

PATTERNS OF INFORMATION SEEKING ON THE WEB: A QUALITATIVE STUDY OF DOMAIN EXPERTISE AND WEB EXPERTISE

CHRISTINE JENKINS
CYNTHIA L. CORRITORE
SUSAN WIEDENBECK

ABSTRACT

This research examines the pattern of Web information seeking in four groups of nurses with different combinations of domain expertise and Web expertise. Protocols were gathered as the nurses carried out information-seeking tasks in the domain of osteoporosis. Domain and Web novices searched breadth-first and did little or no evaluation of the results. Domain expert/Web novices also searched breadth-first but evaluated information more thoroughly using osteoporosis knowledge. Domain novice/Web experts searched in a mixed, breadth-first/depth-first pattern and attempted to evaluate information using general criteria. Domain expert/Web experts carried out depth-first searches, following deep trails of information and evaluated information based on the most varied and sophisticated criteria. The results suggest that there are distinct differences in searching patterns related to expertise. Implications of these findings and suggestions for future research are provided.

Christine A. Jenkins is a Solution Architect with Sprint Corporation.
christine_jenkin@hotmail.com.

Cynthia L. Corritore is a professor at the College of Business Administration, Creighton University, Omaha, NE 68178 USA. cindy@creighton.edu.

Susan Wiedenbeck is a professor at the College of Information Science and Technology, Drexel University, Philadelphia, PA 19104 USA. susan.wiedenbeck@cis.drexel.edu.

Before the mid-1990s expert librarians served as search intermediaries, using services such as Dialog that provided access to databases, a common command-based search interface, and a common presentation format. Because each database was different, skilled intermediaries were needed to deal with the choice of databases, the controlled vocabulary for searching and the field formats. The results were typically bibliographic references, not full texts.

Today when people look for information, they often go directly to the World Wide Web. End-users with no formal search training use graphical Web browsers and search engines to find and retrieve full-text documents. End-users who seek information on the Web represent a revolutionary shift in information search and access. However, the Web, while opening direct online information access to end-users, also poses many challenges to users who seek and evaluate Web-based information.

This study examines the effects of subject domain and Web expertise on information seeking on the Web, asking the questions “Is there a difference between subject-domain novices and experts in Web information seeking?” and “Is there a difference between Web novices and experts in Web information seeking?” Information seeking in the medical domain was chosen because finding appropriate medical information on the Web can be challenging. Although there are many sites that offer medical information, it may be difficult to find the specific information that one wants—as in the case of the interaction of two or more medical conditions. Furthermore, evaluation of the validity of the information is difficult for non-experts (Overall 2002).

BACKGROUND RESEARCH

Literature relevant to user performance in Web-based information seeking is scattered across different areas that tend to stay separate and, as a result, the information is often not well-integrated. Significant bodies of research are found in the areas of information science, human-computer interaction, and hypertext studies. Although the literature is extensive, only a modest amount of research exists on the role of expertise in information seeking on the Web.

Web information seeking: An active stream of research on Web information-seeking involves studying the searching of the public at large. These studies use transaction log data to analyze large numbers of queries by users of Web search engines, most notably studies of the Excite search engine (e.g., Jansen, Spink and Saracevic 2000; Spink, Wolfram, Jansen and Saracevic 2001). Comparing the results of three such studies, Jansen and Pooch (2001) determined that users pose an average of two queries per session, with the queries consisting on average of only two terms. Queries are simple with very little use of Boolean operators and relatively low use of other modifiers. The number of results typically viewed by users is ten or fewer per session.

Research on Web browsing has shown that the browser's Back button is used very frequently in navigation, whereas the history list mechanism is not (Catledge and Pitkow 1995; Cockburn and Jones 1996). In a logging study of 23 participants with intermediate Web-searching experience, Tauscher and Greenberg (1997a, 1997b) found that only a small number of pages are visited frequently (e.g., search engines, individuals' default start-up pages, individual and organizational home pages). However, 58 percent of pages visited by individuals are revisits. Revisited pages usually come from the last six pages visited, and extreme recency is the best predictor of revisitation. Tauscher and Greenberg argue that recency-based approaches to the Back button and the history list fit users' actual browsing pattern better than a stack-based approach.

Researchers in hypertext systems have argued that users become disoriented and lost in non-linear documents (Conklin 1987; Nielsen 1994). The Web compounds the problem because of the immense information space. Loss of a sense of orientation can involve not knowing where one is, where to go next, how to get back to a previous site, what path one has followed or where to look for information (Edwards and Hardman 1989; Park and Kim 2000). Disorientation and loss of sense of location have been attributed to the cognitive overhead of performing many tasks simultaneously in a complex environment (Kim and Hirtle 1995), along with a lack of context (Utting and Yankelovich 1989; Park and Kim 2000). Operationally, "lostness" has been measured in terms of the number of unique and total links visited compared to an optimal path (Larson and Czerwinski 1998). Such research shows that greater depth of traversal increases lostness more than greater breadth.

Concerns about disorientation have inspired the design of tools to aid users of hypertext and the Web. These tools include the following:

- Overview diagrams (Nielsen 1994; Nilsson and Mayer 2002),
- Graphical history maps of visited Web pages (Ayers and Stasko 1995; Hightower, Ring, Helfman, Bederson and Hirtle 1998),
- Social maps of Web sites combining data from many users as a guide to navigation (Wexelblat and Maes 1999),
- The addition of links to Web pages to provide structural and temporal context (Park and Kim 2000), and
- Combinations of different navigational aids in an integrated Web tool (Hascoët 2000).

Other aids are meant to support organizing, evaluating, and making choices among retrieved items. These include schemes for organizing search results (Dumais, Cutrell and Chen, 2001) and for viewing and manipulating hypertext documents to support decision-making (Marshall and Shipman, 1997).

In work addressing both navigation and sense making, Pirolli and his colleagues have done extensive theoretical work and tool development. They model

information-seeking behavior using “information foraging theory,” which describes user behavior in terms of the trade-off of the cost of the activity versus the value of the information (Pirulli and Card, 1995). Information scent is the user’s perception of value and cost, and it guides the user’s behavior (Chi, Pirulli, Chen and Pitkow 2001). Scatter/gather interfaces support information-seeking behavior by creating a cluster hierarchy that gives the user an overview of the document space and the ability to navigate it at the desired level of granularity (Pirulli and Card, 1995; Pirulli, Schank, Hearst, and Diehl, 1996). As this work suggests, efforts to support Web browsing through visualization are becoming increasingly sophisticated. Commercialization of visual Web browsers is also proceeding with introduction of a visual Web browser by the software company, Groxix (<http://www.groxix.com/>), at the end of 2002.

Expertise in information seeking: Studies of expertise in many domains indicate differences between novices and experts in key areas, such as knowledge organization and problem representation (Chi, Glaser and Farr 1988). Several studies address features of novice and expert behavior in information seeking. Borgman (1986) showed the importance of having a good mental model of an information-retrieval system. She found that training end-users with a model of the system leads to greater success in complex, searching tasks.

Based on search outcomes, Sutcliffe, Ennis, and Watkinson (2000) divided a group of medical student searchers into good and poor searchers. They found that the poor searchers had difficulty choosing search terms, used fewer search terms and simple queries, and gave up more easily. The good searchers iterated more and evaluated retrieved material more carefully. Lazonder, Biemans and Wopereis (2000) carried out a study of Web use in which the participants were students who differed in Web experience but not in subject-domain expertise. They found that novice Web-users had more difficulty locating Web sites because of their lack of skill using search engines, but they were equivalent to experienced Web-users in locating desired information within a given Web site.

