experts often utilize their understanding of a process to take shortcuts (deviations from standard procedures) that save time or effort. Also, a role of the human expert is to adapt to situations that were not anticipated in the design of procedures. For these two reasons, it is important not only to consider procedures, but also the constraints (e.g., physical laws, organizational and regulatory constraints) that must be respected to generate new procedures (i.e., to take advantage of local opportunities—shortcuts, or to adapt to unexpected events). In other words, it is important to consider the deep structure that allows the expert to respond generatively or adaptively to situations. From this perspective, the challenge is to discover an objective or situation-referenced theory of meaning.

In sum, attention should be shifted from the surface features of work (e.g., procedures) to the deep structure (e.g., functional constraints such as goals and bounds on safe operation). Designers need to make the deep structure explicit to workers through interfaces, instruction, and training. The deep structure is more easily specified in systems whose functions are bounded by known (and typically invariant) physical constraints (e.g., aircraft or nuclear power plants). In these cases, the structure includes the physical laws and engineering principals that determine the work processes. However, in domains where the constraints are more ambiguous social constructions (e.g., Web search engines or e-commerce sites), the deep structure may be more amorphous and may vary with the idiosyncratic goals of specific users or customers.

Taking End Users Into Account

Avoiding getting users lost. With all the information available to users, the information needs to be structured in a manner that is simple and makes sense to users. Users must be able to comfortably and confidently navigate through the site (Rosenfeld & Morville, 1998). If a user gets lost in a Web site, the user will not be able to find the information for which the user is looking and will have to spend extra time trying to get out of the information structure. Users simply cannot always be expected to build elaborate mental maps of an entire Web site.

Communicating the structure of the site to the users. A possible reason why users may poorly understand how the information in a Web site is structured and organized is that designers do not communicate the structure of the site to users. If users understand how the Web site is organized, then they can more quickly, easily, and confidently move between sections and subsections of the site (Fleming, 1998). There are techniques to help communicate a site's structure (e.g., obvious major section navigation, obvious subsection navigation, navigational breadcrumbs), but these techniques are not used on all sites.

Satisfying customers. The organization and structure of the information must satisfy the needs, objectives, and preferences of users. Because users are not adept at organizing information in a Web site or designing user interfaces, the temptation to let users design the site must be resisted. Although input from users should be used to guide the design process, design experts, such as information architects, should ultimately be responsible for organizing the information in a Web site.

3.2. What Methods Are Being Used to Structure and Organize Information?

Methods Relevant to Web Design

Hyperlinks. With the vast amount of information on the Web, users retrieve the information they desire by navigating through the structure of the Web site.

Hyperlinks allow users to retrieve relevant information by "pointing" them to the location. When a user clicks on a hyperlink, it takes the user automatically to a Web page that displays the information. Therefore, users do not need to know anything about the structure of the site to reach their desired locations.

Extensible Markup Language (XML). XML is an application of the standard generalized markup language that was introduced in 1996. XML documents are composed of storage units of information. XML can be used to specify the content and structure of information in a way that allows particular pieces of information to be retrieved when needed. For example, a piece of information can be identified as a price, date, or product name. It provides a mechanism to impose constraints on the storage and structure of information. Because the content and structure of the information is separated, it allows the use of the information in a variety of contexts, the exchange of database information, the customization of information, and the management of large databases. XML provides a good architecture and schema for structuring and organizing information, however, it is an open standard in which individual implementations can vary (Schmelzer, 2001). Also, XML is a programming language that organizes information at the back-end computer level. Designers must still organize the information for the user interface.

Formats and languages that would support different types of queries. When organizing and structuring information, it is important to do it in a way that optimizes the characteristics of the site's search engine. For example, if the search engine supports semantic queries, then the information in the site should have some type of semantic organization or categorization that facilitates semantic queries. The "Semantic Web" was developed to link information based on its properties. According to the Semantic Web Agreement Group, "The Semantic Web is a Web that includes documents, or portions of documents, describing explicit relationships between things and containing semantic information intended for automated processing by our machines" (2001). For example, the Semantic Web can link all files with the ".doc" extension. The Semantic Web can also be programmed to include unicode and XML, basic assertion languages, schema languages, conversion language, query languages, and so on. Allen (n.d.) noted that a Semantic Web can also be used to gather information about a site from the owner of the site, make comments about a site (e.g., give the site a thumbs-up or thumbs-down), see how your friends rate the site, bookmark pages, and so forth.

Interactive navigation displays. Interactive navigation displays use navigation to tell the user where the user is, how the user got there, how to get back, and where else the user can go (Fleming, 1998). Navigation displays can provide a graphical display of the user's current position in the structure, and provide information about how the user reached a particular Web page and where he or she is in the structure of the Web site. This information can help users learn about the structure of the Web site and its organization. The following are types of interactive navigation displays:

- Site maps: Site maps provide an outline of the Web site and illustrate where particular Web pages are located in the Web site.
- Navigational breadcrumbs: Breadcrumbs are small, hyperlinked page titles usually located at the top of each page, above the title of the current page. These hyperlinks show the page titles that the user came through to get to the current page (Najjar, 2001).
- Persistent major navigation controls: The site should show on each page the navigation controls for the major sections of the site. Each page should show other navigation controls that are important to users (e.g., Contact Us, Security, Privacy). Each page should show the navigation controls to access subsections of that section.

Standards on certain platforms. The benefit of providing a set of standardized conventions (regardless of how usable they actually are) is that they impose a measure of consistency on the user experience. This is perhaps one of the most important tenets of good user interface design. There are user interface design guidelines for the Web, personal computers, such as the Windows system and the Macintosh, UNIX, cellular telephones, and Palm Pilots. These guidelines encourage designers to create user interfaces that are consistent for that platform. Standards are statements (i.e., requirements or recommendations) about interface objects and actions that can be used to address the physical, cognitive, and affective nature of Web interaction. They are written in general and flexible terms because they must be applicable to a wide variety of applications and target user groups.

Table 2 lists several standards developed by specific organizations. Buie (1999) provided recommendations on how to use such standards. These include the following:

- Select relevant standards from existing sources. The standard chosen should match the problem of interest.
- Tailor the selected standards to the specific Web development project.
- Refer to and apply the standards as closely as possible in the Web design.
- Revise and refine the selected standards to accommodate new information and considerations that arise during development.
- Inspect the final design to ensure that the Web design complies with the standards.

For example, on Windows, the organization of the files is visually represented and navigated via the Windows Explorer, with the top-level objects being shown as disks that contain folders and various types of files. Clicking on a folder on the left pane of the Windows Explorer window displays its “content” on the right pane. Each application window has a system menu located at the top of the window. When the user points and clicks on an item in the system menu, that part of the application is brought to the front of the window. Table 3 contains sample guidelines

Table 2: Examples of Standards and Guidelines

Type	*Organization*
International	International Organization for Standardization (ISO 9241; Ergonomic Requirements for Office Work with Visual Display Terminals)
National	American National Standards Institute (ANSI)
	Human Factors and Ergonomics Society (HFES–200; Ergonomics of Software User Interfaces)
	British Standards Institution
	Ente Nazionale Italiano di Unificazione
	Deutsches Institut fur Normung
Government	U.S. Military (MIL–STD–1472D, Human Engineering Design Criteria for Military Systems, Equipment and Facilities; MIL–STK–1472D, HyperCard version)
	National Aeronautics and Space Administration (NASA–STD–3000, Man–Systems Integration Standards Handbook)
	MITRE Corporation (developed for the U.S. Air Force, ESD–TR–86–278, Guidelines for Designing User Interface Software)
Industry	Macintosh Human Interface Guidelines
	Common User Access; IBM HCI Guidelines (CUA IBM)
	Open Software Foundation (Motif Style Guide)
	Microsoft (The Windows Interface Guidelines for Software Design)
Individuals	Deborah Mayhew (Principles and Guidelines in Software User Interface Design Guidelines)
	Lawrence Najjar (1990, Using Color Effectively; 1992, Multimedia User Interface Design Guidelines)
	Jakob Nielsen (1999a, Designing Web Usability: The Practice of Simplicity)
	Wanda Smith (ISO and ANSI Standards for Computer Products)

Note. Based in part on Buie (1999).

from Najjar (2001) to consider when structuring and organizing the page format and navigation for an e-commerce Web site.

Methods for Organizing and Structuring Content

Objects–actions interface model. Shneiderman (1997) proposed that the problem of structuring and organizing information for the Web can be resolved in part by decomposing the complex information into a hierarchy of manageable units. For this purpose, he developed an objects–actions interface model in which objects and actions in the domain of the task are hierarchically decomposed into basic elements so that they can be represented by interface objects and actions (see Figure 3). According to Schneiderman, the objects–actions interface model encourages Web site designers to focus on two components of the task, structured information objects (e.g., networks) and information actions (e.g., searching); and two components of the interface, metaphors for information objects (e.g., bookshelf) and affordances for action (e.g., querying).

Table 3: Example of How to Structure a Web Page and Organize its Navigation: Sample Guidelines to Consider When Structuring and Organizing the Page Format and Navigation for an E-Commerce Web Site

Navigation	*Page Format*
Design navigation that is simple, intuitive, and obvious. Put the navigation controls in the same locations on each page. Use navigation to tell the user where the user is, how the user got there, how to get back, and where else the user can go. Provide "breadcrumb" navigation on the site. Breadcrumbs are small, hyperlinked page titles at the top of each page, usually above the title of the current page. The breadcrumbs show the user the pages that the user viewed to get to the current page in the current major section. Limit the number of major sections in the site to about seven. Provide an intuitive name for each major section. Design the navigation so that the user can browse to any product in five clicks or less. Provide some specialized browse functions to meet specific user needs. For example, allow users to browse through suggested gifts or a special section with discounted products. Provide search entry fields on nearly every page. To improve search, use meta-tag tools	Design the page size to accommodate America Online (AOL) members. There are over 29 million AOL members, and they buy at three times the rate of other users. Because AOL users open their browsers inside the AOL service window, the space available for the browser window is reduced. Therefore, for an 800 × 600 screen resolution, use a 625 × 270 page size. Because download time is the biggest problem with the Web, design your pages to download quickly. To get pages to download in 10 sec or less over a 56K modem, limit the size of each page to 40K or less. Except for the product comparison tool, never require the user to scroll horizontally. Avoid forcing the user to scroll vertically on the home page. Format the pages so that user interface elements are in locations that are familiar to users (e.g., put horizontally across the top of the page the navigation controls for the major sections of the site). Except for the shopping cart page and checkout pages, put a shopping cart summary on every page. To easily accommodate registered members, include sign-in entry fields on the home page and a sign-in hyperlink on every page. So that signed-in members know that they are signed in, show a personalized welcome message at the top of the page. Be sure to put a "Contact Us" hyperlink on every page, perhaps at the bottom. Provide at the bottom of each page a link to a simple referral form (e.g., "Tell a Friend"). Provide links to the privacy and security policies at the bottom of every page.

Note. Based on Najjar (2001).

Although decomposing the domain into a hierarchy with basic end elements, or atoms, sounds reasonable, determining the organization can be difficult. Shneiderman (1997) noted:

> While many 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, etc., and libraries have magazines, videotapes, special collections, manuscripts, etc. 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. (p. 16)

The basic objects and how they are aggregated for a Web site may vary across different domains. According to Shneiderman (1997),

> Atoms can be a birthdate, name, job title, biography, resume 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. (p. 16)

Due to these differences, it is important to have clear definitions that can be conveyed to designers and to users. The information atoms can be organized in several ways: short unstructured lists, linear structures, arrays or tables, hierarchies or trees, faceted retrieval (e.g., photos indexed by date, photographer, location, topic, film type), and networks.

Elemental information actions include scanning a list of article titles, looking for a specific word in an alphabetized list, following a link, and reading a paragraph.

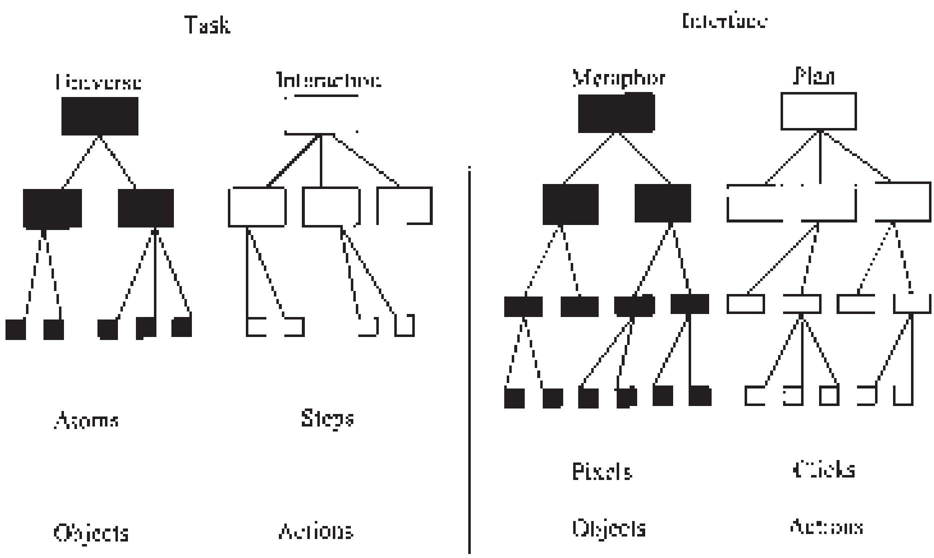

FIGURE 3 Objects–actions interface model: Hierarchical decomposition of the task and interface into basic objects and actions. Based on Shneiderman, *International Journal of Human–Computer Studies*. Copyright (1997) by Harcourt, Inc. Adapted with permission.

These elemental information actions can be combined into aggregate information actions such as locating a term in an alphabetic index and reading abstracts of articles that contain the term. Once the elemental information objects and actions for a task, and their structure, have been determined, these can be mapped into specific interface objects and actions.

Ecological interface design. The ecological approach to human factors starts by considering the constraints in the task environment that are relevant to the user (Vicente, 1999). Within this approach, a theoretical framework for interface design, called *ecological interface design,* has been developed (Vicente & Rasmussen, 1992). This framework is based on Rasmussen's (1986) influential behavioral taxonomy. He proposed that for familiar events behavior occurs in a relatively automatic, skill-based mode. For example, if users interact with a system on a daily basis, the routine commands become automatized and are executed with little thought. Rule-based behavior is used when the user encounters a novel situation to which rules they have learned previously can be applied. When e-commerce customers enter sites that they do not visit frequently, the rules that they have learned from interacting with other sites will guide their behavior. Finally, behavior is knowledge-based when an event is unfamiliar, no simple rule applies, and the user must engage in problem solving. For example, when a customer submits billing and shipping information and the site returns an error message that is not specific, the customer must develop and evaluate hypotheses about where the error occurred.