Marchionini and his colleagues have done the most extensive work dealing specifically with novice/expert behavior in information seeking. Using a small CD-ROM, full-text, hypertext database, Marchionini, Dwiggins, Katz and Lin (1990) compared retrieval by three groups of searchers: domain experts without search training, professional-search intermediaries without domain expertise, and novices with neither domain nor search expertise. All groups favored searching rather than browsing, even though hypertext is designed for browsing. However, the two expert groups differed in their approaches. Domain experts were content-driven, using their greater conceptual knowledge in disambiguating the search problems, choosing the appropriate query terms and evaluating the relevance of retrieved information. Search experts were search-driven, using their expertise in searching procedures to form more sophisticated queries and using all the search features of the system. Compared to domain experts, they had more difficulty understanding the questions asked in the search tasks, carried out more distinct

searches, spent less time reading the material retrieved, and had more uncertainty about the relevance of the material found. Subsequent studies (Marchionini, Dwiggins, Katz and Lin 1993) amplified these results by showing that domain experts use their pattern-matching ability to locate relevant results. Search experts, on the other hand, evaluate results using general heuristics (type of publication, date, etc.).

Evaluating retrieved information: Now that almost anyone can put information on the Web, information seekers often retrieve large numbers of documents. Having retrieved information, they face the problem of determining which documents to use. Users make two kinds of judgments: (1) **Relevance** judgments, which focus primarily on how well a document relates to the user's topic (Barry 1994; Fitzgerald and Galloway 2001), and (2) **Evaluation** judgments, which concern the quality of the document (Fitzgerald and Galloway 2001). Because people searching for information want relevant information of high quality, they are likely to treat these two aspects as part of the same judgment process (Fitzgerald and Galloway 2001). For example, having retrieved a document, a user may first scan to determine whether it is "on topic," then apply further evaluation criteria to decide on the quality of the information. Numerous criteria have been identified that may be used in relevance/evaluation decisions (Barry 1994; Smith, 1997; Fitzgerald and Galloway 2001) such as the topic of the document, the scope and breadth of the information, the author, the context of the information, and the recency of the information.

In addition to these traditional criteria, Smith (1997) suggests Web-related criteria, such as ease of use of the Web site, required software and hardware to view the Web site, provision of a search engine on the site and interactivity of the site.

Contribution of this research: This research differs from prior research on information seeking and expertise in several ways. First, it studies expertise on the Web rather than in the more restricted environments of traditional information-retrieval systems or non-Web-based hypertext systems. Web searching is highly complex in many ways, including the huge search space, the design complexity and variability of sites, and the variety of browsers and search engines. Results from a Web search are difficult for a searcher to evaluate, because of the huge number of sites potentially available and the fact that the provenance of the information may not be clear using traditional heuristics. While one expects many similarities to those found in studies of more restricted hypertext environments, there is the need to be alert to differences that may emerge.

Second, the study is restricted to end-users of the Web—people who are not search intermediaries and, in fact, who have no formal training in information seeking. The norm for information seeking on the Web is for users, whatever their subject-domain expertise, to be self-taught searchers.

Third, the focus is on users who differ in both subject-domain expertise and Web expertise. Neither subject-domain nor searching expertise is easily inferred in studies based on logs of search engine data, where the person behind the data is unknown. Furthermore, some data-logging studies in which the users are known have not distinguished users by expertise (e.g., Catledge and Pitkow 1995; Tauscher and Greenberg 1997a, 1997b). As a result, there is only moderate empirical evidence about novice/expert behaviors on the Web. Such data is needed as a comparison and supplement to Marchionini et al.'s (1990, 1993) data on expertise in seeking information in hypertext documents.

Fourth, the focus is on the pattern of information-seeking activities, not on the nature of the queries. Research on the nature of Web queries has provided much detailed knowledge, but *patterns* of activity during Web information seeking are less known, especially those that occur off-line.

This, then, is a qualitative, observational study of a small number of novice and expert users seeking information on the Web. It is meant to give a first-hand, more holistic view of users at work on the Web, including activities and problems of users with different expertise. It is a qualitative supplement to the valuable data from large-scale logging studies.

METHODOLOGY

Participants: The general area of the study was nursing, chosen because it is a medical area that contains clear specializations within it—allowing a clear separation of the 16 expert and novice information seekers who took part in the study. The participants were female practicing nurses, ages 25 to 61 with an average age of 46. Two of the participants had three-year, hospital-based nursing degrees, twelve had B.S. degrees in nursing, one a masters degree, and one a doctorate in nursing.

The specific domain expertise chosen for the study was the medical condition of osteoporosis. Osteoporosis was chosen for two reasons: (1) it is a common condition and thus somewhat familiar to all nurses and, (2) an expert nurse population was available through two nationally recognized osteoporosis research centers located in schools of nursing accessible to the researchers. Eight nurses recruited from the research centers were classified as domain experts. They had specialized training in osteoporosis, or they had performed research in the field. Eight hospital nurse-generalists with no special knowledge of osteoporosis were also recruited from the same institutions, and they served as domain novices in osteoporosis.

Web expertise was defined as one year or more of Internet/Web use. The participants classified as Web experts reported over five years of computer use. They used computers daily in their jobs, and they had used the Internet, the Web, and search engines for an average of 4.5 years. All but one used the Internet daily for three primary purposes: email, searching for information on the Web, and shopping on the Web.

Nurses classified as Web novices reported between one and five years of personal computer use but less than one year of Internet/Web use. All Web novices used computers daily in their work, and all except one had prior experience using the Web and search engines. One Web novice reported using the Internet daily; all others reported using it about twice a month. Web novices reported using the Internet primarily for email.

In this study domain expertise was crossed with Web expertise to form four groups:

- Domain Novice/Web Novice (DNWN)
- Domain Expert/Web Novice (DEWN)
- Domain Novice/Web Expert (DNWE)
- Domain Expert/Web Expert (DEWE)

There were four participants in each group.

MATERIALS

Two information-seeking tasks were used in the study. In the first task participants were asked to find Web sites for a family member newly diagnosed with osteoporosis that would be shared with the family member. In the second task participants were asked to find sources on the Web that they would use to make a 10-minute, professional presentation on the topic of steroid use and its effects on osteoporosis. In terms of Daft and Weick's taxonomy of organizational scanning (Daft and Weick 1984; Choo, Detlor and Turnbull 1998), the tasks are formal searches, because they involve a deliberate effort to find information about a specific issue. In terms of Ellis's model of information-seeking behavior (Ellis 1989; Choo et al. 1998), the main information-seeking behavior is differentiating, meaning a person filters and selects among the sources scanned.

The first task was considered easier than the second, because it asked for general information to be used by a non-specialist. The second task required more knowledge because of its more specialized topic and audience. Nevertheless, both tasks were meant to be doable by all participants, and this was verified in pilot testing.

Procedures: Each participant took part in an individual, two-hour session. The experimenter and the participant were present in a private room containing a PC with a 15-inch monitor and a mouse. The PC had high-speed Internet access and was equipped with the latest versions of both Internet Explorer and Netscape. Two short practice tasks were conducted to familiarize the participant with the equipment, the software, the information-seeking tasks, and thinking-aloud.

During the practice tasks, the Web novices were given help in using the search engine, navigating, and examining and saving results.

During the study, participants were allowed to use their choice of browser, Netscape or Internet Explorer, so that they would have the comfort of working with their usual tool. On opening their chosen browser, participants all began at the same university home page. From there they could proceed as they chose, by typing URLs, browsing or searching with any Web search engine. No direction was provided.

The task of finding information for a family member was always done first, followed by the task of finding sources for a presentation. In both tasks, the participant searched for appropriate Web sites and listed his/her five best sites in rank order.

Data were primarily collected using the think-aloud method. This method is a well-established means of data collection about cognition and task behavior (Ericsson and Simon 1993). The objective is to capture the user's thoughts concurrent with performing a task. The method consists of asking people to 'think aloud' while they solve a problem. The resulting data are largely qualitative. In this study, the method allowed for more direct insight into the knowledge and strategies used by the participants, as well as their behavior and affect. A limitation of the method is that it generates large amounts of verbal and video data that are very time-consuming to analyze. Therefore, the think-aloud method is used to study small numbers of participants in depth.

In this study, participants talked aloud, verbalizing their thoughts, during the tasks. Videotaping was used to record data. The audio portion of the videotape contained the voice of participant speaking aloud, whereas the video portion showed the computer screen. Notes were also made during the session to record non-verbal behaviors. The experimenter prompted participants to speak aloud whenever they were silent for 20 seconds. The prompts were neutral, e.g., "What are you thinking?"

After each search task, a retrospective review of the search task was conducted. This review consisted of an interview supplemented by the videotape. The participant and the experimenter viewed the just-recorded videotape together. The participant clarified events and comments on the tape in response to the experimenter's questions.

RESULTS

The data analysis was based primarily on the verbal protocols. Given the large amount of data generated, three participants from each group were chosen for in-depth analysis. The videotape of the fourth participant was scanned for a general sense of its content and for behaviors that were not seen in other members of the group. The protocols chosen for in-depth analysis were transcribed and annotated with descriptions of non-verbal behavior, actions from the videotape, and clarifications made in the retrospective interview.

The protocols were then divided into segments and encoded, distinguishing four types of statements as proposed by Ericsson and Simon (1993): intentions, cognitions, planning and evaluation. The protocols were analyzed independently by two of the authors, with 85% agreement on the classifications. Differences were resolved by discussion and by reviewing the videotapes.

Reporting on the task time and task outcomes is followed by a qualitative description of the patterns of Web information seeking of novices and experts, as represented in the protocol analysis and the participants' online actions.

Search time: Table 1 shows the mean task times (in minutes) and standard deviations for the four groups of participants. The DEWE group searched more quickly than the other groups. The two Web novice groups completed Task 2 faster than Task 1, while the task times within the two Web expert groups did not differ. Standard deviations were smallest for the DEWE group and largest for the two Web novice groups. Statistical analysis was not carried out because of the small number of participants.

Search outcomes: There was substantial overlap in the sites chosen by the four groups. Some of the most often-chosen sites were well-known, not-for-profit organizations—such as the National Institutes of Health (NIH) and the National Osteoporosis Foundation. For-profit sites like WebMD were also included. A binary evaluation of the sites chosen (appropriate/not appropriate) was made by a physician who was an osteoporosis expert. The evaluator was blind to the design of the experiment. All of the sites were judged “appropriate” as responses to the search tasks.

Qualitative results by group: The search patterns and characteristics of each of the four groups are presented below.

In this presentation, the “hub” is usually the main search results page returned by a search engine. Because the main interest is in the *pattern* of information-seeking activities, rather than in the nature of the queries, results are discussed in terms of “seeks” not “queries,” with seeks defined as any movement away from the hub—i.e., clicking a link on the results page to go to one of the sites listed. From that location, the seek may continue if a participant clicks other links to browse further away from the original hub. A seek ends when the participant returns to the hub. A seek may be followed by another seek from the same hub (results page) to examine more results, or by entering a new or revised query in the search engine.

1) *Domain Novice Web Novice:* All DNWN participants lacked detailed knowledge about how to search the Web, and they were vague about the concept of a search engine, even though two of the three DNWNs had used a search engine before. As a result, they were uncertain how to begin searching or, once in a search engine, they had difficulty figuring out how to use it. Sample questions

TABLE 1: MEANS IN MINUTES AND STANDARD DEVIATIONS (IN PARENTHESES) OF SEARCH TIMES ACROSS GROUPS AND TASKS

Task 1		Task 2	
DNWN 29 (15.6)	DEWN 32 (18.9)	DNWN 25 (18.0)	DEWN 28 (14.4)
DNWE 28 (12.6)	DEWE 22 (7.6)	DNWE 28 (12.6)	DEWE 22 (2.9)

thus included: “What is a search engine?” and “What do you call them? Search engine?” Statements such as “I don’t know how to get in to find what I want” were also common.

DNWNs used the default search engine that was available via a link on the university home page from which they started. Once they began the first task, DNWNs experienced difficulties with the mechanics of searching and navigation. They used few browser features, and they had difficulty understanding the browser features that they tried. For example, one participant wanted to save Web links, but she failed in her attempts to add the links using the Favorites’ button in Internet Explorer. Participants became confused during navigation, not knowing how to get back to the search engine’s results page, even though they were only one link away from it. They sometimes forgot about navigation options that they had previously discovered and used, for example, the ability to return to a previous page using the Back button.

Navigation was a hard task in itself, as expressed in the verbalizations: “I go back, forward, back, forward. Where am I?” Overall, the focus of the DNWN group was on the mechanics of information seeking, rather than on the task of finding information about osteoporosis. This focus was reflected in the verbal protocols by frequent verbalizations about how users were getting information and where they were going, but few verbalizations about the osteoporosis tasks. This focus on mechanics continued throughout the session, but it became less pronounced in the second task.

Graphics were a common source of confusion. All DNWN participants indicated in the retrospective review that colors and pictures on the Web pages distracted them and made searching difficult. The participants looked for meaning

in colors and images that did not necessarily exist. This was reflected by comments in the protocols. For example, on seeing stylized images on a Web page, one participant asked, "The little red guys [images on the screen], what do they mean?"

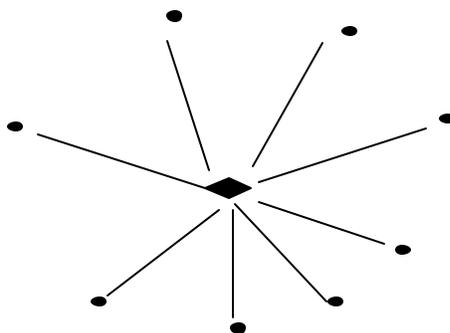
In Task 1, two of the three participants failed to retrieve results about osteoporosis during their first few attempts to use the search engine. This failure was the result of confusion about entering queries in the search engine and invoking the search. After these initial difficulties, participants did retrieve osteoporosis results. DNWNs completed an average of eight seeks before finishing Task 1. In Task 2 they carried out an average of about eight seeks as well. In their information seeking the DNWNs used only basic terms, such as "osteoporosis" and "osteoporosis treatment."

The DNWNs were timid information seekers. This is evident in their pattern of activities, which often showed self-imposed constraints on navigation. In the most extreme case, one participant was very reluctant to click on any link on the search results page. She carefully read the brief text associated with the links on the search engine results page, and in most cases she made her decision about whether to add the site to her "top 5" without visiting the Web page itself. (In fact, on seeing this behavior the experimenter verified that she *did* know she could click on the links.) The pattern of the other two participants was a classical hub-and-spoke (Catledge and Pitkow 1995; Tauscher and Greenberg 1997a 1997b).