The basic idea behind ecological interface design is that the information that needs to be represented in a good display varies as a function of the appropriate behavior mode. The principle for skill-based behavior is that the user should be able to respond directly to the display, with the structure of the displayed information being independent of the required actions. For rule-based behavior, the interface should signal the appropriate rules to be applied with a consistent mapping between the constraints of the work environment and the information provided by the display. Finally, for knowledge-based behavior, the work environment should be represented in the form of an abstraction hierarchy (a hierarchy of information relevant to problem solving at several levels of abstraction) that provides an externalized mental model in support of problem solving.

Information theory. One of the oldest means for quantifying the structure of information is information theory, which assumes that information can be characterized into bits, or units, as a function of uncertainty (Garner, 1962). Research has shown that the time it takes to understand information is highly correlated to the number of bits subsumed by the information (Fitts & Posner, 1967). Garner (1962, 1974) devoted his career to studying the role of structure in perceptual identification. He developed ways to describe perceptual displays and ascertain the difficulty in perceiving them. Among the concepts stressed by Garner (1974) was the importance of inferred subsets. An inferred subset is the set of stimuli that the observer considers to be likely to occur. Specifically, Garner (1974) showed that the smaller the inferred subset of alternatives from which a specific pattern is derived, the greater the redundancy and the easier the pattern is to perceive.

Information theory can be used to quantify how much information a response carries about the stimulus (i.e., information transmitted) and the rate at which this transfer occurs. This can be compared to the information transferred in particular models of human performance and to maximum possible information transfer. Although not used as widely as it once was, information theory is still used to quantify the efficiency of displays (e.g., Tan, Durlach, Reed, & Rabinowitz, 1999) and neural coding (Borst & Theunissen, 1999), and to develop hierarchical decision trees for machine learning (Ben-David & Utgoff, 1995).

Discourse processing and propositional representation. Kintsch (1974, 1998) spent the last 30 years studying discourse comprehension in detail, proposing a variety of ways to characterize the ideas represented in sentences and paragraphs. Van Dijk and Kintsch (1983) put forward a theory of textual information processing known as strategic information processing. They suggested that a user relies on general background knowledge and initial surface cues of text to make meaning out of the text. These surface cues are extracted from words, sentences, paragraphs, headings, and so on. At the same time, the reader is attempting to apply whatever genre knowledge he or she may have, such as the structure of news stories, to make sense of and strategically process the incoming information to more efficiently build up a mental model indicated by the semantics. Van Dijk and Kintsch emphasized the importance of "structural cues" for information organization, from both bottom-up (e.g., sensory input) and top-down (e.g., high-level expectancies) processing because they help novices learn and experts be more efficient.

Van Dijk and Kintsch (1983) concluded that the meaning of text is represented best by networks of propositions, each of which is a predicate–argument schema. An atomic proposition is of the form PREDICATE [ARGUMENT, ARGUMENT, …], for example GIVE [agent: MATT; object: VIDEOGAME; goal: SCOTT]. Complex propositions are formed from compounds of atomic propositions, and texts consist of one or more compound compositions that can be related in several different ways. Propositions derived directly from a text make up the textbase, which combines with propositions contributed from the individual's long-term memory to form a situation model. Van Dijk and Kintsch also distinguished between the sentence-by-sentence information of the text, called the *microstructure,* and a hierarchically organized set of propositions representing the global text, called the *macrostructure.* An appropriate summary of a text should reflect the text's macrostructure. Propositional analysis has been shown to provide a useful way for analyzing many aspects of text structure, as well as for describing how people comprehend, remember, and respond to questions about texts (Kintsch, 1998). However, a limitation with respect to practical application is that propositional representations have to be constructed through hand coding and therefore are difficult to use for large applications.

Latent semantic analysis. The limitation of coding propositions by hand can be overcome by using statistical methods to analyze the meaning of text through the use of words in a variety of contexts. Latent semantic analysis is one such formal model developed to simulate human understanding of how words and combinations of words form meaning in a passage (Landauer, 1998; Landauer, Foltz, & Laham, 1998; see also the closely related Hyperspace Analog to Language Model; Burgess, Livesay, & Lund, 1998). Latent semantic analysis represents words and larger text units (e.g., propositions) as vectors in a high-dimensional semantic space that is constructed from statistical computations performed on a large corpus of text. The analysis is based on considering the frequencies with which words occur in various discourse contexts (e.g., in the same sentence, in sentences with the same words, etc.). Each word occupies a position in the semantic space and is associated with other words through links (common labels or features). The closer two words are located in this space, the higher the association is between them. With this structure, latent semantic analysis can determine similarity relations between words, sentences, paragraphs, and essays, even if the word sets themselves are not taken from the original corpus. This has been demonstrated in a number of ways, a few of which are mentioned here.

Latent semantic analysis has been used to provide measures of the coherence of texts by comparing the similarity for adjoining text segments (Foltz, Kintsch, & Landauer, 1998). The general idea is that a text is more coherent when the adjoining segments are more similar in meaning on average than when they are less similar. Foltz et al. demonstrated that the measures of text coherence predicted the relative comprehension shown by readers of the texts. Therefore, latent semantic analysis provides an automated method for evaluating text coherence. Latent semantic analysis also was able to predict learning from four texts of different difficulty as a function of the match between the reader's conceptual knowledge and the level of difficulty of the text (Wolfe et al., 1998). Learning was best when the text was neither too easy nor too difficult relative to the reader's knowledge. These examples illustrate that latent semantic analysis provides a valuable tool for organizing and structuring conceptual knowledge in a way that will be appropriate to the user's knowledge level. The site, http://LSA.colorado.edu/, contains documentation, demonstrations, programs, and manuscripts to assist designers in becoming familiar with this sophisticated technique.

Multivariate techniques. Multivariate statistical techniques, such as multidimensional scaling and the Pathfinder networks, can be used to analyze and organize information. Multidimensional scaling provides a representation of the dimensions on which people organize a set of concepts. In Pathfinder networks, the concepts are represented by nodes and their relations by weighted links. A Pathfinder analysis tends to provide more information about pairwise relations between concepts, whereas multidimensional scaling analysis provides more information about global relations (Cooke, Durso, & Schvaneveldt, 1986).

Identifying methods to "test and evaluate" structure and content. Testing for structure is harder than it sounds because there are many components to the Web site, and each component may need to be tested separately. The test for some components is easy. To test whether a link works, click on the link. Other components are harder to test. To test whether the structure is conveyed in a meaningful manner to users, many narrative reports must be taken and integrated, if possible. Furthermore, even if each component can be tested individually, the components may have interactive effects when placed in the system context.

Methods Involving Users

Concept sorting. As discussed in the knowledge elicitation section, concept sorting can be used to determine relations between fixed sets of concepts. When users are asked to sort concept words into discrete categories, the categories can be used to help determine how the information should be organized. For example, the major categories can be organized by frequency of use or importance on the homepage, whereas the concepts within the categories can be organized into subsections.

Other user-testing methods. A Web site should be organized and structured in a way that will allow end users to achieve their goals. As a result, end users should be involved in the design process. Methods such as usability testing (Mayhew, 1999; Nielsen, 1993), iterative design and testing cycles (Shneiderman, 1992), wants and needs analysis, and group task analysis can be used to help structure and organize a Web site. These methods are described in detail in other parts of the article, and they will not be elaborated on here.

An Example of Structuring and Organizing Web-based Designs

The Web has evolved from a simple browsing environment to an interactive, goal-oriented, task-driven application and script-enabling platform. More and more Web applications are being developed to allow users to create, edit, and manipulate objects in the browser, just like traditional graphical user interface-based applications. The subsequent statement from Zhu (2001) describes how a Web-based application, SharePoint™ Team Services from Microsoft®, is structured:

> It [SharePoint] is a Web site, it should look like a typical Web site. The underlying data structure of the Web site is a list structure, and there are three types of lists that users can use: Documents, discussions and lists. A list that holds documents is called a document

> library. A list that holds discussions is called a discussion board. A list that holds any other items is called a custom list. Users can create multiple document libraries, discussion boards and lists. Within each list, users can create multiple items. (p. 839)

The design of SharePoint is based on assumptions about different user activities and the relative frequencies with which users engage in these activities. Because SharePoint is designed to help people collaborate and share information, it is assumed that users will most often consume information, rather than add or edit information. The Web site is therefore optimized for browsing behavior, and edit controls are hidden and only visible after a deliberate action from the user.

There are four basic types of pages: the homepage, the list and item view page, the customization summary page (the page that shows all the options for a list), and the customization controls page. The list view page and the item view page have editing controls in a special toolbar area above the content area. Each page has four areas: navigation bar, title area, left pane, and content area. The Web site has one main navigation bar on the top that is always visible, which contains links to the different sections on the Web site: documents, discussions, lists, create, site settings, and help. On the list view page, the left pane contains links that filter the list (also called views); and on the homepage, the left pane contains quick launch links.

3.3. Summary and Future Research

The problem of organizing large amounts of information in an efficient and manageable manner is not new, and this problem has been magnified with the introduction of the Internet. The Internet allows users to publish information on the Web from anywhere at any time. Once this information is posted to the Web, it can be modified, moved to a different location, or even deleted. This ability to constantly change the information stored and add to the Web results in the need to design an interface that structures the information in a way that can be reorganized quickly and is manageable.

To improve the way information is structured and organized for the Web, future research should focus on determining the best tools for organizing, updating, and maintaining information. Hyperlinks were initially used to solve the problem of having to update and store the same material in different locations. In recent years, there has been increased interest in using mark-up languages like XML to help structure and organize material because XML can specify the nature of the information using predefined categories or elements, or be customized in a manner that is consistent with the needs and goals of the designers. Therefore, the potential of using XML as a tool for organizing and structuring information in the future is great, and many companies that were mentioned in the introduction of this article are currently using XML.

Interactive navigation displays seem to help users and should be available when possible. Site maps and navigation breadcrumbs can be used to help convey the structure of the site to users so that they can develop a general "feel" for where items may be located in the site. Persistent and consistent major navigation controls are important because the same functions are incorporated into many Web sites causing users to expect the controls to be placed in a location that is consistent with their previous experience. However, research is needed to evaluate the alternative ways of conveying the structure of the site to users. In addition, evidence needs to be gathered about the relative benefits of interactive navigation displays for Web sites of different sizes and types.

Progress has been made in developing methods for organizing and structuring content. A major problem that remains is how to determine the fundamental elements of information for specific tasks, and research along this line needs to be continued. Information theory has not been used much for interface design in recent years, but because it offers a way to quantify information and has been successfully used in a variety of disciplines, such as electrical engineering and neurobiology, further exploration of its use for interface design seems warranted. Perhaps the most promising approach at present is the statistical approach exemplified by latent semantic analysis. The value of this approach for HCI is that analysis can be performed automatically on demand. Initial tests of its usefulness for applications have been positive, but much remains to be examined.

Finally, end users should always be involved in the design of a Web site because the successfulness of the site is often measured by the ability of the end users to interact with it to achieve a goal or task. User testing should occur throughout the design process, and input from users should be incorporated. For example, card-sorting techniques can be used initially to determine items that should be grouped and categorized together. Once these categories are incorporated into the Web site, usability tests can be performed to determine whether some items were overlooked or should be moved to different locations. As the design process continues, iterative testing cycles can be used to determine whether modifications to the structure and organization of the site affect performance. Because of the time and cost involved in conducting user testing, detailed case studies of the relative importance of user input at different stages of the design process should be conducted.

4. RETRIEVAL OF INFORMATION: IMPROVING WEB SEARCH

Web search engines serve as catalogs of the Web. They index the Web pages by deploying a special computer program called a *spider* or *robot.* Users can enter search terms to query against the index database. The search engine returns a list of Web pages along with short descriptions. If the search is effective, the returned results of the query will consist of the information that the user needs; however, if the search fails (by not returning the desired hit or returning too many irrelevant hits), the user must engage in another, different attempt to retrieve the information.

4.1. What Are the Current Problems With Search Techniques?

In this section, problems preventing efficient and effective retrieval of information are discussed. These problems can be characterized as problems involving the search engines themselves or problems involving the end users interacting with the search engines.

Search Engines

If the search criteria programmed into a system are too strict, users will not be able to locate the items for which they are looking. On the other hand, if the search criteria are too lenient, users will be bombarded with irrelevant hits. The following are some of the major problems with search engine designs.

Poor back-end design. Many search engines are not designed to be flexible. That is, these engines do not accommodate typographical errors, synonyms and variant forms of keywords, recognize Boolean operators, and so forth. For example, it is known that users make typographical errors: They may not know the proper spellings of low frequency words, or they may make a typing error in the entry. Search engines should include a spell checker, and commonly misspelled words should be linked to the correct words. In addition, users do not always use the same keyword to look for similar products, and search engines should provide a good thesaurus to help return relevant matches. For example, if the goal of a search is to find a person who treats injuries, users may use any one of the following keywords: doctor, physician, medicine, hospital, clinician, healer, and so on.

Irrelevant hits. The returned Web pages are usually ranked based on some type of relevance measure designed for the search engine, although the ranking is often inaccurate because of the difficulty of defining relevance. The large volume and lack of structure of the information make it very difficult to index Web pages efficiently. These factors significantly reduce the power of Web search engines to retrieve useful information requested by a user. Inefficient indexing and inaccurate search queries can easily result in millions of hits for a single search query (Fang & Salvendy, 2000). These hits can be relevant to the task at hand or completely irrelevant to it. Because users must work their way through the long lists of hits, this task is very time consuming, and information may not be retrieved because users are overloaded with information or the information is "hidden" by the irrelevant hits surrounding it.

Inconsistent functions. Each search engine works differently. When a user enters the same search terms into different search engines, the search engines produce different results. This inconsistency can frustrate and confuse users.

Difficult advanced search options. Advanced search usually provides features such as drop-down menus or Boolean logical operators that allow the user to narrow the span of the search. This feature is recommended for experienced users who know specifically their search target. This option should be made available, but a link should be provided to lead users back to the simple search engine (Nielsen, Molich, Snyder, & Farrell, 2000).