Information seeking radiated from a central hub (see Figure 1). The hub chosen depended on the Web site's architecture. In most cases, the hub was the results page returned by the search engine. The participants clicked a link on the results page, scanned the information on the Web page that appeared, and then returned immediately to the results page. However, in some cases the Web page that appeared did not contain osteoporosis information, but rather it was an index to the site, containing links to the actual information. In this case, one of the participants made this index page the temporary hub and moved out and back from the index page to the information-content pages in a hub-and-spoke pattern. The other participant who encountered an index page simply returned to the search engine results page and chose another link.

In either case, the participants very seldom browsed further afield from their hub by clicking links in a Web page that would have taken them deeper into a site. In only one case did a DNWN click a link to another Web site, and when she saw a different style page appear, she immediately returned to her previous location.

This pattern of single hops away from a hub and back was very dominant. The pattern of search strongly resembled a breadth-first search in which participants systematically explored the edges of the hub to visit vertexes (Web sites) reachable from the hub, but did not explore in depth along any path.

FIGURE 1: DNWN SEARCH PATTERN (DOTS REPRESENT WEB SITES)

DNWN participants were also characterized by minimal evaluation of the information on their chosen Web pages. As indicated above, one participant made most of her choices without visiting the Web pages to see their content. She made comments such as, “Oh, I got a result; that’s enough.” The other two participants were also low on evaluation. They scanned pages with osteoporosis content but did not read the content in detail or systematically explore the site for further possible information.

They also made few comments about the validity of the material they found for the purpose stated in the task. As might be expected of domain novices, they did not employ osteoporosis-related, technical terminology in querying or in their verbalizations. The difficulty in dealing with the domain was especially evident in the second task, in which it was clear in the protocols that these participants did not understand well the relationship of steroid use and osteoporosis.

2) *Domain Expert Web Novice*: Starting the first task was difficult for the DEWN participants, all of whom used the search engine accessible by a link from the experiment’s homepage. In spite of their self-reported experience with the Web, DEWNs were not clear about what a search engine was, and they were not sure how to use the Web to find information, as reflected in comments like, “What is a search engine?”, “I don’t know what to put in here” and “I don’t know where to go.” They had a difficult time orienting themselves in the search space. Like the DNWNs, they became confused during navigation and could not understand how they had arrived at their current location. Problems using the browser and search engine resulted in a focus on the mechanics of searching much more than on the osteoporosis task, even though they were osteoporosis experts.

Like the DNWN participants, the DEWNs were also distracted by images. This is illustrated by comments such as “The lady image on the NOF [National Osteoporosis Foundation] Web page is appealing, and she looks friendly” or “I liked NOF because of the purple colors.” In addition, all of the DEWN participants were bothered by the distraction of advertisements on the Web pages.

The DEWN’s first few seeks dead-ended with no osteoporosis results due to their difficulty using the browser and search engine, although later searches

yielded results. Like the DNWNs, they carried out an average of eight seeks in both Task 1 and Task 2. The DEWNs used a somewhat wider variety of search terms than the DNWNs, but the queries were still simple words or phrases, such as “osteoporosis,” “general osteoporosis,” “steroids,” “asthma,” and “osteoporosis and steroids.”

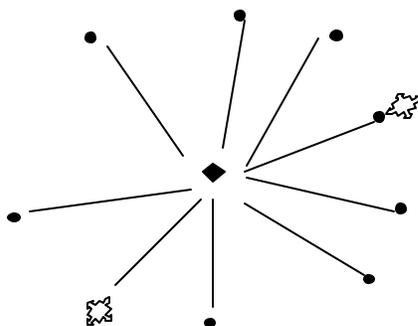
The DEWNs’ information-seeking pattern was quite similar to that of the DNWNs, corresponding to a breadth-first search in which they systematically explored the edges radiating from the hub to discover Web sites reachable from the hub. (See Figure 2.)

Again, the hub for a seek could be either the search engine results page or an index page in a Web site, depending on the architecture of the site. Nevertheless, the patterns of the DEWNs were more evolved than those of the DNWNs in that the first two seeks were strictly hub-and-spoke, with a few of the later searches characterized by more exploration.

For example, a DEWN would move to a page one level from the hub and then, unlike a DNWN, investigate it more deeply by clicking an internal link that she thought seemed interesting. Thus, the decision to move further into a site appeared to be related to these participants’ domain expertise. Still, within a particular Web site, the DEWNs never followed more than one additional internal link in the site and never clicked a link to take them to another site. They always returned to the hub via the Back button.

DEWN participants clearly used their expertise in evaluation. While they commented a great deal about their information-seeking difficulties, when they obtained results of their seeks, they inspected them closely and evaluated how well they fulfilled the task. They never simply accepted results on the search engine results page without viewing the page. The DEWNs had a much clearer sense of what they were looking for: “Ok, let’s go to facts, prevalence, symptoms, risk, we got that, detection” and “This is more prevention and not what we want.” These kinds of evaluative statements were lacking in the DNWN group. In their evaluations, the DEWNs used domain terminology. They also evaluated the type of material: “That’s not research; let’s try something else.” Their expertise was particularly evident when they evaluated information with respect to the audience. This can be seen in Task 1 (information for a family member diagnosed with osteoporosis): “I don’t know if that would throw people off, seeing that it was ‘endocrine’ and not understanding that being bone”; “Endocrine Web would be too confusing for someone.” In the second task the DEWNs clearly understood the relationship between osteoporosis and steroids. Two of them were dissatisfied with the amount of technical medical information about this topic on the Web, and they felt that they could get better information using traditional, non-Web-based medical sources.

FIGURE 2: DEWN SEARCH PATTERN
(DOTS REPRESENT WEB SITES; SNOWFLAKES REPRESENT FOLLOWING LINKS WITHIN A WEB SITE)

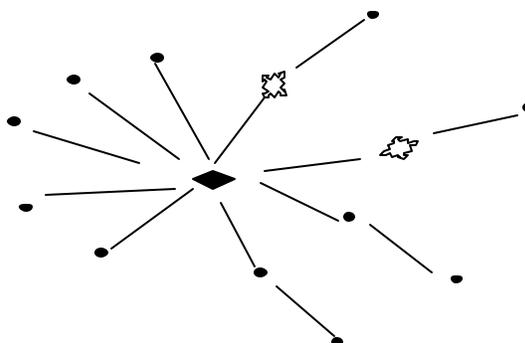


3) *Domain Novice Web Expert*: DNWEs did not exhibit confusion during information seeking. They all used the search engine on the home page of the experiment for some of their seeks, but two of them also traversed to other search tools, namely Ask Jeeves and Google. All of the DNWE participants mentioned “ease of navigation” as a key feature of successful completion of a task. They did not get confused in using Web tools, nor did they get lost in traversing the Web. The DNWE participants tended to be very focused in their information seeking. These participants expressed concern about getting the job done expeditiously and were critical of features that interfered with efficiency. For example, one participant verbalized dislike of search engines that required many steps in the query process, or that returned results in an unfamiliar format. Another expressed frustration when a search engine was not provided to search within a particular Web site.