Search engines do not incorporate an objective theory of meaning. There are not adequate theories or descriptions of the workspaces that allow organization and classification of concepts to reflect the important functional constraints (i.e., meaningfulness). An abstraction hierarchy provides a generic structure for building these kinds of descriptions (Rasmussen, Pejtersen, & Goodstein, 1994; Vicente, 1999). However, discovery of the meaningful dimensions of a work domain often requires a global analysis that looks at work in both natural and historical contexts. This is because functional constraints can be integrated within tools and standard work practices. As such, the constraints become implicit, and it becomes difficult for experts and work analysts to parse the deep structure from the "syntactical" aspects of the work. This problem is illustrated by Hutchins's (1995) study of navigation on a U.S. navy ship that analyzed the performance of the navigation crew in detail.

Search strategies used by expert information retrievers are not incorporated. Most search engines do not incorporate the search strategies used by experts. For example, when someone wants to find out information regarding a particular topic, a librarian is a very good person to ask. Although the librarian may not know much about the specific topic being queried, the librarian's knowledge about the structure of the library catalog system will enable the librarian to easily find the information relating to the topic. Therefore, the strategies that expert information retrievers, such as librarians, engage in to help return relevant matches should be given more consideration.

End Users

Lack of knowledge about how users search. The design of Web search engines and interfaces could benefit substantially from the extensive research in the field of information sciences (e.g., Salton & McGill, 1983). That body of research has revealed some understanding of how users search. For example, string search and Boolean logic have long been used in information retrieval in bibliographic systems such as library catalog systems, large online database systems, and CD-ROM reference systems. Despite their popularity and wide availability in commercial systems, string search and Boolean logic have been criticized by information scientists for the following reasons (Bookstein, 1985): (a) Boolean logic is hard to use or understand, especially for expressing complicated expressions; and (b) retrieval output is hard to control. A common or high-frequency search term will result in a large number of documents or references retrieved. Because no other information is given about the relevance of these output documents, users will have tremendous difficulties in reviewing all the results. On the other hand, a phrase query with several search terms connected by the "and" operator may yield no search results at all. Although various techniques such as expanding a search by specifying synonyms provided by a thesaurus or to narrow down a search by adding constraints can be used to obtain a moderate-sized result set, these techniques usually require special training and experience.

The search function is not always immediately apparent to users. When users enter a Web site, they have certain expectations of the general structure of the Web site. With regard to search engines, users expect to find a search engine right away, especially in the form of a text box. Nielsen et al.'s (2000) study found that "users told us that when they looked for the search function, they looked for 'one of the little boxes.' Tabs and links to a separate Search page just didn't work for them" (p. 7). Nielsen et al. also indicated that a Web site should have the search feature available on every page in the site so that users can use the feature at any time.

Error messages are difficult to interpret. Perhaps the most frustrating aspect of a search engine is that after users enter a query, they receive a message that is not comprehensible. Error messages such as "Error: Syntax not found" or "Error: 404 Page Not Found" should not be used. Rather, designers should have the search engine provide constructive feedback (Nielsen et al., 2000). For example, if users entered the query "Oh be one Kenobi," the search engine can display a message indicating that there are no matches for "Oh be one Kenobi," but try "Obiwan Kenobi." If the search engine is programmed to recognize that this character is related to the Star Wars series, the text message should indicate "Star Wars character: Obiwan Kenobi," to explain to the user why this hit may be relevant instead of just returning a Star Wars match.

Perspective of user is not incorporated. Users do not have a standardized way of searching for information, making it hard for designers to program search engines to mimic search behaviors. As mentioned before, different users may use different keywords to describe the same item, and the same user may use different keywords at different points in time. The large range of users and their idiosyncratic search techniques are often not studied because of the time involved in analyzing users' search behaviors. Therefore, a development team focuses entirely on designing search engines with no consideration of users' search techniques, making it difficult for users to obtain the information they need in an effective manner.

Problems relating to browsing versus searching. There are two basic ways to search for information on the Web: browsing and keyword searching (Marchonini, 1995). Browsing seems to be more intuitive than keyword searching, especially for naïve users. Even when they do not know how to conduct a keyword search, users should be able to browse the content of a Web page or the information categories that are usually organized as a hierarchy. However, as the amount of information increases, browsing becomes less efficient and search becomes more practical.

Widening gap between users and search engines. The development of more advanced technology is progressing at a faster rate than the average user can handle. For example, the average user may just now be comfortable with the passive interface design in which the user must enter a command to execute a function. However, the software industry has already been utilizing "agents" to document user activity and execute commands independently from the user by predicting what the user would probably do based on the user's past action. Although the development of agents is needed to help users deal with vast amounts of information, some users are not ready to "give up" control of a process they are just beginning to understand to an agent about which they have little or no understanding or experience.

4.2. What Methods Are Being Used to Improve Search?

Statistical Techniques

Natural language processing. Computers cannot always interpret users' queries in the same way a human can. A major problem that engineers need to overcome is to construct a natural language processing system that can mimic that of humans. To understand a discourse, the computer must be able to evaluate the coherence of alternative interpretations. Theories of natural language processing assume that coherence is based on meaning, structure, and intention (Mc Kevitt, Partridge, & Wilks, 1999): Analysis of meaning is based on semantic relations between words, analysis of structure is based on examining syntactic markers, and analysis of intention has the purpose of determining the intentions and goals of the user.

One potential advantage of the natural language processing approach is that it attempts to determine the intent of the user, and when coupled with user profiling, provides even better results. However, at present, intent cannot be easily inferred from users' queries because most users mainly use keywords for search. The average length of queries is approximately 2.5 terms, which reduces the ability of natural language processors to determine what the user means. In addition, users do not always use Boolean operators ("and" vs. "or") correctly.

Latent semantic analysis. Latent semantic analysis, described in the organizing and structuring information section, can be used to help facilitate Web search. In fact, this analysis method was developed originally to improve access to textual information in HCI (Deerwester, Dumais, Landauer, Furnas, & Harshman, 1990; Dumais, Furnas, Landauer, & Deerwester, 1988). With latent semantic analysis, sentences or concepts can be analyzed to find similar topics based on associative information. Therefore, strings of text can be matched to content on the basis of semantic relatedness instead of keyword overlap. Dumais et al. applied latent semantic analysis to a database of 1,033 medical reference abstracts and showed 13% better average precision in the documents retrieved than for keyword matching. Chute (1991) applied latent semantic analysis to the entries in the National Library of Medicine's Unified Medical Language System for the classification and retrieval of patient diagnoses. Natural language diagnosis entries and inquiries were entered into the resultant semantic space, and the system provided output of the top five most similar entries from the semantic space. Chute concluded that the results suggest that classification and retrieval from large databases can benefit from coding entries and queries within a statistically derived semantic space.

Soto (1999) showed that latent semantic analysis can be used to assist in selection of menu labels for applications with which most users will be relatively unfamiliar. He used latent semantic analysis to compute the semantic similarity between task descriptions and labels in the menu system for an application. Participants performed 12 tasks by exploration, in which they engaged in search through the interface for labels that would lead them to the task solution. Performance of the tasks was faster when the labels were similar to the task description than when they were not. Soto concluded that latent semantic analysis can eventually be used along with other techniques to test the usability of computer applications automatically.

Adaptive Search and Agent-Based Technologies

Adaptive search uses software agents to profile users' past behaviors, and uses these data to predict users' goals. This helps users to achieve their goals (e.g., narrowing down the search hits; Choi & Yoo, 1999; Glance, 2001). Agent-based search tends to work well if users are trying to find something similar to what they searched for previously, but it is not very effective if users are looking for something different or new. Agents also work well when users are searching for related items within the same session (e.g., the user searches for roller skates and then helmets). However, because there is not enough data regarding what users are looking for when they are searching, agents may perform some action for the user that was not warranted or desired. Furthermore, agents' actions are based on the information the agents are logging, and if the wrong information is collected, then the agent impedes rather than facilitates the user's search.

Agent-based technologies can also personalize Web searches. Based on the individual differences between users, it is virtually impossible to design a uniform search engine that is suitable for all users. It seems that a good solution is to personalize the search engine based on users' characteristics. The problem with current software agents is that the information collected for profiling a user might not capture the user's characteristics very well. Therefore, the recommendations made by agents may be misleading. To address the issue of what information must be collected, more studies on human search behavior need to be undertaken in the future.

An example that suggests that this approach works is Amazon.com. On this Web site, the system uses the user's wish list, previous purchases, and items at which the user looked to build the user profile. The site then recommends other books or items that might interest the user.

Indexing Search

Use of meta-tag tools. Most search engines index the words in a Web page and use the first few words as a description. Meta tags allow designers to control how the page is indexed by using a meta tag to specify additional keywords to index and to provide a short description that can be displayed with each search hit. Meta tags allow designers to specify additional key words or synonyms that describe the contents of a site, but will not display on the user's Web page. Not all search engines support meta tags (Sullivan & Sherman, 2001), and it is very labor intensive to create helpful meta tags.

Listing Web sites in order of relevance. Returned Web pages are usually ranked based on some relevance measures designed for the search engine, although the ranking may be inaccurate due to the difficulty of defining relevance. Relevance is a fundamental, though not completely understood, concept for documentation, information science, and information retrieval. Studies on relevance fall into the following categories (Mizzaro, 1997): foundations, kinds, surrogates, criteria, dynamics, expression, and subjectiveness. Besides topicality, some other factors also affect a user's relevance judgment. Taylor (1986) proposed the following six criteria adopted by users in expressing dichotomous relevance judgments: ease of use, noise reduction, quality, adaptability, time saving, and cost saving. There have been other types of suggested relevance categories (see Fang & Xu, 2001), but most of these studies are exploratory and preliminary, and there is no consensus concerning the criteria adopted by users.

Improved Search Engines

Database search engines. Database search engines attempt to provide accurate, high-quality hits by having the search engine process an indexed database for a Web site. For example, the database search engine used by Fact City is a software called FactLink (http://www.factcity.com/services.shtml). FactLink manages data by organizing information in the database into data modules that consist of the original information and added information. FactLink also organizes the database into content modules that consist of metadata. The metadata help describe the content of the databases to FactLink's search engine, which uses an XML-based interface to search and display information. The search engine accepts free-form queries from users and returns results that are integrated from multiple content modules in the form of tables, sentences, and so forth. The benefits of Web site database search engines are that when an item is added to a site, a description of the product (including other names it is referred to as and related product information) can be added quickly and easily into a database. Web site database search engines are also commercially available (EasyAsk, Fact City, iPhrase, Mercado, and Requisite Technology).

Meta-search engines. Meta-search engines submit the user's keywords to several search engines and their databases of Web pages (InFoPeople Project, 2002). For example, MetaCrawler searches the databases of Alta Vista, Excite, Lycos, LookSmart, About.com, and so on. Ixquick searches Alta Vista, AllTheWeb, Excite, Hotbot, MSN, Yahoo, and so on. Users have to download Copernic Basic, but it searches multiple search engines and Web sites: AltaVista, Direct Hit, Excite, FAST Search, FindWhat, GoTo.com, HotBot, Lycos, Mamma.com, MSN Web Search, Netscape Netcenter, Open Directory Project, and Yahoo. Meta-search engines make it quicker and easier for users to utilize multiple search engines.

OminiSearch. The OminiSearch system utilizes a method that incorporates new dynamic sources into an integration system in which useful data objects are extracted from dynamic Web pages without programmer intervention (Buttler et al., 2001). It maintains the capability to extract relevant information even when faced with the evolving structure of dynamically changing Web sites. The OminiSearch architecture consists of four main components: a search execution engine; the Omini object extraction engine; an autonomous Web crawler to discover and categorize new sites; and a context catalog to store discovered Web sites, their relevant contexts, and associated information extraction rules. After entering a search request, users choose from a list of contexts for their search (e.g., general Web search, text search, article search, etc.). As described by Buttler et al.,

> The search execution engine then requests a list of Web sites that match the chosen context for the search from the context catalog. Then, for each Web site it constructs the appropriate URL based on the user search request and the Web site search interface description. The URL's are passed to the Omini object extraction engine, which retrieves the Web pages and returns extracted objects to the search execution engine for formatting. Results from each Web site are grouped together and sent back to the user. If a page format or query interface change was detected and the object extraction process failed, new rules are automatically discovered and an update is sent to the context catalog. (p. 1)

The autonomous Web crawler component in this method constantly searches the Web looking for new sources that have searchable interfaces and are related to the contexts stored in the catalog. The core of the OminiSearch system is its object extraction engine, which converts the extracted and normalized Web page into a tag tree, allowing only the relevant parts of the document to be extracted.

Content integration for e-business. Stonebraker and Hellerstein (2001) from Cohera Corporation indicated that content integration can be used to help search engines function more effectively. According to Stonebraker and Hellerstein, content integration deals with cross-business integration of information. For example, a distributor must integrate the items from individual catalogs provided by each supplier, and a travel company wanting to sell airline tickets must keep track of which seats are available for the different airlines and flights. Information that needs to be searched can be categorized as constantly changing or slowly changing. Product availability and price for many e-businesses are often considered volatile or constantly changing. For these types of information, search based on demand is needed in which the information is updated and integrated from all sources involved. Other information such as hotel addresses and airport facilities can be searched in advance and stored so that resources can be focused more to retrieving changing information that needs to be retrieved instantaneously. Stonebraker and Hellerstein also indicated that search engines should support search on all fields of the available information. In addition, given the different types of users that need to access the information, thesauruses and fuzzy searches (approximate matching of queries) should be incorporated into the system. Stonebraker and Hellerstein's company has created the Cohera Content Integration System that provides information retrieval, taxonomies, and automatic classification capabilities to help solve some of the problems of content integration.

Powerful search engines. Very powerful search engines can make it easy for users to get good search results. The best example of this kind of search engine is Google.com. The Google search engine uses a proprietary process to rank pages by how likely the pages are to contain the desired search information. Users do not have to worry about capitalization or use operators (such as "and" or "or"). As a result, for very little user effort, Google produces very helpful search results.

Search Tools

Thesauruses, dictionaries, and alternate spellings. Thesauruses, dictionaries, and alternate spellings of words should all be incorporated into a search engine. Different users do not always think of the same item in the same manner, especially when they use only one or two words to describe the item. Thesauruses can help solve the problem of ambiguity and allow users to search more efficiently. Users do not always spell keywords correctly, and the inclusion of dictionaries will reduce the number of "no matches found" error messages. Users from different countries spell words differently (e.g., color vs. colour). Including mechanisms to detect and match alternative spellings of words increases the ability of a large population of users to quickly and accurately access information. For example, Clear Ink's SpellWeb provides a solution that allows users to enter two queries, and the search engine decides which one is correct or more popular (http://www.spellweb.com/).

Providing limited search options. To eliminate the problem of users entering queries that do not yield desired results, designers can choose to present users with limited search options. With this method, when users utilize the search function, a drop-down menu appears offering them a limited number of categories from which to choose. For example, if the Web site was designed to sell flowers, the drop-down menu may consist of the following categories: specific flowers, holidays, special deals, and so on. Designers can also use multiple limited search option menus. After the user selects "specific flowers," another drop-down menu may appear with the following categories: lilies, roses, orchids, and so forth. After the user makes the selection "lilies," another drop-down menu may appear with the following categories: Casablanca lilies, stargazer lilies, yellow lilies, and so on. Although this process narrows the search and ensures that the users' queries are always valid, this can be a tedious process, and users may not want to go through the entire series of steps. Furthermore, this method seems to work better if users know specifically for what they are looking (Denning, 1992).

Ask users the goals for search. One way to help users find what they are looking for is to ask them explicitly to indicate the goals and desired outcomes of their search. This method is currently used by the "Ask Jeeves" search site. Ask Jeeves allows users to enter a question and attempts to provide answers to the questions with the goal of conducting an interaction based on understanding users' goals. For example, if a user is planning to visit the state of California and is interested in places to visit, the question that could be submitted is, "Where is the best place to go in California?" Once this query is submitted, the system presents different fixed ways to phrase the question, as well as a list of sites that might be relevant.

Designing for Effective Problem Solution

Keeping smart, experienced people in the loop. Today, the most reliable solution for engineering adaptive systems that are capable of responding to situations that were not anticipated in the design phase is to include "smart" human operators within the control loop. For example, nuclear power plants operators, air traffic control operators, or airline pilots are often required to "complete the design" by creatively responding to situations that were not anticipated by system designers. The ability to do this often reflects implicit knowledge of the work domain. The type of knowledge that these experts bring to the system is similar to "survey" knowledge in the context of wayfinding (Rossano, West, Robertson, Wayne, & Chase, 1999). In contrast, normative procedures or the rules implicit in many automated systems reflect a kind of "route" knowledge. The survey knowledge of experts allows them to adapt (i.e., find an alternative solution or path to the goal state) when the route (implicit in normal operating procedures or rule-based solutions) is blocked by events that were not anticipated in the system design. The survey knowledge of experts should be incorporated into procedures or rules to assist the user when problems arise. One goal of knowledge elicitation is to tap into the implicit survey knowledge of experts so that this knowledge can be reflected in interfaces and training protocols. For example, well-designed graphical interfaces can sometimes be designed to provide survey knowledge about a work domain. In this sense, they function as a map that shows the deep structure of the domain. Such maps may allow operators to respond adaptively when normal procedures do not satisfy the goals.

Cognitive systems engineering. Cognitive systems engineering has emerged as a new approach to human–machine systems that considers work as a system that has emergent properties that are not easily appreciated from within the traditional disciplines (e.g., Flach & Rasmussen, 2000; Vicente, 1999). It attempts to discover the deep structure within work domains. To contrast cognitive systems engineering with classical approaches to human factors, the term *use-centered* has been used (as opposed to *user*-centered; Flach & Dominguez, 1995). Cognitive systems engineering attempts to discover the functional constraints on work through its analysis, and then to make these constraints explicit to human operators through the design of interfaces and training protocols. It views technology as tools that mediate work. Therefore, for example, the focus of cognitive systems engineering is on human–work interaction, rather than on HCI. The computer and other technologies are evaluated in terms of their ability to provide direct perception and direct manipulation of work-related "objects."

The BookHouse interface example. The BookHouse (Pejtersen, 1989; Rasmussen et al., 1994) is a good example where the constraints underlying the functional task are not well understood. In this case, the goal was to help people find interesting books in a public library. To solve this problem, the interface was organized to reflect strategies that librarians use to assist patrons. Different areas within a virtual room were designed to support different search strategies (e.g., search by genre, search by author or title, or browse). The graphical interface allows patrons to implement the expert strategies by manipulating graphical objects. In this case, the patrons' spatial manipulation skills were translated into intelligent search strategies. An assumption of this approach was that the strategies used by expert librarians were responsive to the deep structure reflected in the organization of libraries and in the goal to find something interesting.

4.3. Summary and Future Research

Current search engines are for the most part inadequate and inconsistent with each other. Furthermore, lack of knowledge about how users search for information results in users' perspectives not being incorporated into the search engine design. In general, user testing of all types with respect to search seems to be deficient (e.g., Hagen, Manning, & Paul, 2000; Nielsen & Tahir, 2001). This deficiency needs to be remedied.

The problems associated with search engines have resulted in a tremendous amount of research looking at specific methods to remedy the problems. The statistical techniques for analyzing natural language have considerable potential for improving search by taking meaning and intentions into account. Agent-based technologies can personalize search based on previous search behaviors. Indexing tools can help organize and structure Web pages, and improved search engines can improve the hit rate. For example, Google.com shows that a powerful search engine can provide a simple user interface that gets helpful hits on the first page of the search results. The search engine should be designed to work hard to figure out what the user means and to prioritize the search results. Search tools can also help reduce irrelevant hits and retrieve sites that are related to the users' queries. The methods discussed with respect to designing for effective problem solution suggest that the search process should take into account the users' goals.

The methods that seem to have the most promise for returning relevant hits for each of the major categories are

- Latent semantic analysis.
- Agent-based technologies.
- Meta tags.
- Development of powerful search engines.
- Thesauruses.
- A cognitive systems engineering perspective.

With regard to users, future research should focus on observing, evaluating, and comparing novice and expert search behavior for given domains. Currently, there is simply not enough research on analyzing what users do when they search for information and whether these search techniques and strategies are generalizable across different domains. More research should be conducted to find out what strategies users adopt to get a hit with the top five queries when they have a definite search goal. Furthermore, studies should also be conducted to determine the role of experience in facilitating search techniques. With experience, users should learn to become better searchers, and determining what can be done to facilitate learning to search will be beneficial.

One research problem with browsing is to determine how to design the information hierarchy. Numerous studies have been undertaken to address the hierarchy design of computer menu systems (Jacko & Salvendy, 1996; Jacko, Salvendy, & Koubek, 1995), but these findings were based on the assumption that the response time of each step is constant and users have some basic training before they use the system. Because this assumption does not hold for Web applications, the findings about computer menu systems may not be applicable to Web applications. Another research question that needs to be understood is when and how to apply browsing and keyword searching in a Web application.

5. PRESENTATION OF INFORMATION

Even if software engineers develop search engines that elicit the appropriate information with a logical and manageable structure, effective and efficient retrieval of this information requires that the information be displayed in a manner that is compatible with the user's goals and models. The question of interest, then, is what is the best way to present information? Web design generally begins with an initial definition of the design and evolves into a detailed design, from which iterative cycles of evaluation and improvement transpire (Martel, 1998). For example, when clients approach designers to develop a Web site for their business, the designer often analyzes the wants and needs of the clients. Often designers will look to existing products within their own product lines and competitors' products (Stanney et al., 2001). The information obtained can provide ideas to fuel truly innovative Web designs that uniquely meet clients' needs. The designer then develops an objective for the site that drives the design process.

Because clients often demand that the Web site be constructed in a short amount of time, the designer must define the target users in terms of their objectives, work environments, and lowest common technology used (Najjar, 2001). This information will allow designers to develop prioritized functional requirements that meet the users needs. The designer then begins to structure the site, organize its functions, and design the user interface. Iterative user testing is then performed to ensure the effectiveness of the site.

5.1. What Are the Current Problems That Prevent Effective Information Display?

Lack of Attention to Usability Issues

Despite the demonstrated cost effectiveness of a user-centered approach, most user interfaces are designed without user input. Therefore, the information is organized and presented in a way that is intuitive for designers, but not for end users (e.g., Tullis, 1993). There are many reasons for this inattention to usability in design:

- Presently, there is an inadequate number of interface designers with expertise in usability. There are few graduate programs in applied user interface design.
- There is no emphasis on performing usability evaluations early and throughout the design process. Usability evaluation is often not incorporated until the late design stages, making changes either more costly or impossible. According to C.-M. Karat (1990), for each dollar spent to fix a problem during early design stages, $10 must be spent to fix the same problem in later stages.
- A common misconception in the business community is that time and financial costs involved in usability evaluations are high, and the return on investment is questionable. Yet, IBM's Ease-of-Use Web site states, "The rule of thumb: Every dollar invested in ease of use returns $10 to $100" (I.B.M., n.d.). C.-M.. Karat (1997) showed that an early focus on usability evaluations actually advanced a product's release date. Other benefits of usability evaluations include lower training costs, increases in productivity, reduction in customer support costs, and increases in customer satisfaction (Bias & Mayhew, 1994; Mayhew, 1999). Obviously, these benefits are not being made known to those responsible

for allocating resources to usability evaluations. User interface design professionals must spend more time educating the business community on these benefits.

- Usability testing is sometimes conducted only on a small group of "average" users rather than on a range of diverse users (Stephanidis & Emiliani, 1999; Weber 1999). Therefore, it is difficult to generalize the test results to the end users. Although this may be appropriate for products aimed at a specific target group, Web applications must be usable to users with a wide range of backgrounds and abilities.
- There is reliance on guidelines (which are not always followed), rather than on usability testing.

Not Enough Time to Iterate the User Interface Design

When designers get feedback from users, they can improve the way information is presented on Web pages. Designers can use a series of iterative design–feedback–redesign steps to create pages that are easy to use and understand. Unfortunately, many Web projects do not allocate enough time for these iterative design cycles.

Lack of Standardized Guidelines

As mentioned earlier, guidelines provide useful information regarding how information should be presented. Although some general guidelines (e.g., Nielsen, 1997a) are well known to user-interface designers, many guidelines for specific user tasks (e.g., data displays, Smith & Mosier, 1986; e-commerce user interface design, Najjar, 2001) are hidden in conference proceedings and other specialized publications. Moreover, guidelines are sometimes contradictory. Guidelines are not always followed (e.g., putting too much information on a page, poor cueing for important areas, etc.) when they should be and are sometimes enforced in special cases where they should not be (e.g., horizontal scrolling should be allowed when information needs to be compared, Najjar, 2001). Although adherence to design guidelines can improve user interactions with information displays, the guidelines need to be organized around specific tasks and displays for maximum benefits. Finally, guidelines often focus on the visual presentation of information, whereas new interaction platforms and devices increasingly make use of other or additional modalities (e.g., speech output, voice input). For these additional modalities, no comparable sets of established and accepted guidelines exist.

Designing for Multiple Cultures

One problem for global e-commerce is that the Web site must be designed in a manner that can be translated into various languages. The varying protocols of multiple cultures and languages pose a difficult design challenge for any company. Furthermore, because the users are heterogeneous, a site's information should not lose its meaning through translation, and the translated version should incorporate cultural differences. For example, white is the color of mourning in Japan. Therefore, to avoid offending Japanese users, do not use white text.

Inaccessibility

Among the ingredients of success of the emerging information society, accessibility is considered to be of paramount importance. The issue of accessibility concerns the right of all users to be "connected" (i.e., to obtain and maintain access) to a society-wide pool of information resources and interpersonal communication facilities, given the varieties of context (Stephanidis & Emiliani, 1999). Designing for disabled populations often produces benefits to people without disabilities (Vanderheiden,1997), yet very few Web sites currently comply with guidelines for allowing disabled persons to use Web sites. In the United States, the Americans with Disabilities Act (U.S. Department of Justice, 1990; 28 C.F.R. Sec. 36.303; 28 C.F.R. Sec. 35.160) requires that Web pages be "accessible to people with disabilities, unless doing so would result in a fundamental alteration to the program or service or in an undue burden."

Currently, a topic under discussion in Europe is the participation of all people in the knowledge society (eEurope, 2002 initiative of the European Commission, 2000). In its e-accessibility action line the initiative focuses mainly on the provision of access to the Web. If Web pages are properly designed (e.g., by following the W3C–WAI guidelines of the World Wide Web Consortium, 1997–2002), they are accessible by users with diverse sensory abilities (e.g., blind or visually impaired users), but also through different access devices and software applications (e.g., nonvisual browsers like the AVANTI Web browser; see Stephanidis, Paramythis, Akoumianakis, & Sfyrakis, 1998).

Lack of Research on How Information Should be Presented to Facilitate Web Searching

Web searching is an interactive process. The user enters a search query, reviews the returned search results, and then may refine the initial search query. During this process, what and how to present the search results plays an important role because the information not only potentially gives the user what he or she wants, but also provides clues that could be used to refine the search. The information displayed in the results page is not very helpful (Haskin, 1997; Hearst, 1997). Few studies have been undertaken on these issues.

5.2. What Are the Well-Known, Existing Methods Being Used to Remedy These Problems?

Generally, information displays are evaluated with usability testing. Usability testing is perceived by many as time consuming and expensive, despite evidence to the contrary (e.g., C.-M. Karat, 1994; C.-M. Karat, 1997; Nielsen & Coyne, 2001; for a review, see Bias & Mayhew, 1994). Another method for evaluating information displays uses design guidelines. Designers are encouraged to use them in the design process, and usability engineers often evaluate designs for guideline violations. The most cited Web user interface design guidelines are, "top ten mistakes of Web design," "top 10 new mistakes of Web design," and "ten good deeds in Web design," proposed by Nielsen (1997a, 1999b, 1999c, respectively). Although following these guidelines can improve the usability of Web sites, complying with the guidelines does not eliminate the need for usability evaluations. The following sections discuss some of the current methods used to present information effectively.

User Testing

As indicated earlier, it is important to involve users in the evaluation of a product to promote effective and efficient user interactions with the product. User-testing methods assess the adequacy of systems in terms of their compatibility with the characteristics of the potential users and their effectiveness in helping users achieve their goals. Some companies have their own usability labs. It is important that companies hire qualified usability experts to design the interface to convey information effectively. Furthermore, user testing is normally conducted using prototypes to identify problems during product design. A prototype is a model of the complete system or interface that looks and works like the proposed system, but is usually incomplete and easier to change (e.g., Najjar, 2000). Prototypes are easily altered to accommodate design changes identified during usability testing for testing alternatives. Prototypes may be used in all the types of usability tests described later. The earlier user needs and feedback are used in the design process, the more quickly and cheaply design improvements can be made. The following is a description of the most common methods of user testing for assessing information presentation issues. The first six are only described briefly because they are minor variations of methods previously described for use in knowledge elicitation.