The DNWEs carried out rapid seeks, an average of 6.5 to 7 in each task. DNWEs were at ease with the technology, but they showed their lack of subject knowledge while seeking information. This is illustrated by their willingness to try searches without much planning, for example, “Ahh...let’s just give this a try, well, what do you know, get lucky.”

The actual pattern of their seeking resembled a hybrid breadth-first/depth-first search, as shown in Figure 3. In a given seek all the DNWE participants chose a single Web site that acted as the hub of their information seeking. They began at that hub and followed a link from the hub to a site. At that point, they usually either returned to their hub or followed internal links within the site, exploring deeper along a relevant information trail if they felt they were on the right track. Sometimes, however, they extended their seek by traversing a link from the current site to another external site, resulting in the more ‘bushy’ patterns shown in Figure 3. The DNWEs navigated with more confidence than the

FIGURE 3: DNWE SEARCH PATTERN (DOTS REPRESENT WEB SITES; SNOWFLAKES REPRESENT FOLLOWING LINKS WITHIN A WEB SITE)



Web novice groups, moving freely and deeply in the Web space and showing no concerns about where they were or how they would get back to the hub.

The DNWEs were quick at finding results, using only simple terms such as “osteoporosis.” They had some difficulty evaluating results. They were aware that their goal was both to find relevant, credible sites and to evaluate them with respect to the tasks. The DNWEs expressed credibility entirely in terms of the organization sponsoring the Web site. They verbalized that they preferred to stay on Web sites of organizations with which they were familiar. For example, one participant commented, “My perception is that the National Institutes of Health is just more cutting edge information.” Another said, “I have never liked the National Institutes of Health, but I like the Mayo Clinic.” Nevertheless, in one of her seeks, a participant was observed to accept uncritically as a task solution the first item listed on the search-engine results page. As she explained, “This site [the first listed] has a lot of bearing for the search because it is on top.” Thus, the DNWEs tended to use general heuristics, both domain-oriented and search-oriented, in their evaluation. An interesting finding in this group was that DNWE participants were concerned about their ability to print pages and make copies to give to others. Not surprisingly, they preferred long scrolling pages over short ones with multiple links since they could be printed easily and also because “reading them was like reading a book.” Thus, while they may not have evaluated the content closely in terms of the information needs of the consumer, they did evaluate the format for its ease of use.

4) *Domain Expert Web Expert*: Within the Web environment DEWE participants were adept. They were familiar with browser and search engine conventions. They used a variety of search engines, such as Google, Ask Jeeves, Yahoo, and Alta Vista. All DEWE participants navigated in multiple ways. These included entering a query in a search engine, typing a known URL into the browser address box directly, using the Back and Forward buttons, and following links internal to sites as well as between sites. They knew how to return to their

chosen search engine site after following a trail that originated in a search engine results page, without multiple clicks on the Back button. They did not get lost, and their only point of even slight confusion was in fully understanding the information-seeking task—something that the other groups did not question. All DEWEs used bookmarks to mark sites that they found as they worked on a task. DEWEs were highly confident in their abilities. One of the participants said, “I am not intimidated by searching.” They were strongly focused on the domain-oriented goal of the tasks and expressed their information-seeking goals using domain terms.

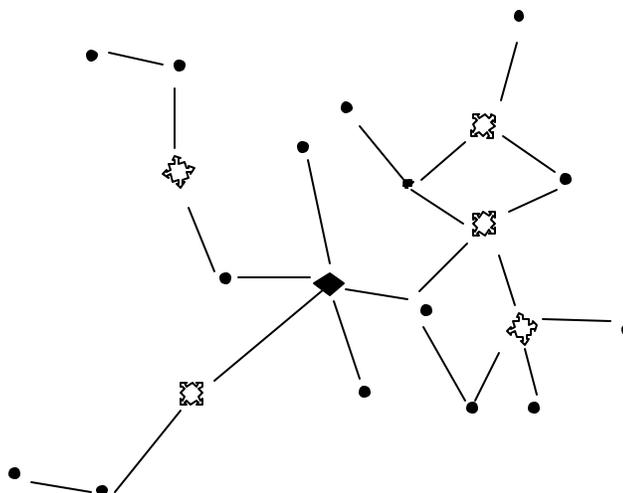
The DEWEs’ seeks were executed quickly. All of the DEWE participants performed five searches for each task, one for each of the five required Web sites they put on their list of their “top five.” They used more domain terminology and more complex searches than the other groups. For example, in Task 2 they entered the names of specific drugs and classes of steroid drugs that affect osteoporosis. They used Boolean operators and modifiers in forming queries. These behaviors were not seen in the other three groups.

Although they performed only five seeks for each task, their seeks were more extensive and took them further into the Web than the other groups. All three DEWEs browsed more deeply, following a much longer trail of information. Their seek pattern resembled a depth-first search in which they began at the search engine results page, then followed links in the site and from one site to another, following an information trail, until they found a Web site to add to their list, as in Figure 4. In a seek, the trail of links from the hub varied from one to six links. They then returned to a search-engine results page. Consistent with a depth-first search, they created many lower-level hubs within a given seek and followed many paths from these hubs.

DEWEs expressed confidence about moving very deeply along a trail in their search for information. They did not express confusion about their location or how to get back to the central hub. Although the DEWEs were similar to each other in working quickly and moving deeply, their actual paths through the Web and the specific sites visited showed much variability, something else that distinguished them from the other groups. They were persistent in seeking the best results for the task, regardless of where it led them. This persistence appeared to be driven by their domain expertise, but it was facilitated by their mastery of Web technology.

Not only were all the DEWEs fast at information seeking, but they were also fast at discerning whether information was valuable. DEWEs scanned Web pages and made quick decisions about the quality of the content, using a mix of scientific criteria and heuristics. All of these participants expressed their reasons for accepting or rejecting information. They verbalized a preference for organizations that were not-for-profit. One also expressed a preference for sites that were international in nature. As she explained, “I will always choose an

FIGURE 4: DEWE SEARCH PATTERN (DOTS REPRESENT WEB SITES; SNOWFLAKES REPRESENT FOLLOWING LINKS WITHIN A WEB SITE)



international site over a local one as it has a broader research focus.” The quality of organizations was also important to them, as evidenced by statements such as, “[This organization] is the state of the science.” Although DNWEs and DEWNs also expressed preferences for certain organizations, they did not describe the criteria that made the sites credible in the explicit way the DEWEs did. Interestingly, DEWE participants did not hesitate to use sites with which they were unfamiliar. They used factors such as the terminology used on the site, the credibility of researchers referenced at the site, and the nature of the presentation of the material to determine the credibility of the Web site.

A clear demonstration of domain expertise applied to tasks was by a DEWE participant who showed high name recognition for researchers in the field. She used her name recognition ability to evaluate information on a site, “I know that author, so I’ll use that one.” She used her knowledge of the field in evaluating the currency of the information, as well, “I know Dr. Whyte [a well-known osteoporosis researcher]; this information is old. He died a few years back.”

DEWE participants were offended by both advertisements and commercial backing of a Web site. They also expressed a preference for text over pictures. DEWEs were concerned about the end-user of the information and evaluated with that in mind. They stated that they knew what information “would be good for a patient,” or words to that effect. Like the DNWEs, they were concerned about printing and physical presentation of results. With respect to tools, it is interesting to note that DEWEs all expressed a general preference for using MEDLINE rather than the Web for medical information searches. One participant

observed that because it is a tool for clinical experts, “It’s organized for searching, has nice scientific reports and is tested for credibility.”