Naturalistic and field observations. One method for evaluating how users interact with a product is to observe users unobtrusively in their environments (see Proctor & Van Zandt, 1994). This method focuses on real-world interactions with the system and usually requires usability engineers to visit the site they are studying to collect information. This technique is often used early in the design cycle, when information is gathered about the context regarding the existing products. One advantage of this approach is that evaluators can determine whether users actually use the system (Eason, 1984; Nielsen, 1993).

Interviews. Interviews are formal methods of gathering information. As described in the knowledge elicitation section, there are two types of interviews: structured and unstructured. With structured interviews, users answer a set of predetermined questions; and with unstructured interviews, users provide answers or descriptions about a task or interface. The type of interview to be conducted depends on the goals of the interviewer.

Contextual inquiries. Contextual inquiry is an exploratory interview method. It is aimed at providing the usability engineer with background information on the use of a product. It helps develop a context regarding how the product is being or will be used, how users interact with the product, and so on. Contextual inquiries combine the methods of interview and observation. Contextual inquiries focus on understanding users' work habits and their interactions with the system.

Focus groups. Focus groups usually consist of a group of users that are brought together to discuss different issues regarding the features of a system. The group interaction allows for the exchange of ideas as well as feedback from the group. The group is usually moderated by a usability engineer who keeps the discussion on track. If users involved in the focus group have not actually used the system for a period of time, they must rely on the demonstrations provided by the moderator. Because watching the moderator or another user perform or interact with the system is not the same as actually interacting with the system, the feedback provided by the focus group may not identify all the usability issues that would have been identified if the user explored the system.

Questionnaires. Questionnaires are often used to gather information about user preferences for a system (e.g., Chin, Diehl, & Norman, 1987). Questionnaires can be sent quickly to a wide range of users and have the potential for gathering a significant amount of meaningful data. However, the return rate for questionnaires is often low and what users say they do and what they actually do often differ (e.g., Nielsen, 2001a).

Journals and logs. Journals and logs are notes kept by users that log their activities and thoughts. The details of the logs are dependent on the conscientiousness of the user. Users may not log all their activities or thoughts, and usability engineers might not understand parts of users' reports. This method is also used in the early stages of development to obtain information on what and why users perform certain tasks.

Heuristic analysis. With heuristic analysis, several usability evaluators interact and examine the interface to judge its compliance with a set of established usability rules of thumb (e.g., Nielsen, 2001b). Nielsen (1993) showed that three to five evaluators working independently will find 80% of the usability problems. A group of independent evaluators is preferred because different evaluators will focus on different problems, increasing the likelihood that the collective list of problems will be more complete. The evaluators also assign usability severity ratings to each problem. This provides designers with a hierarchy of problems and their severity. Heuristic analyses are conducted in early design phases because they provide quick methods to find major usability problems. They are also used to evaluate competing sites.

Walkthroughs. Walkthroughs are system evaluation methods whereby designers interact with the system from the user's point of view to find usability problems. There are two major types of walkthroughs: cognitive and pluralistic. A cognitive walkthrough is a method in which a usability engineer constructs scenarios of users interacting with the system. The usability engineer then works through the scenarios while assuming the role of the user. The engineer notes problems that prevent users from completing their tasks. A pluralistic walkthrough is a method that is conducted by one or more usability engineers, system designers, and users. The users interact with the system while the usability engineer records problems with the product. The system designer's role is to answer questions regarding the system (Wharton, Rieman, Lewis, & Polson, 1994).

Usability testing. Usability testing is a technique for assessing the usability of a prototype or incomplete version of a Web site. In a usability test, a representative user is given a scenario and performs a set of tasks on the Web site. The tasks are designed to elicit problems with the interface. The session is usually videotaped. User performance measures (e.g., such as time on task and number and severity of errors) are recorded and defined as indicators of usability (Nielsen, 1993; Scerbo, 1995). The user also completes questionnaires on the usability of each task and the usability of the Web site as a whole. The users are run separately, and each testing session is mediated by a usability engineer who either stays in the room with the user or moves to an observation room where the user is observed through a one-way mirror or video screen. While the user is interacting with the interface, the user may be asked to think aloud. The user is instructed to verbalize his or her thought processes, feelings, preferences, and problems while performing the task. The usability engineer mediates the session by asking questions relevant to the task. Sometimes a method of joint discovery is used where two users interact with the system. This allows examination of situations in which users are working with other people to complete a task and provides a more natural method of thinking aloud. Early testing of prototypes and mockups can find potential problems when the cost and time to make fixes is lower than completed Web sites (Mayhew, 1999; Nielsen, 1993; Scerbo, 1995).

Ensure That Critical Information Items That Support the Web Design Process Have Been Addressed

Stanney et al. (2001) suggested that the following activities be completed to ensure that the critical information items that support the Web design process have been addressed:

- Identify necessary goals, functions, objectives, desires, behavioral patterns, and usage contexts to be met by Web design.

- Become familiar with users' roles (both situated and formal), practices, tools, equipment, organizational factors, and vernacular of Web environment.
- Characterize user profiles in terms of psychological characteristics, underlying cognitive processes and knowledge structures, experience, job and task characteristics, information needs, and physical attributes.
- Develop information models via knowledge elicitation (e.g., interviews, protocol analysis, questionnaires, concept sorting, focus groups, critical incident reports, abduction, narrative, observation and ethnographic study, expert systems, expert diary, log files, semantic analysis, wants and needs analysis, and group task analysis) that are relevant, accurate, adaptable to changes in user and group behavior, and generalizable.
- Acknowledge individual and group differences within the target user population.
- Develop an information architecture that reflects formal and informal practices, goals, activities, and organizational structures and has highly usable navigation structure.
- Conduct a competitive analysis and identify means of organizing and structuring information in a way that promotes efficient and effective retrieval, yielding an advantage over competitors.
- Set usability goals that promote efficient and effective information retrieval.

Ecological Interface Design

Ecological interface design, described earlier, is based on the idea that the interface should explicitly display the constraints of a work domain to the user. Vicente (1999) stated, "According to EID [ecological interface design], to properly support such knowledge-based behavior, an interface should display the physical and functional properties of the work domain in the form of a multilevel representation based on the abstraction hierarchy" (p. 63). For interface design, the user's behavior mode should be kept as low as possible, because response times will be shorter and fewer errors will be made.

Jungk, Thull, Hoeft, and Rau (1999) designed an interface to support the decision making of anesthetists based on ecological interface design. An anesthetist's decision-making process is based on a concept of partly abstract physiological parameters (e.g., depth of anesthesia) derived by combining parameter relations and additional context information. However, current displays typically show vital parameters as trends along a timeline and do not visualize the concept. Jungk et al. developed a new interface based on the ecological principle of displaying all information necessary for decision making in a single display. Their display visualized 35 relevant parameters for anesthesia monitoring according to their function (e.g., temperature, respiratory mechanics and gases, etc.). Their studies suggest that the ecological interface may provide better support for the anesthetists' concept of decision making than do traditional trend displays.

Self-Adaptation

In HCI, self-adaptation of a system's interactive behavior has been proposed as a framework for providing accessibility and high interaction quality to all potential users, on the basis of each user's individual abilities, requirements and preferences, as well as of the context in which interaction takes place and of the adopted interaction technology (Stephanidis, 2001). For more than a decade, adaptation has been discussed as an approach for coping with the problems of diversity in the HCI field (Browne, Totterdell, & Norman, 1990; Brusilovsky, 1996; Stephanidis, 2001).

Adaptation implies the capability on the part of the system to capture and represent knowledge concerning alternative instantiations suitable for different users, contexts, purposes, and so on, as well as of reasoning about those alternatives to arrive at adaptation decisions (Savidis, Akoumianakis, & Stephanidis, 2001). Furthermore, adaptation implies the capability to assemble, coherently present, and manage at run time the appropriate alternatives for the current users, purposes of use, and context (Savidis & Stephanidis, 2001). Adaptation is a multifaceted process that can be analyzed along three main axes, namely the source of adaptation knowledge, the level of interaction at which it is applied, and the type of information on which it is based (Stephanidis, 2001; see Figure 4).

Regarding the source of adaptation knowledge, one can identify two complementary classes: knowledge available at start-up, that is, prior to the initiation of interaction (e.g., user profile, platform profile, usage context); and knowledge derived at run time (e.g., through interaction monitoring, inspection of the computing environment). Adaptation behavior based on the former type of knowledge is termed adaptability and reflects the capability on the part of the interface to tailor itself automatically to the initial interaction requirements, as these are shaped by the information available to the interface. Adaptation behavior based on the latter type of knowledge is termed *adaptivity* and refers to the capability on the part of the interface to dynamically derive further knowledge about the user, the usage con-

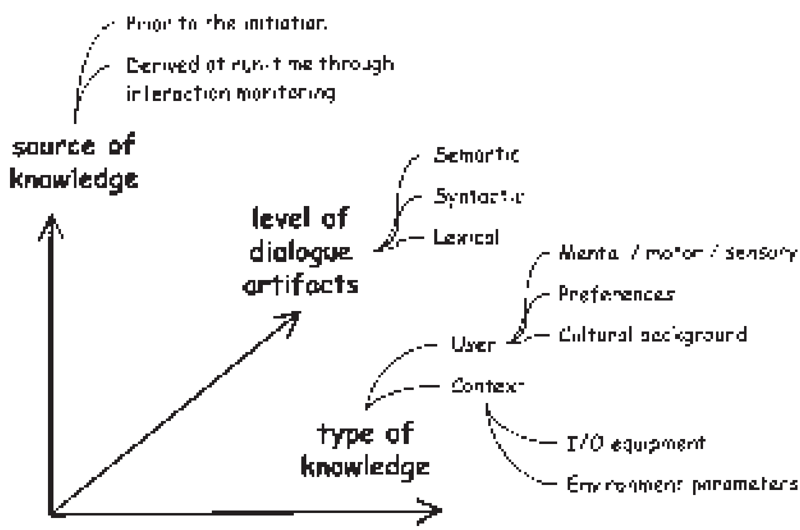

FIGURE 4 Adaptation dimensions in self-adapting interfaces.

text, and so forth, and use that knowledge to further modify itself to better suit the revised interaction requirements.

The second axis of analysis of adaptation concerns the level of interaction at which adaptations are applied. In particular, it is possible to design and effect adaptations at all three levels of interaction:

1. The semantic level (e.g., by employing different metaphors to convey the functionality and facilities of the underlying system).
2. The syntactic level (e.g., by deactivating alternative dialogue patterns, such as, for example, object function vs. function object interaction sequencing).
3. The lexical level (e.g., grouping and spatial arrangement of interactive elements, modification of presentation attributes, alternative input–output devices).

The third main axis of analysis concerns the type of information being considered when deciding on adaptations. Exemplary categories of information that can be employed include the following: design constraints, as these are defined by user characteristics (e.g., abilities, skills, requirements, preferences, expertise, cultural background); platform characteristics (e.g., terminal capabilities, input–output devices); task requirements (e.g., urgency, criticality, error proneness, sequencing); and so on. Furthermore, information that can only be acquired during interaction can equally participate in the decision process (e.g., identifying the user's inability to successfully complete a task, inferring the user's goal or intention, detecting modifications to the run-time environment).

Conceptual Design Methods

Conceptual design methods can determine the type of information users expect with a certain interface. This information is organized into sequences or hierarchies that can be used to for displaying information in a manner that is consistent with the users' mental models of the system.

Statistical Ranking Methods

Statistical ranking methods can be used to rank information from servers and databases according to the probability that they contain relevant information for a given user query (e.g., see the sections on latent semantic analysis).

Cognitive–User Modeling

The term *user model* has many different meanings (e.g., user's mental model of a system). Here, models of the user's knowledge of the system and task are referred to, which are also known as cognitive models. Booth (1989) cited the following advantages of cognitive–user models:

- The modeling process itself can help in matching system features to the needs of the user.
- User models can suggest metaphors (e.g., desktop) to improve user learning and understanding.
- User modeling can guide design decisions and choices by making the assumptions about the users explicit.
- User models can provide a method for evaluating alternative design solutions. Ideally, quantitative comparisons (e.g., learning or performance times; cognitive complexity) can be made between proposed designs.

Developing user models is essentially a process of task analysis. Unlike most task analysis methods, however, cognitive modeling provides a structured format (usually based on principles of human cognition) for the description of user performance. The most widely known family of cognitive modeling techniques is the goals, operators, methods, selection (GOMS) model first described by Card, Moran, and Newell (1983). A GOMS model is a description of the user's task goals, methods needed to accomplish specific goals and selection rules for applying methods when multiple methods exist. Methods are sets of (often observable) operators that the user performs.

Since Card et al.'s (1983) initial GOMS model of simple text editing, several variations on the basic GOMS model have been developed (e.g., CPM–GOMS, NGOMSL, John & Kieras, 1996b). Successful applications of the GOMS modeling range from an analysis of a telephone operator's directory assistance workstation to a nuclear power plant operator's associate (for a review, see John & Kieras, 1996a).

Cognitive modeling is often criticized for the amount of time required to develop models. However, more practical approaches have been developed (e.g., Kieras, 1997). Moreover, John and Kieras (1996a) pointed out that this perception of the time required for GOMS modeling is based on the early development of complete research-based GOMS models. In practice, modeling can be performed on a critical subset of the user tasks, significantly reducing development time.

The Field of Semiotics

Semiotics is the study of signs (Chandler, 2001). It includes not only signs of the type usually of concern in human factors (i.e., visual displays of information) but also photographs, drawings, words, and body language. In short, a sign is anything that stands for something else. Contemporary semiotics is concerned with sign systems and how meanings are made, conveyed, and maintained by these systems. Semiotics encompasses syntax (the structural relations among signs), semantics (the relation of signs to what they signify), and pragmatics (the relation of signs to the persons interpreting them). The general idea is that people make meanings through the creation and interpretation of signs.

Lytje (2001) took a cognitive linguistic perspective to user interface design. He suggested a semiotic analysis of the software interface, viewing the interface as an interplay between "1) the structural characteristics of interface signs, 2) the meaning of the signs, and 3) the implementation of the signs" (p. 229). At what Lytje called the *expression plane,* the signs of the user interface are symbolic representations that enable the user to interact with different kinds of information. At the content level, cognitive schemas are activated when interacting with a software product, allowing interpretation of the sign. The general idea is that the contextual background must be made explicit so that the appropriate schemas are activated to allow the user to understand the meaning of the sign.

Standardizing User Interfaces

One of the reasons why information is not presented in an effective manner is that Web user interfaces are so different from each other that users cannot create a mental model that can be used in all situations. For example, the navigation controls for Web sites are often in different locations. By producing a set of generalizable guidelines for a given domain, users will be able to transfer their knowledge and expectations from one user interface to another with little interference. Although there has been an emphasis in recent years to allow user customization based on individual needs, establishing basic standard guidelines will still be beneficial (e.g., Nielsen 1999a).

Framing the Problem as One of "Genre"

Most hypertext research has sought to solve the problems of users navigating information space by either (a) providing better navigation tools, or (b) providing better labels and indexing—one is a spatial help, the other is a semantic help. For working with digital information spaces spatial and semantic cues are needed, at the same time. To do this, one needs to understand both the surface cues and the deep structure of different genres of Web sites, Web applications, and whatever comes next. A genre, by definition, is a set of conventions, but is more than the surface conventions; it is also the deep structure (Vaughan, 2000).

Taking Different Users Into Account When Presenting Information

Although the Web is available to anyone, it is unlikely that organizations intend to conduct business with everyone. Companies routinely select target markets for their products as part of strategic planning. It is no different when the selling outlet is the Web instead of a brick and mortar store. Unlike brick and mortar environments, however, companies selling online may collect, store, analyze, and use data on individuals that purchase or even just visit a site. Although it is possible that demographic or psychometric characteristics impact site preferences, little is known about that impact or its extent.

Goodhue (1995) proposed that users of information systems possess characteristics that determine their propensity to use information systems and their skill at using them. These characteristics include those within the individual, such as cognitive skills, as well as social norms and environmental characteristics. Characteristics of consumers have been identified and investigated to the point that a *Handbook of Marketing Scales* exists for guidance (Beardon & Netemeyer, 1999). Commerce on the Web incorporates aspects of both these environments, using information systems and marketing, because use of the Web for shopping is discretionary. Preliminary studies have identified areas of importance to online consumers in general. They are transaction security, information quality, and information quantity (Lightner, 2001). In addition, the presence of a combination of both text and pictures for product descriptions is preferred, with textual descriptions alone preferred over pictures alone (Lightner & Jackson, 2001). Evidence suggests that differences in preferences exist between demographic groups as well, with younger respondents valuing sensory impact more than older respondents (Lightner, 2001).

Discussions concerning consumer groups include lifestyle factors and level of comfort with technology as possibly determining interface preferences. Difficulties arise when determining user preferences for a technology that is relatively new. As consumers gain experience, their preferences may change. This evokes the dilemma of reacting to a user's changing preferences or directing users to become proficient at a site that is designed with efficiency in mind. Because commercial Web sites are complex software systems that are resource intensive to develop and maintain, this dilemma has real implications for company planning.

Designing for special populations.

It is necessary to consider the special characteristics of users in the information presentation to get a better performance in Web browsing and searching tasks. The Americans With Disabilities Act specifies that, to accommodate disabled persons, every image needs to include a textual alternative. However, there is much information in complex images that cannot be expressed verbally. Consequently, preprocessing images to improve their comprehension may also be of benefit to users with disabilities.

Yang and Lehto (2001) conducted research regarding what types of images on the Web should be processed and how this processing should be done to increase the accessibility for the visually disabled users. The results showed that the use of image processing can improve image understandability for visually challenged Web users. Segmentation and edge detection algorithms in the image processing technology, which are of particular interest due to their relatively small computational requirements, were used in this study to simplify Web images. Blind participants interpreted photographic images (typical of those found on the Web), converted into tactile form, more accurately after image processing. Segmentation was somewhat more effective than edge detection. The understandability of less complicated Web images, such as symbols, was not significantly increased by either method. Furthermore, blind users performed better on Web browsing tasks after Web pages were simplified using the segmentation method.

5.3. Summary and Future Research

A key factor for displaying information successfully is to involve users testing throughout the design process. Currently, this does not seem to be the common practice. Furthermore, the user interfaces for most Web sites are not designed to allow equal access for all users. To remedy the problems associated with the way information is displayed, future research should focus on determining research tools, techniques, and methodologies that incorporate user feedback earlier in the design process (e.g., Najjar, 2000). Currently, most user testing is conducted late in the design process. However, many of the methods described previously focus on gathering user input earlier in the design phase. Detailed evaluations of the contributions of the alternative methods applicable at early stages need to be performed, and emphasis should be placed on developing and evaluating methods applicable to earlier stages, when it is quicker, easier, and cheaper to make design improvements. To facilitate the incorporation of usability into the user interface design process, research that examines ways to quickly and effectively train user interface design experts on usability will also be beneficial.

In the process of incorporating user feedback into product design, researchers need to develop research-based user interface design guidelines for the Web and promote the use of these guidelines. These guidelines can then serve as a foundation for developing a catalog of Web user interface design templates that can be customized for Web sites. Furthermore, these templates can be designed in a way that is compliant with the Americans With Disabilities Act.

Because there is an increasing tendency to allow users to customize the way information is presented or to program the system to adapt to users preferences based on their past actions, studies need to be conducted to determine the costs and benefits of these options. For example, to determine the degree of adaptability the system should perform, appropriate user modeling based on an individual's tolerance level should be conducted. Because individuals have limits regarding how much information or restrictions can be given to them, and the quality of the adaptations made by the system, users may want more or less control over the system's settings.

6. CONCLUSIONS

Content preparation and management has developed into one of the central concerns regarding efficient use of the Internet and of organizational intranets. The challenge is finding ways to prepare and manage content so that the full range of users will be able to access the information that they need quickly and easily. Content preparation is critical to the success of companies' electronic endeavors, resulting in the growth of numerous businesses devoted to providing this service.

This white paper has reviewed several areas of concern with respect to content preparation and management. The first area is that of how to elicit the knowledge and information that needs to be contained in a particular Web site or application. Many knowledge elicitation techniques were described, which can be used effectively to extract knowledge from various user groups. However, knowledge elicitation arose originally as a topic of concern in the development of expert systems and only recently has emerged as a focus for development of Web-based applications. Consequently, the most sophisticated methods are directed toward elicitation of expert knowledge rather than toward content development. The need exists for a concentrated research effort with the specific goal of developing and analyzing knowledge elicitation techniques for content preparation in various domains.

The second area considered was that of the organization and structuring of content for the Web. It was indicated that the organization and structure of information should reflect the context, content, and users, and numerous methods relevant to Web design were described. The major issue that needs to be resolved is to determine the fundamental elements of information for specific tasks. Significant strides toward addressing this issue have been made with respect to natural language processing, and research should focus on how best to implement the current methods, such as latent semantic analysis to improve natural language interfaces. Research also needs to focus on developing similar methods for determining and organizing meaning elements from images and the physical environment.

The third area investigated was improving Web search through making information retrieval more effective. Considerable knowledge is already available for ways to improve the effectiveness of search. After evaluating these methods, it was proposed that several methods with the most promise should be the focus of future research. Of particular concern with regard to the issue of search is a relative lack of research on what strategies are used by various users when engaging in the task.

The last area considered involved the design of the user interface to provide accessibility to all and to maximize usability. The major issues pertaining to this area involve mapping the structure and organization of information to the interface display, conveying the information in a manner that promotes successful interactions with users, and designing for all potential user groups. In general, there is inadequate consideration of users throughout the design process. Although much is known about usability testing in general, more thorough evaluations of the contributions of alternative methods and the best way to incorporate usability information in the design process are needed.

At a global level, this white paper illustrates that there is a wealth of knowledge in many disciplines that bears on issues pertaining to content preparation and management. This review has only touched on the range of methods and theories that may be of benefit to content design. It is clear that integration of this information in a way that will significantly advance the state of the art in content preparation and management will require a concerted multidisciplinary effort involving academia and industry.

REFERENCES

Allen, J. (n.d.). *Making a semantic Web.* Retrieved October 1, 2001, from http://www.netcrucible.com/semantic.html

Armitage, J. (2001). Viant. *Interactions, 8*(2), 75–79.

Bailey, R. W. (1993). Performance versus preference. In *Proceedings of the Human Factors and Ergonomics Society 37th Annual Meeting* (pp. 282–286). Santa Monica, CA: Human Factors and Ergonomics Society.

Bear, E. G., Teasley, B., & Carroll, L. P. (2001). MONKEYmedia. *Interactions, 8*(2), 63–69.

Beardon, W. O., & Netemeyer, R. G. (1999). *Handbook of marketing scales* (2nd ed.). Thousand Oaks, CA: Sage.

Bell, G. (2001). Looking across the Atlantic: Using ethnographic methods to make sense of Europe. *Intel Technology Journal, 5*(3). Retrieved January 27, 2002 from http://developer.intel.com/technology/itj/q32001/articles/art_1.htm

Ben-David, A., & Utgoff, P. (1995). Monotonicity maintenance in information-theoretic machine learning algorithms. *Machine Learning, 19,* 29–43.

Beyer, H., & Holtzblatt, K. (1998). *Contextual design: A customer centered approach to systems designs.* San Francisco: Kaufmann.

Bias, R. G., & Mayhew, D. J. (Eds.). (1994). *Cost-justifying usability.* New York: Academic.

Bookstein, A. (1985). Probability and fuzzy-set applications to information retrieval. In M. E. Williams (Ed.), *Annual review of information science and technology* (Vol. 20, pp. 117–151). White Plains, NY: Knowledge Industry Publications.

Boose, J. H. (1986). Rapid acquisition and combination of knowledge from multiple experts in the same domain. *Future Computing Systems, 1,* 191–216.

Boose, J. H., & Bradshaw, J. M. (1988). Expertise transfer and complex problems: Using AQUINAS as a knowledge-acquisition workbench for knowledge-based systems. In J. H. Boose & B. R. Gaines (Eds.), *Knowledge acquisition tools for expert systems: Knowledge-based systems* (Vol. 2, pp. 39–64). London: Academic.

Boose, J. H., & Gaines, B. R. (Eds.). (1988). *Knowledge acquisition tools for expert systems: Knowledge-based systems* (Vol. 2). London: Academic.

Booth, P. (1989). *An introduction to human–computer interaction.* Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

Borst, A., & Theunissen, F. E. (1999). Information theory and neural coding. *Nature Neuroscience, 2,* 947–957.

Browne, D., Totterdell, P., & Norman, M. (Eds.). (1990). *Adaptive user interfaces.* London: Academic.

Brusilovsky, P. (1996). Methods and techniques of adaptive hypermedia. *User Modeling and User-Adapted Interaction, 6,* 87–129.

Buie, E. (1999). HCI standards: A mixed blessing. *Interactions, 6*(2), 36–42.

Burgess, C., Livesay, K., & Lund, K. (1998). Explorations in context space: Words, sentences, and discourse. *Discourse Processes, 25,* 211–257.

Buttler, D., Liu, L., Pu, C., Paques, H., Han, W., & Tang, W. (2001, May). *OminiSearch: A method for searching dynamic content on the Web.* SIGMOD Demonstrations Session 1: XML and Web (ACM SIGMOD/PODS 2001).

Byrne, J. G., & Barlow, T. (1993). Structured brainstorming: A method for collecting user requirements. In *Proceedings of the Human Factors and Ergonomics Society 37th Annual Meeting* (pp. 427–431). Santa Monica, CA: Human Factors and Ergonomics Society.

Card, S. K., Moran, T. P., & Newell, A. (1983). *The psychology of human–computer interaction.* Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

Carroll, J. M. (1999). Five reasons for scenario-based design. In *Proceedings of the Hawaii International Conference on System Sciences* (p. 123). Los Alamitos, CA: IEEE Computer Society.

Chandler, D. (2001). *Semiotics: The basics.* London: Routledge.

Chi, M. T. H., Glaser, R., & Farr, M. J. (Eds.). (1988). *The nature of expertise.* Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

Chin, J. P., Diehl, V. A., & Norman, K. (1987). Development of an instrument measuring user satisfaction of the human–computer interface. In *Proceedings of the ACM CHI '88* (pp. 213–218). New York: ACM Press.
Choi, Y. S., & Yoo, S. I. (1999). Multi-agent learning approach to WWW information retrieval using neural network information retrieval agents. In *Proceedings of the 1999 International Conference on Intelligent User Interfaces* (pp. 23–30). New York: ACM Press.
Chute, C. G. (1991). Classification and retrieval of patient records using natural language: An experimental application of latent semantic analysis. *Annual International Conference of IEEE Engineering in Medicine and Biology Society, 13,* 1162–1163.
Cooke, N. J. (1994). Varieties of knowledge elicitation techniques. *International Journal of Human–Computer Studies, 41,* 801–849.
Cooke, N. J., Durso, F. T., & Schvaneveldt, R. W. (1986). Recall and measures of memory organization. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 12,* 538–549.
Crandall, B., & Calderwood, R. (1989). *Clinical assessment skills of experienced neonatal intensive care nurses* (Contract No. 1 R43 NR01911 01). Yellowsprings, OH: Klein Associates Inc.
Creative Good. (1999–2001). *The dotcom survival guide.* Retrieved June 12, 2000, from http://www.creativegood.com/survival
Crow, D., & Yang, M. (2001). Reactivity. *Interactions, 8*(2), 71–74.
Cymfony (2001). *Cymfony-about.* Retrieved August 26, 2001, from http://www.cymfony.com/about.html
Dauer, J. A. (1990). Successful knowledge elicitation. *Information Executive, 3*(1), 54–56.
Dayton, T., McFarland, A., & Kramer, J. (1998). Bridging user needs to object oriented GUI prototype via task object design. In L. E. Wood (Ed.), *User interface design: Bridging the gap from user requirements to design* (pp. 15–56). New York: CRC.
Deerwester, S., Dumais, S. T., Landauer, T. K., Furnas, G. W., & Harshman, R. A. (1990). Indexing by latent semantic analysis. *Journal of the American Society for Information Science, 41,* 391–407.
Delbecq, A. L., & Van den Ven, A. H. (1971). A group process model for identification and program planning. *Journal of Applied Behavioral Sciences, 7,* 466–492.
Denning, P. (1992). Electronic junk. *Communication of the ACM, 25*(3), 163–165.
DiPiazza, J. S. (1990). Interweaving knowledge extracting, organizing, and evaluating: A concrete design for preventing logic and structure bugs while interviewing experts. *Journal of Automated Reasoning, 6,* 299–317.
Dumais, S. T., Furnas, G. W., Landauer, T. K., & Deerwester, S. (1988). Using latent semantic analysis to improve information retrieval. In *Proceedings of CHI '88: Conference on Human Factors in Computing* (pp. 281–285). New York: ACM Press.
Eason, K. D. (1984). Towards the experimental study of usability. *Behaviour and Information Technology, 3,* 133–143.
Ericsson, K. A., & Simon, H. A. (1984). *Protocol analysis: Verbal reports as data.* Cambridge, MA: MIT Press.
European Commission. (2000, June). *eEurope 2002—An information society for all—Draft action plan.* Prepared by the European Commission for the European Council in Feira, Brussels.
Fang, X., & Salvendy, G. (2000). Keyword comparison: A user-centered feature for improving web search tools. *International Journal of Human–Computer Studies, 52,* 915–931.
Fang, X., & Xu, S. (2001). Design for better information searching. In M. J. Smith, G. Salvendy, D. Harris, & R. J. Koubek (Eds.), *Usability evaluation and interface design: Cognitive engineering, intelligent agents, and virtual reality* (Vol. 1, pp. 848–852). Mahwah, NJ: Lawrence Erlbaum Associates, Inc.
Fitts, P. M., & Posner, M. I. (1967). *Human performance.* Belmont, CA: Brooks/Cole.