DISCUSSION

As has been shown in other domains, expertise appears to affect information seeking on the Web (Chi et al. 1988). This study has identified two distinct patterns of information seeking that varied with experience: one resembling a *breadth-first* search and one a *depth-first* search. These two patterns were previously identified by Catledge and Pitkow (1995) and Tauscher and Greenberg (1997a, 1997b), but their work, based on log data, did not distinguish under what conditions these patterns are likely to emerge.

In these protocols, one finds that Web novices tended to use a breadth-first pattern of information seeking. This pattern was characterized by an unwillingness to stray more than one click beyond a hub. One could refer to their method of searching as timid and attribute it in part to their difficulty staying oriented in the search space. One expects Web novices to become disoriented while following long trails of links. Instead, they became disoriented even one link from their home base. Given how easily they became confused, they appear to have adopted a strategy of staying in a small virtual space to minimize getting lost. The breadth-first pattern was seen most strongly in DNWNs, although it was also prevalent in the DEWNs, suggesting that the advantage of domain expertise did not ameliorate the Web-search disorientation and resultant timid searching. The observed tendency to employ a pattern of searching that had greater breadth than depth to minimize getting lost is consistent with the concept of “lostness” specified by Larson and Czerwinski’s (1998), who predicted that greater depth of traversal increases lostness more than does greater breadth of traversal.

The DEWEs, who never appeared to exhibit disorientation in the tasks, all consistently employed a strongly depth-first search pattern. They boldly moved out from their search results page hub, following links across multiple sites in their search for appropriate information. They did not constrain their search space, and they used the hypertext nature of the Web to its fullest. However, they too used the search results page as a kind of ‘home base’ to which they returned frequently. In between these two extremes, we saw the use of a hybrid configuration by the DNWEs. Their search pattern contained both breadth-first and depth-first elements. This hybrid pattern was not seen so much in the DEWNs, suggesting that Web expertise may be more central in facilitating depth-first searching than domain expertise.

In general, being a Web novice led to diminished task performance. More confusion was experienced, less planning was evidenced and content evaluation was less prominent. Such poor performance in information seeking is consistent with the findings of Sutcliffe et al. (2000). We postulate that these difficulties may be tied to the lack of an appropriate mental model of the Web and Web searching,

as discussed by Borgman (1986). One would expect the lack of a good mental model to lead to an inability to predict the effects of one's actions, a lack of planning, and disorientation in the search space, all of which were seen. In addition, the lack of Web expertise and a poor mental model is likely to have led to a high cognitive load.

Such an increased cognitive load has been suggested as a reason for disorientation in a complex environment (Kim and Hirtle 1995; Park and Kim 2000). While seeking information, the Web novice groups were overwhelmed by simply trying to keep track of their location, and they repeatedly got lost. This is consistent with similar findings of Conklin (1987) and Nielsen (1989) in non-Web-based hypertext systems, although they did not differentiate between novices and experts. Consequently, the Web novices were reluctant to stray far from their central starting point. Additionally, Web novices were all very focused on the mechanics of searching the Web, with little mention of the task domain during searching. In our study, Web novices appeared to attempt to reduce their cognitive load by restricting the depth of their information seeking. During a two-hour session, Web novices became more capable at using the search engine and at finding their way back to their hub. In their later seeks, DEWNs began to occasionally browse a bit deeper; however, overall they maintained their restricted hub-and-spoke pattern, like the DNWNs. From those data, it is not clear how much more experience it would have taken them to gain greater fluidity in Web navigation.

In contrast, Web experts—whether domain novices or experts—seemed to have less cognitive load. They were able to keep the domain of the task in mind during the search process, stay oriented in the search space and consider the needs of their user—the hypothetical osteoporosis patient. For example, they expressed concern about whether the level of the material they found was suitable for the end-user, as well as if it was of the appropriate format for their patient. This finding is consistent with the work of Marchionini et al. (1990, 1993), who also found that experienced search intermediaries had the needs of their end-users in mind while searching.

An interesting finding was the low content evaluation by domain novices. Unlike many novice/expert studies, in this study the domain novices were *not* completely without relevant domain knowledge. They were practicing nurses and, as such, had familiarity with both the medical condition of osteoporosis and methods of evaluating the merit of medical information. It appears that domain novices were satisfied simply to have found something relevant to the task, accepting it without further evaluation. While DNWEs did evaluate more than DNWNs, their evaluative comments were about the format of the information rather than about the content. In contrast, domain experts all evaluated the value of the Web site and its content in some way before “accepting” the Web site as a solution to a task. Thus, it seems that domain expertise was essential for significant information evaluation, in spite of the fact that all the participants were trained nurses who had at least basic familiarity with the condition of

osteoporosis. When both domain and Web expertise were present, there seemed to be a synergistic effect—as in the DEWEs. They had the resources to question the meaning and intent of the tasks. Once having clarified the task for themselves, they searched with ease, guided to appropriate information by domain knowledge and then evaluated the retrieved information using their domain knowledge.

While this study did not set out to study credibility, participants frequently referred to the role of Web site graphics, the hosting organization and advertising in their perception of information and site credibility. Web site graphics had a different effect on participants, depending on their expertise. Web experts disliked graphics and even used them in part as an indicator of the credibility of the information on the site. It appears that fewer graphics led to a perception of greater information credibility. In contrast, Web novices—regardless of their level of domain expertise—reported simply being distracted by graphics and color. Furthermore, they made incorrect interpretations of the graphics. Information credibility was also influenced by the nature of the hosting organization. A for-profit hosting company or a Web site containing advertisements led to a perception of less credible information than information from a Web site hosted by a not-for-profit organization with no advertisements. A similar finding was reported by Fogg and his colleagues (2001).

CONCLUSIONS AND IMPLICATIONS

Studying information seeking on the Web is of great importance because more and more people use the Web as one of their primary information sources, and their information-seeking episodes are often conducted without the benefit of domain expertise, Web expertise or professional search assistance. Continued study of patterns of activity and problems in Web information seeking is needed to develop systems to support the enormous and varied user base. In this endeavor, many different methodologies add value to the body of results.

The results of this study suggest that expertise affects both seeking and evaluation of information on the Web. While these findings cannot be broadly generalized because of the small sample size and the single domain, they are in accord with the general results of Marchionini and his colleagues' studies (1990, 1993) of small hypertext collections.

In general, users seeking information on the Web possess a mix of domain and Web expertise. Commonly, they are not experts in Web searching. An implication of the lack of Web searching expertise is that systems to facilitate Web information seeking are very important in order to decrease the cognitive load of searching and thus increase the cognitive resources available for content evaluation. An important question for such systems is whether they reduce the difficulty of searching for different groups of users. One could suggest that a valuable line of study would be to empirically test (or re-test) systems discussed earlier in this paper (for example, see Nielsen 1994; Ayers and Stasko 1995; Hightower et al. 1998), including novice users, expert users, and gradations in

between. Chi et al. (2001) have proposed a system that infers the expertise of a user. This information could then be used to present an appropriate information-seeking interface. This seems a particularly important direction of research for novices, because they encounter the most difficulty.