Flach, J. M. (2000). Discovering situated meaning: An ecological approach to task analysis. In J. M. Shraggen, S. F. Chipman, & V. L. Shalin (Eds.), *Cognitive task analysis* (pp. 87–100). Mahwah, NJ: Lawrence Erlbaum Associates, Inc.

Flach, J. M., & Dominguez, C. O. (1995). Use-centered design. *Ergonomics in Design, 3*(3), 19–24.

Flach, J. M., & Rasmussen, J. (2000). Cognitive engineering: Designing for situation awareness. In N. B. Sarter & R. Amalberti (Eds.), *Cognitive engineering in the aviation domain* (pp. 153–179). Mahwah, NJ: Lawrence Erlbaum Associates, Inc.

Fleming, J. (1998). *Web navigation: Designing the user experience.* Sebastopol, CA: O'Reilly.

Foltz, P. W., Kintsch, W., & Landauer, T. K. (1998). The measurement of textual coherence with latent semantic analysis. *Discourse Processes, 25,* 285–307.

Fulton-Suri, J. (April, 1999). The next 50 years: Future challenges and opportunities. Paper presented at *The Ergonomics Society Annual Conference 1999* in Leicester, England.

Gaines, B. R., & Boose, J. H. (Eds.). (1988). *Knowledge acquisition tools for expert systems: Knowledge-based systems* (Vol. 1). London: Academic.

Garner, W. R. (1962). *Uncertainty and structure as psychological concepts.* New York: Wiley.

Garner, W. R. (1974). *The processing of information and structure.* Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

Gillan, D. J., Breedin, S. D., & Cooke, N. J. (1992). Network and multidimensional representations of the declarative knowledge of human–computer interface design experts. *International Journal of Man–Machine Studies, 36,* 587–615.

Glance, N. S. (2001). Community search assistant. In *Proceedings of the 2001 International Conference on Intelligent User Interfaces* (pp. 91–96). New York: ACM Press.

Goodhue, D. L. (1995). Understanding user evaluations of information systems. *Management Science, 41,* 1827–1844.

Gott, S. P. (2001). Bringing system requirements into focus: The power of knowledge structures. In M. J. Smith, G. Salvendy, D. Harris, & R. J. Koubek (Eds.), *Usability evaluation and interface design: Cognitive engineering, intelligent agents, and virtual reality* (Vol. 1, pp. 1400–1404). Mahwah, NJ: Lawrence Erlbaum Associates, Inc.

Gott, S. P., & Morgan, S. (2000). Front end analysis: From unimpressive beginnings to recent theory-based advances. In S. Tobias & J. D. Fletcher (Eds.), *Training and retraining: A handbook for business, industry, government, and the military* (pp. 148–170). New York: Macmillan.

Graetz, K. A., Barlow, C. B., Proulx, N., & Pape, L. J. (1997). Facilitating idea generation in computer-supported teleconferences. In *Proceedings of Group '97: International Conference on Supporting Group Work* (pp. 317–324). New York: ACM Press.

Groen, G. J., & Patel, V. L. (1988). The relationship between comprehension and reason in medical expertise. In M. T. H. Chi, R. Glaser, & M. J. Farr (Eds.), *The nature of expertise* (pp. 287–310). Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

Hackos, J. T., & Redish, J. C. (1998). *User and task analysis for interface design.* New York: Wiley.

Hagen, P. R., Manning, H., & Paul, Y. (2000, June). *Must search stink?* Retrieved January 27, 2002, from http://www.forrester.com/ER/Research/Report/0,1338,9412,00.html

Haskin, D. (1997). *The right search engine.* Retrieved February 25, 1998, from http://www.internetworld.com/print/monthly/1997/09/report.html

Hearst, M. A. (1997). *Interfaces for searching the web.* Retrieved February 25, 1998, from http://www.sciam.com/0397issue/0397hearst.html

Hutchins, E. (1995). *Cognition in the wild.* Cambridge, MA: MIT Press.

IBM. (n.d.). *Cost justifying ease of use.* Retrieved January 27, 2002, from http://www-3.ibm.com/ibm/easy/eou_ext.nsf/Publish/23

InFoPeople Project. (2002). *Search tools chart.* Retrieved August 26, 2001, from http://www.infopeople.org/search/chart.html

Jacko, J. A., & Salvendy, G. (1996). Hierarchical menu design: Breadth, depth and task complexity. *Perceptual & Motor Skills, 82,* 1187–1201.
Jacko, J. A., Salvendy, G., & Koubek, R. J. (1995). Modeling of menu design in computerized work. *Interacting With Computers, 7,* 304–330.
John, B. E., & Kieras, D. E. (1996a). Using GOMS for user interface design and evaluation: Which technique? *ACM Transactions on Computer–Human Interaction, 3,* 287–319.
John, B. E., & Kieras, D. E. (1996b). The GOMS family of user interface analysis techniques: Comparison and contrast. *ACM Transactions on Computer–Human Interaction, 3,* 320–351.
Jungk, B., Thull, B., Hoeft, A., & Rau, G. (November 11-18, 1999). Evaluation of an ecological interface for the anaesthesia workplace by eye-tracking. In *Proceedings of the 10th ESCTAIC Meeting Glasgow.* Retrieved January 27, 2002, from http://www.eur.nl/fgg/anest/esctaic/meeting/jungk2.html
Karat, C.-M. (1990). Cost–benefit analysis of usability engineering techniques. In *Proceedings of the Human Factors Society* (pp. 839–843). Santa Monica, CA: Human Factors Society.
Karat, C.-M. (1994). A business case approach to usability cost justification. In R. Bias & D. Mayhew (Eds.), *Cost-justifying usability* (pp. 45–70). New York: Academic.
Karat, C.-M. (1997). Cost-justifying usability engineering in the software life cycle. In M. Helander, T. Landauer, & P. Prabhu (Eds), *Handbook of human–computer interaction* (2nd ed., pp. 767–778). Amsterdam: North-Holland.
Karat, J. (1997). Software evaluation methodologies. In M. Helander, T. K. Landauer, & P. Prabhu (Eds.), *Handbook of human–computer interaction* (2nd ed., pp. 689–715). Amsterdam: North-Holland.
Kieras, D. E. (1997). A guide to GOMS model usability evaluation using NGOMSL. In M. Helander, T. Landauer, & P. Prabhu (Eds.), *Handbook of human–computer interaction* (2nd ed., pp. 733–766). Amsterdam: North-Holland.
Kintsch, W. (1974). *The representation of meaning in memory.* Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.
Kintsch, W. (1998). *Comprehension: A paradigm for cognition.* New York: Cambridge University Press.
Klein, G. (1999). Applied decision making. In P. A. Hancock (Ed.), *Human performance and ergonomics* (pp. 87–107). San Diego, CA: Academic.
Lafrehiere, D., Dayton, T., & Muller, M. (1999). Tutorial 27: Variations of a theme: Card-based techniques for participatory analysis and design. In *Proceedings of CHI '99: Human Factors in Computing Systems* (pp. 1–81). Pittsburgh: ACM SIGCHI.
Landauer, T. K. (1998). Learning and representing verbal meaning: The latent semantic analysis theory. *Current Directions in Psychological Science, 7,* 161–164.
Landauer, T. K., Foltz, P., & Laham, R. D. (1998). An introduction to latent semantic analysis. *Discourse Processes, 25,* 259–284.
Lehto, M. R., Boose, J., Sharit, J., & Salvendy, G. (1992). Knowledge acquisition. In G. Salvendy (Ed.), *Handbook of industrial engineering* (2nd ed., pp. 1495–1545). New York: Wiley.
Lightner, N. J. (2001). Strategies for designing usable interfaces for Internet applications. In M. J. Smith, G. Salvendy, D. Harris, & R. J. Koubek (Eds.), *Usability evaluation and interface design: Cognitive engineering, intelligent agents, and virtual reality* (Vol. 1, pp. 1387–1389). Mahwah, NJ: Lawrence Erlbaum Associates, Inc.
Lightner, N. J., & Jackson, J. (2001). Adaptive Web interfaces for electronic commerce. In M. J. Smith, G. Salvendy, D. Harris, & R. J. Koubek (Eds.), *Usability evaluation and interface design: Cognitive engineering, intelligent agents, and virtual reality* (Vol. 1, pp. 868–872). Mahwah, NJ: Lawrence Erlbaum Associates, Inc.

Lytje, I. (2001). A cognitive linguistic perspective on the user interface. In M. J. Smith, G. Salvendy, D. Harris, & R. J. Koubek (Eds.), *Usability evaluation and interface design: Cognitive engineering, intelligent agents, and virtual reality* (Vol. 1, pp. 228–232). Mahwah, NJ: Lawrence Erlbaum Associates, Inc.

Marchonini, G. (1995). *Information seeking in electronic environments.* London: Cambridge University Press.

Martel, A. (1998). Application of ergonomics and consumer feedback to product design at whirlpool. In N. Stanton (Ed.), *Human factors in consumer products* (pp. 107–126). London: Taylor & Francis.

Mayhew, D. J. (1999). *The usability engineering lifecycle: A practitioner's handbook for user interface design.* San Francisco, CA: Kaufmann.

Mc Kevitt, P., Partridge, D., & Wilks, Y. (1999) Why machines should analyse intention in natural language dialogue. *International Journal of Human–Computer Studies, 51,* 947–989.

Miller, G. A. (1992). *The science of words.* New York: Freeman.

Mizzaro, S. (1997). Relevance: The whole history. *Journal of the American Society for Information Science, 48,* 810–832.

Najjar, L. J. (1990). *Using color effectively (TR 52.0018).* Retrieved January 27, 2002, from http://mime1.gtri.gatech.edu/mime/papers/colorTR.html

Najjar, L. J. (1992). *Multimedia user interface design guidelines (TM 52.0046).* Retrieved January 27, 2002, from http://mime1.gtri.gatech.edu/mime/papers/multiTR.html

Najjar, L. J. (2000). Conceptual user interface: A new tool for designing e-commerce user interfaces. *Internetworking, 3.3.* Retrieved January 27, 2002, from http://www.internettg.org/newsletter/dec00/article_cui.html

Najjar, L. (2001). E-commerce user interface design for the Web. In M. J. Smith, G. Salvendy, D. Harris, & R. J. Koubek (Eds.), *Usability evaluation and interface design: Cognitive engineering, intelligent agents, and virtual reality* (Vol. 1, pp. 843–847). Mahwah, NJ: Lawrence Erlbaum Associates, Inc.

Nielsen, J. (1993). *Usability engineering.* Cambridge, MA: Academic.

Nielsen, J. (1996). *The death of file systems.* Retrieved August 26, 2001, from http://www.useit.com/papers/filedeath.html

Nielsen, J. (1997a). *Top ten mistakes in Web design.* Retrieved August 26, 2001, from http://www.useit.com/alertbox/9605.html

Nielsen, J. (1997b). Usability testing. In G. Salvendy (Ed), *Handbook of human factors and ergonomics* (2nd ed., pp. 1543–1568). New York: Wiley.

Nielsen, J. (1999a). *Designing Web usability: The practice of simplicity.* Indianapolis, IN: New Riders.

Nielsen, J. (1999b). *Ten good deeds in Web design.* Retrieved August 26, 2001, from http://www.useit.com/alertbox/991003.html

Nielsen, J. (1999c). *The top ten new mistakes in Web design.* Retrieved August 26, 2001, from http://www.useit.com/alertbox/990530.html

Nielsen, J. (2001a). *First rule of usability? Don't listen to users.* Retrieved August 5, 2001, from http://www.useit.com/alertbox/20010805.html

Nielsen, J. (2001b). *Ten usability heuristics.* Retrieved October 1, 2001, from http://www.useit.com/papers/heuristic/heuristic_list.html

Nielsen, J., & Coyne, K. P. (2001, February 15). A useful investment. *CIO Magazine*. Retreived January 27, 2001 from http://www.cio.com/archive/021501/et_pundits.html

Nielsen, J., Molich, R., Snyder, C., & Farrell, S. (2000). *E-commerce user experience.* Fremont, CA: Nielsen Norman Group.

Nielsen, J., & Tahir, M. (2001, February). Building web sites with depth. *Webtechniques, 2001*(2). Retrieved January 27, 2002, from http://www.Webtechniques.com/archives/2001/02/nielson

Peirce, C. S. (1878a). Deduction, induction, and hypothesis. *Popular Science Monthly, 13,* 470–482.

Peirce, C. S. (1878b). How to make our ideas clear. *Popular Science Monthly, 12,* 286–302.

Pejtersen, A. M. (1989). *The BOOK HOUSE: Modeling users' needs and search strategies as a basis for system design* (Riso–M–2794). Roskilde, Denmark: Riso National Laboratory.

Proctor, R. W., & Van Zandt, T. (1994). *Human factors in simple and complex systems.* Boston, MA: Allyn & Bacon.

Rasmussen, J. (1986). *Information processing and human–machine interaction: An approach to cognitive engineering.* Amsterdam: North-Holland.