If a search engine could reliably infer expertise, then a search engine results page for novices could display results in a format graphically similar to the hub-and-spoke in Figure 1. In this format, Web sites chosen from links on the search engine results page could be displayed in the lower part of the screen with a miniature image of the search results diagram appearing in the upper part of the screen. Thus, novices would always have their hub in sight, to which they could return with a click, while having the option of exploring the “new” site in the lower portion of the screen. Their current path could be highlighted to facilitate orientation in their search space, reducing their cognitive load and helping them to acquire a more accurate mental model. It might allow them to begin to employ more depth in searching, since they could easily see their starting point and so stay better oriented. By decreasing their cognitive load, such a system would make it more likely that novices would keep their task goals in mind and evaluate their results to the best of their ability. It would also eliminate the implicit hierarchy of a list of search results, which some novices take to indicate relevance or quality in a strong sense.

Although systems to reduce the burden of Web searching are important, they are only part of the support needed by users. Domain novices, even with good search and browsing tools, may simply lack the knowledge necessary to make good information choices. This suggests that system support for making choices also needs to be integrated into Web tools. People need help, in real time as they search, to determine the credibility of information.

The problems of credibility and trust on the Web have come into prominence as the Web has grown (e.g., Fogg and Tseng 1999; Shneiderman 2000). The problem for the information seeker becomes two-pronged: on the one hand, the non-domain expert lacks the knowledge to make decisions and, on the other hand, the provenance of much information on the Web is notoriously obscure. The fact that nurse generalists in this study did little content evaluation bodes ill for the “true” domain novice looking for medical information on the Web. For example, consider a non-medical professional looking for information on their new medical diagnosis of diabetes. The findings here suggest that such users, whether Web novices or experts, will tend not to evaluate content deeply—nor even to use content-based heuristics—but to accept the information rather uncritically.

Emerging responses to the challenge of evaluating Web-based information include cues to credibility in Web sites (e.g., author, credentials), expert ratings, and reputation systems. Reputation systems (also known as recommender systems, collaborative filtering or social navigation systems) provide a mechanism for judging who is trustworthy, when parties lack a personal history of knowledge or past history with each other (Resnick and Varian 1997; Dieberger et al. 2000;

Resnick et al. 2000). In online auction sites, such as eBay, a reputation system allows buyers and sellers to rate each other, and the aggregated ratings provide a meaningful history that can be used by another person to judge the risk of an interaction. Reputation systems are aimed mostly at commercial transactions, where two parties share the direct knowledge they have of each other. Applying the concept directly to judgments of the credibility of information is difficult because, for much information, only an expert can make a valid judgment. Recommender systems are in use in sites that sell information, such as Amazon.com, but the reader of Amazon book recommendations has very scant information about who the recommender is and whether the recommender can be trusted.

These findings also re-emphasize the importance of content over graphics. Practitioners (Spool 1998; Nielsen 1999) have argued repeatedly that graphics must support the content but be minimal so as not to distract the user. Novices appear to be particularly affected by graphics and color, so these elements should be employed in a way that correctly emphasizes the content and properly directs the novice's attention. For experts, however, the interface should minimize graphics (even advertising) to a more text-oriented format. This might also increase the credibility of the information at the site. Realizing that many users may not evaluate information makes it important to include visual, easily identified cues to credibility and quality such as authorship, credentials and independent ratings.

A limitation of this study is the small number of participants. Replication with a larger sample and in other domains would be useful to identify if similar patterns of information seeking are present. A larger sample would also allow for quantitative analysis. In addition, it is becoming evident that it is necessary to consider other individual differences in information seeking (Chen, Czerwinski and Macredie 2000). Some areas of apparent importance are cognitive style (Palmquist and Kim 2000), visual ability (Chen 2000) and associative memory (Chen 2000). These factors are of importance in themselves, but it is also necessary to determine how they interact with each other and with expertise. Considering universal access and the "digital divide," further research is also necessary on Web browsing, reading Web documents, and Web information evaluation in disadvantaged populations.

The roles of immersion and engagement have also been suggested as important in hypertext documents (Douglas and Hargadon 2000), but work on this topic to date has not focused on information-seeking tasks. However, innovative features of search and browsing tools and of Web sites should contribute to immersion and engagement, possibly leading to deeper information seeking and evaluation. Another line of research suggested by this study is research into the elements of credibility, confidence, and trust related specifically to information seeking. This study could build on related work by Fogg et al. (2001). Finally, the authors also suggest an empirical study of which factors Web users consider in

prioritizing their search results. The role of expertise, both in the domain and the Web, would be of interest in all these studies.

REFERENCES

- Ayers, E.Z. and Stasko, J.T. 1995. Using Graphic History in Browsing the World Wide Web. In *4th International World Wide Web Conference*, 11–14 December, 1995, Boston, MA. Available at: <http://www.w3.org/conferences/WWW4/papers2/270>. Accessed on: November 12, 2002.
- Barry, C.L. 1994. User-Defined Relevance Criteria: An Exploratory Study, *Journal of the American Society for Information Science*, 45(4), p. 149–159.
- Borgman, C.L. 1986. The User's Mental Model of an Information Retrieval System: An Experiment on a Prototype Online Catalog, *International Journal of Man-Machine Studies*, 24, p.47–64.
- Catledge, L. and Pitkow, J. 1995. Characterizing Browsing Strategies in the World Wide Web. In *Proceedings of the 3rd International World Wide Web Conference*, Darmstadt, Germany. Available at: http://www.igd.fhg.de/archive/1995_www95/papers/80/userpatterns/UserPatterns.Paper4.formatted.html. Accessed on: November 13, 2002.
- Chen, C. 2000. Individual Differences in a Spatial-Semantic Virtual Environment, *Journal of the American Society for Information Science*, 51(6), p. 529–542.
- Chen, C., Czerwinski, M. and Macredie, R. 2000. Individual Differences in Virtual Environments—Introduction and Overview, *Journal of the American Society for Information Science*, 51(6) p. 499–507.
- Chi, E.H., Pirolli, P., Chen, K. and Pitkow, J. 2001. Using Information Scent to Model User Information Needs and Actions on the Web. In *Proceedings of the Conference on Human Computer Interaction 2001*, NY: ACM Press p. 490-97.
- Chi, M., Glaser, R. and Farr, M.J. (eds.). 1988. *The Nature of Expertise*. Hillsdale, NJ: Erlbaum.
- Choo, C.W., Detlor, B. and Turnbull, D. 1998. A Behavioral Model of Information Seeking On the Web: Preliminary Results of a Study of How Managers and IT Specialists Use the Web. In *Proceedings of the 61st ASIS Annual Meeting '98*, Vol. 35, Medford, NJ: Information Today, Inc. p. 290-302.
- Cockburn, A. and Jones, S. 1996. Which Way Now? Analysing and Easing Inadequacies in WWW Navigation, *International Journal of Human-Computer Studies*, 45(1), p. 105–129.
- Conklin, J. (1987). Hypertext: An Introduction and Survey, *Computer*, 20(7), p. 17–41.
- Daft, R.L. and Weick, K.E. 1984. Toward A Model of Organizations as Interpretation Systems, *Academy of Management Review*, 9(2), p. 284–295.