Rasmussen, J., Pejtersen, A. M., & Goodstein, L. P. (1994). *Cognitive systems engineering.* New York: Wiley.

Rieman, J. (1993). The diary study: A workplace-oriented research tool to guide laboratory efforts. In *Conference proceedings on human factors in computing systems* (pp. 321–326). Amsterdam.

Rieman, J. (1996). A field study of exploratory learning strategies. *ACM Transactions on Computer–Human Interaction, 3*(3), 189–218.

Rosenfeld, L., & Morville, P. (1998). *Information architecture for the World Wide Web.* Sebastopol, CA: O'Reilly.

Rossano, M. J., West, S. O., Robertson, T. J., Wayne, M. C., & Chase, R. B. (1999). The acquisition of route and survey knowledge from computer models. *Journal of Environmental Psychology, 19,* 101–115.

Salton, G., & McGill, M. J. (1983). *Introduction to modern information retrieval.* New York: McGraw-Hill.

Savidis, A., Akoumianakis, D., & Stephanidis, C. (2001). The unified user interface design method. In C. Stephanidis (Ed.), *User interfaces for all—Concepts, methods and tools* (pp. 417–440). Mahwah, NJ: Lawrence Erlbaum Associates, Inc.

Savidis, A., & Stephanidis. C. (2001). The unified user interface software architecture. In C. Stephanidis (Ed.), *User interfaces for all—Concepts, methods and tools* (pp. 389–416) Mahwah, NJ: Lawrence Erlbaum Associates, Inc.

Scerbo, M. W. (1995). Usability testing. In J. Weimer (Ed.), *Research techniques in human engineering* (pp. 72–111). Englewood Cliffs, NJ: Prentice Hall.

Schmelzer, R. (2001). *The pros and cons of XML.* Retrieved January 27, 2002, from http://www.zapthink.com/reports/proscons.html

Sell, P. S. (1985). *Expert systems: A practical introduction.* New York: Wiley.

Semantic Web Agreement Group. (2001). *What is the semantic web?* Retrieved January 7, 2002 from http://swag.semanticweb.org/whatisSW

Shadbolt, N., & Burton, M. (1995). Knowledge elicitation: A systematic approach. In J. R. Wilson & E. N. Corlett (Eds.), *Evaluation of human work* (2nd ed., pp. 406–440). London: Taylor & Francis.

Shneiderman, B. (1992). Iterative design, testing, and evaluation. In B. Sheiderman (Ed.), *Designing the user interface: Strategies for effective human–computer interaction* (pp. 471–504). New York: Addison-Wesley.

Shneiderman, B. (1997). Designing information-abundant Web sites: Issues and recommendations. *International Journal of Human–Computer Studies, 47,* 5–29.

Shraggen, J. M., Chipman, S. F., & Shalin, V. L. (Eds.). (2000). *Cognitive task analysis.* Mahwah, NJ: Lawrence Erlbaum Associates, Inc.

Smith, S. L., & Mosier, J. N. (1986). *Guidelines for designing user interface software* (Report No. ESD–TR–86–278). Bedford, MA: MITRE Corporation, Electronic Systems Division.

Soto, R. (1999). Learning and performing by exploration: Label quality measured by latent semantic analysis. In *Proceedings of CHI '99* (pp. 15–20). Pittsburgh: ACM.

Stanney, K. M., Maxey, J., & Salvendy, G. (1997). Socially-centered design. In G. Salvendy (Ed.), *Handbook of Human Factors and Ergonomics* (2nd ed., pp. 637–656). New York: Wiley.

Stanney, K. M., Smith, M. J., Carayon, P., & Salvendy, G. (2001). Human–computer interaction. In G. Salvendy (Ed.), *Handbook of Industrial Engineering* (3rd ed., pp. 1192–1236). New York: Wiley.

Stephanidis, C. (2001). User interfaces for all: New perspectives into HCI. In C. Stephanidis (Ed.), *User interfaces for all—Concepts, methods and tools* (pp. 3–20). Mahwah, NJ: Lawrence Erlbaum Associates, Inc.

Stephanidis, C. (2001). Adaptive techniques for universal access. *User Modeling and User Adapted Interaction, 11,* 1–2.

Stephanidis, C., & Emiliani, P. L. (1999). Connecting to the information society: A European perspective. *Technology and Disability Journal, 10,* 21–44.

Stephanidis, C., Paramythis, A., Akoumianakis, D., & Sfyrakis, M. (1998). Self-adapting Web-based systems: Towards universal accessibility. In C. Stephanidis & A. Waern (Eds.), *Proceedings of the 4th ERCIM Workshop on "User Interfaces for All."* Retrieved January 27, 2002, from http://ui4all.ics.forth.gr/UI4ALL-98/stephanidis2.pdf

Stonebraker, M., & Hellerstein, J. M. (2001). Content integration for e-business. In *Proceedings of the 2001 ACM SIGMOD International Conference on Management of Data* (pp. 552–560). New York: ACM Press.

Sullivan, D., & Sherman, C. (1996–2002). *Search engine features for Webmasters.* Retrieved January 27, 2002, from http://searchenginewatch.com/webmasters/features.html

Takahashi, D. (1998 October, 27). Technology companies turn to ethnographers. *Wall Street Journal.*

Tan, H. Z., Durlach, N. I., Reed, C. M., & Rabinowitz, W. M. (1999). Information transmission with a multifinger tactual display. *Perception & Psychophysics, 61,* 993–1008.

Taylor, R. S. (1986). *Value-added processes in information systems.* Norwood, NJ: Ablex.

Tedeschi, B. (1999, August 30). Good Web site design can lead to healthy sales. *New York Times e-commerce report.* Retrieved January 27, 2002, from http://www.nytimes.com/library/tech/99/08/cyber/commerce/30commerce.html

Tullis, T. S. (1993). Is user interface design just common sense? In G. Salvendy & M. J. Smith (Eds.), *Human-computer interaction: Software and hardware interfaces* (pp. 9–14). New York: Elsevier.

U.S. Department of Justice. (1990). *Americans With Disabilities Act.* Retrieved January 27, 2002, from http://www.usdoj.gov/crt/ada/adahom1.htm

Vanderheiden, G. C. (1997). Design for people with functional limitations resulting from disability, aging, or circumstance. In G. Salvendy (Ed), *Handbook of human factors and ergonomics* (2nd ed., pp. 2010–2052). New York: Wiley.

Van Dijk, T. A., & Kintsch, W. (1983). *Strategies of discourse comprehension.* New York: Academic.

Vaughan, M. W. (2000). *Identifying regularities in users' conceptions of information spaces: Designing for structural genre conventions and mental representations of structure for web-based newspapers.* Unpublished doctoral dissertation, Indiana University.

Vaughan, M. W., Candland, K. M., & Wichansky, A. M. (2001). Information architecture of a customer Web application: Blending content and transactions. In M. J. Smith, G. Salvendy, D. Harris, & R. J. Koubek (Eds.), *Usability evaluation and interface design: Cognitive engineering, intelligent agents, and virtual reality* (Vol. 1, pp. 833–837). Mahwah, NJ: Lawrence Erlbaum Associates, Inc.

Vicente, K. J. (1999). *Cognitive work analysis: Toward safe, productive, and healthy computer-based work.* Mahwah, NJ: Lawrence Erlbaum Associates, Inc.

Vicente, K. J., & Rasmussen, J. (1992). Ecological interface design: Theoretical foundations. *IEEE Transactions on Systems, Man, and Cybernetics, 22,* 589–606.

Vu, K.-P. L., Hanley, G. L., Strybel, T. Z., & Proctor, R. W. (2000). Metacognitive processes in human–computer interaction: Self-assessments of knowledge as predictors of computer expertise. *International Journal of Human–Computer Interaction, 12,* 43–71.

Weber, H. (1999). Design for all—A sketch of challenges for HCI designers. In H. J. Bullinger & J. Ziegler (Eds.), *Human–computer interaction: Communications, cooperation, and application design* (pp. 777–781). Mahwah, NJ: Lawrence Erlbaum Associates, Inc.

Welbank, M. (1990). An overview of knowledge-acquisition methods. *Interacting With Computers, 2*(1), 83–91.

Wharton, C., Rieman, J., Lewis, C., & Polson, P. (1994). The cognitive walkthrough method: A practitioner's guide. In J. Nielsen & R. L. Mack (Eds.), *Usability inspection methods* (pp. 105–140). New York: Wiley.

Williams, K., & Kotnour, T. (1993). *Knowledge acquisition: A review of manual, machine-aided and machine learning methods.* (Interim Report Contract No. N00014–91–J–1500). Arlington, VA: Office of Naval Research.

Wixon, D., & Jones, S. (1992). Usability for fun and profit: A case study of the design of DEC RALLY version 2. In C. Lewis & P. Polson (Eds.), *Human computer interface design: Success cases, emerging method and real world context*. New York: Springer Verlag.

Wolfe, M. B. W., Schreiner, M. E., Rehder, B., Laham, D., Foltz, P. W., Kintsch, W., et al. (1998). Learning from text: Matching readers and texts by latent semantic analysis. *Discourse Processes, 25,* 309–336.

World Wide Web Consortium (W3C). (1997–2002). *Web Accessibility Initiative (WAI).* Retrieved January 27, 2002, from http://www.w3c.org/wai

Yang, Y., & Lehto, M. (2001). Increasing access of visually disabled users to the World Wide Web. In M. J. Smith, G. Salvendy, D. Harris, & R. J. Koubek (Eds.), *Usability evaluation and interface design: Cognitive engineering, intelligent agents, and virtual reality* (Vol. 1, pp. 853–857). Mahwah, NJ: Lawrence Erlbaum Associates, Inc.

Zhu, W. (2001). Designing and evaluating a Web-based collaboration application: A case study. In M. J. Smith, G. Salvendy, D. Harris, & R. J. Koubek (Eds.), *Usability evaluation and interface design: Cognitive engineering, intelligent agents, and virtual reality* (Vol. 1, pp. 838–842). Mahwah, NJ: Lawrence Erlbaum Associates, Inc.

APPENDIX

Description of Participants

Session Organizers

Robert W. Proctor is Professor of Psychology at Purdue University. His research focuses on basic and applied aspects of human performance, including stimulus-response compatibility, multiple task performance, measurement of attentional demands, and skill acquisition.

Kim-Phuong L. Vu is a doctoral student in Cognitive Psychology at Purdue University. Her research focuses on basic and applied aspects of human performance, including display-control compatibility, metacognition, human factors, and human-computer interaction.

Gavriel Salvendy is Professor of Industrial Engineering at Purdue University and Head of the Department of Industrial Engineering at Tsingua University in China. He is a member of the National Academy of Engineering.

Participants (in alphabetical order):

Helmut Degen is a project engineer for Siemens Corporate Technology's User Interface Design Center.

Xiaowen Fang is an assistant professor in the School of Computer Science, Telecommunications and Information Systems at DePaul University. His research focuses on user-centered design of Web search tools, usability of e-business, and usability of mobile commerce.

John M. Flach is a Professor in the Psychology Department at Wright State University. His interests are in general issues of coordination and control in cognitive systems. Specific research topics include visual control of locomotion, visual interface design, decision making, motor control, and dynamic tactile displays.

Sherrie P. Gott is a senior research associate of the Center for Health Economics and Policy, at University of Texas Health Center at San Antonio.

Douglas Herrmann is Chair of the Psychology Department at Indiana State University at Terre Haute. His areas of research include memory and cognition, including investigations that focus on age and cultural differences.

Heidi Krömker is the head of Siemens Corporate Technology's User Interface Design Center.

Nancy J. Lightner is an assistant professor at the Darla Moore School of Business, Department of Management Science.

Kem Lubin is a usability specialist with a background in multimedia, communication and information design. She works at Siemens Corporate Research, Inc.

Lawrence Najjar is an information architect at Viant. He designs user interfaces that are simple and easy to use.

Leah Reeves is a doctoral student in Industrial Engineering at the University of Central Florida. Her research interests involve both basic and applied aspects of human-machine interaction, including virtual and multi-modal environment systems and the Internet, as well as consumer products.

Arnold Rudorfer is the lab manager for Siemens User Interface Design Labs, US office located in Princeton, NJ.

Kay Stanney is an associate professor with the University of Central Florida's Industrial Engineering & Management Systems Department. Her research focuses include: Human factors issues in systems design; human-virtual environment interaction; the use of information visualization techniques to enhance human performance; and product usability testing and evaluation.

Constantine Stephanidis is Deputy Director of FORTH-ICS, and Head of its HCI & AT Laboratory. He is also a member of the Faculty at the Department of Computer Science and member of the Senate of the University of Crete. He has been engaged, as Prime Investigator, in pioneering research work and, in 1995, he introduced the concept of "User Interfaces for All" as a socio-technical goal in the context of the Information Society.

Thomas Z. Strybel is a professor of psychology and human factors at California State University Long Beach. His research interests include auditory and visual display design, selective attention, human-computer interaction, usability evaluation, and human factors.

Misha W. Vaughan is a Senior Usability Engineer in the Usability and Interface Design department at Oracle Corporation. Her experience includes the design and testing of software for Customer Relationship Management (CRM), Enterprise Resource Planning (ERP), and Business Intelligence (BI), as well as a dissertation in human-computer interaction completed at Indiana University.

Huifang Wang is a Senior User Centered Design Engineer at Cisco Systems. Her interest is in user research, design and usability evaluation of commercial software. In the area of graphical user interfaces, her primary experiences include Windows Desktop Applications, small form GUI such as IP Phone displays. In the area of speech user interface, her experiences include the basic technology research of speech recognition and TTS and their implementation in desktop dictation software, embedded devices (such as car client, PDAs) and telephony applications.

Harald Weber is visiting researcher at FORTH-ICS. His research focuses on Inclusive Design, new technologies and organisational change, vocational and social integration of people with disabilities, occupational health and safety, and Universal Access. His new affiliation from 1st October 2001 is Institute of Technology and Work (ITA), Kaiserslautern, Germany

Yanxia Yang is a Senior Member of Technical Staff at Cadence Design Systems, Inc. She focuses on HCI, User Interface Design and Usability Evaluation for software development. She graduated with her Ph.D. degree from the School of Industrial Engineering at Purdue University.

Wenli Zhu is a usability engineer at Microsoft Corporation. Her research interests focus on the application aspects of human-computer interaction, in particular, software usability, Web-based applications and information retrieval, and the organizational and social impact of technology.