- Dieberger, A., Dourish, A., Höök, K., Resnick, P. and Wexelblat, A. 2000. Social Navigation: Techniques for Building More Usable Systems, *Interactions*, 7(6), p. 36–45.
- Douglas, Y. and Hargadon, A. 2000. The Pleasure Principle: Immersion, Engagement, and Flow. In *Proceedings of Hypertext 2000*, NY: ACM Press p. 153-160.
- Dumais, S., Cutrell, E. and Chen, H. 2001. Optimizing Search by Showing Results in Context. In *Proceedings of the Conference on Human Computer Interaction 2001*, NY: ACM Press p. 277-283.
- Edwards, D.M. and Hardman, L. 1989. Lost in Hyperspace: Cognitive Mapping and Navigation in a Hypertext Environment. In R. McAleese (eds.), *Hypertext: Theory and Practice*, Oxford: Intellect Books p. 105-125.
- Ericsson, K.A. and Simon, H.A. 1993. *Protocol Analysis: Verbal Reports As Data*. Cambridge, MA: MIT Press.
- Ellis, D. 1989. A Behavioural Model for Information Retrieval System Design, *Journal of Information Science*, 15(4/5), p. 237–247.
- Fitzgerald, M.A. and Galloway, C. 2001. Relevance Judging, Evaluation, and Decision-Making in Virtual Libraries: A Descriptive Study, *Journal of the American Society for Information Science and Technology*, 52(12), p. 989–1025.
- Fogg, B.J., Marshall, J., Laraki, O., Osipovich, A., Varma, C., Fang, N., Paul, J., Rangnekar, A., Shon, J., Swani, P. and Treinen, M. 2001. What Makes Web Sites Credible? A Report on a Large Quantitative Study. In *Proceedings of the Conference on Human Computer Interaction 2001*, NY: ACM Press p. 61-68.
- Fogg, B.J. and Tseng, H. 1999. The Elements of Computer Credibility. In *Proceedings of the Conference on Human Computer Interaction 1999*, NY: ACM Press p. 80-96.
- Hascoët, M. 2000. A User Interface Combining Navigation Aids. In *Proceedings of Hypertext 2000*, NY: ACM p. 224-225.
- Hightower, R.R., Ring, L.T., Helfman, J.I., Bederson, B.B. and Hirtle, J.D. 1998. Graphical Multiscale Web Histories: A Study of Padprints. In *Proceedings of Hypertext 1998*, NY: ACM p. 58-65.
- Jansen, B.J. and Pooch, U. 2001. A Review of Web Searching Studies and a Framework for Future Research, *Journal of the American Society for Information Science and Technology*, 52(3), p. 235–246.
- Jansen, B., Spink, A. and Saracevic, T. 2000. Real Like, Real Users, and Real Needs: A Study and Analysis of User Queries on The Web, *Information Processing and Management*, 36(2), p. 207–227.
- Kim, H. and Hirtle, S. 1995. Spatial Metaphors and Disorientation in Hypertext Browsing, *Behaviour and Information Technology*, 14(4), p. 239–250.

- Larson, K. and Czerwinski, M. 1998. Web Page Design: Implications of Memory, Structure and Scent for Information Retrieval. In *Proceedings of the Conference on Human Computer Interaction 1998*, NY: ACM. p. 25-32.
- Lazonder, A.W., Biemans, H.J.A. and Wopereis, I.G.-J.H. 2000. Differences Between Novice and Experienced Users In Searching Information on the World Wide Web, *Journal of the American Society for Information Science*, 51(6), p. 576–581.
- Marchionini, G., Dwiggins, S. and Lin, X. 1990. Effects of Search and Subject Expertise on Information-Seeking in a Hypertext Environment. in *Proceedings of the 53rd Annual Meeting of the American Society for Information Science*, Washington, DC: American Society for Information Science p. 129-142.
- Marchionini, G., Dwiggins, S., Katz, A. and Lin, X. 1993. Information-Seeking In Full-Text End-User-Oriented Search Systems: The Roles of Domain and Search Knowledge, *Library and Information Science Research*, 15, p. 35–69.
- Marshall, C.C. and Shipman, F.M. 1997. Spatial Hypertext and the Practice of Information Triage. In *Proceedings of Hypertext1997*, NY: ACM p. 124-133.
- Nielsen, J. 1999. *Designing Web Usability*. Indianapolis, IN: New Riders Publishing.
- Nielsen, J. 1994. *Multimedia and Hypertext: The Internet and Beyond*. Boston: Academic Press.
- Nilsson, R.M. and Mayer, R.E. 2002. The Effects of Graphic Organizers Giving Cues to the Structure of a Hypertext Document on Users' Navigation Strategies and Performance, *International Journal of Human-Computer Studies*, 57, p. 1–26.
- Overall, K. 2002. Interview on Station WHYY. Philadelphia, June 7, 2002.
- Palmquist, R.A. and Kim, K.-S. 2000. Cognitive Style and On-Line Database Search Experience as Predictors of Web Search Performance, *Journal of the American Society for Information Science*, 51(6), p.558–566.
- Park, J. and Kim, J. 2000. Effects of Contextual Navigation and Aids on Browsing Diverse Web Systems. In *Proceedings of the Conference on Human Computer Interaction 2000*, NY: ACM p. 257-264.
- Pirolli, P. and Card, S. 1995. Information Foraging in Information Access Environments. In *Proceedings of the Conference on Human Computer Interaction 1995*, NY: ACM p. 51-58.
- Pirolli, P., Schank, P., Hearst, M. and Diehl, C. 1996. Scatter/Gather Browsing Communicates the Topic Structure of a Very Large Text Collection. In *Proceedings of the Conference on Human Computer Interaction 1996*, NY: ACM p. 213-220.
- Resnick, P. and Varian, H.R. 1997. Recommender Systems, *Communications of the ACM* 40(3), p.56–58.
- Resnick, P., Zeckhauser, R., Friedman, E. and Suwabara, K. 2000. Reputation Systems, *Communications of the ACM* 43(12), p.45–48.

- Shneiderman, B. 2000. Designing Trust into Online Experiences, *Communications of the ACM*, 43(12), p. 57–59.
- Smith, A. 1997. Testing the Surf Criteria for Evaluating Internet Information Resources, *The Public-Access Computer Systems Review*, 8(3).
Available at: <http://info.lib.uh.edu/pr/v8/n3/smit&n3.html>.
Accessed on: November 13, 2002.
- Spink, A., Wolfram, D., Jansen, B. J. and Saracevic, T. 2001. Searching the Web: The Public and their Queries, *Journal of the American Society for Information Science*, 53(2), p. 226–234.
- Spool, J. (1998). Real Implications for Usability Testing, *Webreview.com*, Sept. 23, 1998. Available at:
http://www.Webreview.com/1998/09_25/strategists/09_23_98_2.shtml.
Accessed on: November 13, 2002.
- Sutcliffe, A.G., Ennis, M. and Watkinson, S.J. 2000. Empirical Studies of End-user Information Searching, *Journal of the American Society for Information Science*, 51(13), p.1211–1231.
- Tauscher, L. and Greenberg, S. 1997a. How People Revisit Web Pages: Empirical Findings and Implications for the Design of History Mechanisms, *International Journal of Human-Computer Studies*, 47(1), p.97–137.
- Tauscher, L. and Greenberg, S. 1997b. Revisitation Patterns in World Wide Web Navigation. In *Proceedings of the Conference on Human Factors in Computing, CHI'97* (p. 399–406), NY: ACM.
- Utting, K. and Yankelovich, N. 1989. Context and Orientation in Hypermedia Networks, *ACM Transactions on Information Systems*, 7(1), p.58–84.
- Wexelblat, A. and Maes, P. 1999. Footprints: History-Rich Tools for Information Foraging. In *Proceedings of the Conference on Human Computer Interaction 1999*, NY: ACM p. 270-77 .