Find What You Need, Understand What You Find.

Gary Marchionini and Ryen W. White

1. Introduction

The developments of the fields of Human-Computer Interaction (HCI) and Information Retrieval (IR) have followed parallel streams with both achieving significant impact in the early part of the 21st century. The intersection of these two areas engages an active community of researchers who have influenced user interfaces for World Wide Web (WWW) sites and search engines (Marchionini, 2006). The roots of this confluence of research were nourished by pioneers in several areas, including hypertext and later in digital libraries. Of particular importance is the work of Ben Shneiderman and his collaborators at the University of Maryland's Human-Computer Interaction Laboratory (HCIL). Shneiderman not only served as an inspirational leader for others aiming to understand how people search for and use digital resources, but he also is a hands-on pioneer in developing novel user interfaces that support the information-seeking process. Although his work has significant impact in other HCI areas as well, the focus on this paper is the current state of effective and usable systems for information seeking.

The development of personal computing in the late 1970s inspired computer scientists and psychologists to collaborate on making computing accessible to non-specialists. The development of alternatives to command line interactions was led by innovations such as graphical user interfaces and pointing devices. Shneiderman's seminal paper in 1983 (Shneiderman, 1983) provided a theoretical framework for the concept of direct manipulation, which posited principles for designing interactive user interfaces. This work laid the foundation for many of the later HCIL designs for search systems and the general rubric of dynamic query systems (Shneiderman, 1994). Advances in computational power allowed theories of hypertext posited by Bush (1945) and Nelson (1983) to be implementable, and the 1980s saw the emergence of hypertext systems such as Guide at the University of Kent, NoteCards at Xerox PARC, and HyperTies at HCIL. One of the important aspects of HyperTies was the concept that Shneiderman called "embedded menus," that foreshadowed the inline hyperlinks in today's Web pages (Koved & Shneiderman, 1986). The development of Apple's HyperCard system advanced the applications of hypertext and Shneiderman and his colleagues focused attention on information structuring and evaluation, including one of the earliest papers on link analysis (Botafogo & Shneiderman, 1991). While these high-profile activities were underway, Shneiderman built upon his database background to address issues of retrieval. In collaboration with colleagues in psychology, a series of studies of library catalogs were undertaken and eventually led to a program of installing and testing touch panel workstations at the Library of Congress. The emergence of the WWW sparked new design challenges and led to the concept of dynamic query interfaces in the 1990s that anticipated the highly interactive AJAX techniques of the early years of the new century.

Inspired by the development of HyperTies and other efforts at HCIL, Marchionini and Shneiderman (1988) presented a framework for information seeking that distinguished classical query-based search strategies from highly interactive, browse-based search strategies, putting the focus on user-control over the search process. Two main themes of such systems are the involvement of an informed information seeker who takes active control over the search process, and the view that search is an iterative process that is embedded in real problems rather than a discrete, self-contained activity. These themes suggest that good search systems must put the user in control, provide support for all the subactivities of the search process across all iterations, including helping people understand and make use of search results. Thus our overall goal is to provide effortless searcher control services to achieve fluid and productive searcher experience. This paper uses an information-seeking framework to discuss progress to date in realizing such a goal, with focus on a small number of important recent results.

The paper is organized as follows. The information-seeking process is first described and the different subactivities are used to discuss the state of support and in some cases illustrate progress with results from recent studies by the authors. We conclude with thoughts about integrated systems that support information seeking and implications for future development.

2. Information-Seeking Process

Information seeking is taken to be a human activity that is part of some larger life activity. It might take place in a few seconds or over a lifetime, may be highly discrete or it may be integrated into the rhythms of daily life. Colloquially, information seeking and search are synonymous, however, we make the distinction that information seeking is a uniquely human activity and search can be undertaken by both machines and humans. Thus, most of what is termed information retrieval or information search in the WWW are actually the search episodes in a human's information-seeking activity that leverage information technology. Although information seeking is driven by human needs and behaviors and thus highly variable, there are several common subactivities that may be supported by good technical design. Ultimately, well-designed search systems aim to support these subactivities and the overall information seeking process. At present, most search systems focus on one or a few of these subactivities. As long as they are compatible with other kinds of information processing applications that support the larger goals that motivate search, this is adequate, although we look for more comprehensive systems in the future.

There is a variety of frameworks for information-seeking behavior (e.g., Ellis, 1989; Ingwersen & Jarvelin, 2005; Kulthau, 1991; Wilson, 1997) and here we adopt one that emerged as the first author collaborated with Shneiderman in the early years of the HCIL. Marchionini (1995), has described the information-seeking process as a set of activities that people undertake in a progressive and diversely iterative manner. The information seeker first **recognizes a need** for information and **accepts the challenge** to take action to fulfill the need. These subactivities are primarily cognitive and affective respectively and traditionally foreshadow actions that involve search systems. A **problem formulation** activity follows acceptance and involves the information seeker conceptualizing the bounds of the information need, imagining the nature and form of information that will meet the need, and identifying possible sources of information pertinent to the need. This activity typically does not involve a search system. Once the information need has been formulated sufficiently to take action, a search system is used to **express the information need**. Problem expression is strongly constrained by the system's user interface and thus this is an activity that has attracted considerable attention from the design community. In all but simple lookup situations, people tend to express and re-express their need over several iterations depending on what transpires as they interact with the system. Every expression act generates some kind of response from a search system and the information seeker engages in one or more **examination of results** activities. This activity tends to take the most time of all the information seeking activities as people read/view/listen to intermediate and primary content. Typically there are many results to consider and so there are many sub-iterations within this activity. Often, examination of the results does not yield the sought information or sufficient information, and the information seeker re-expresses the need or reformulates the problem. System support for these **reformulations** is also an active area of research. At some point, the information seeker makes a decision to stop the search and **use** the found information. Most search systems do not address information use. However the trend in interactive systems is to collapse the temporal gaps and distinctions in these subactivities so that they are tightly coupled or concurrent. Thus far, good progress has been made in integrating need expression, results examination, and reformulation activities. This paper uses these information seeking activities as an organizing framework for surveying the state of development with special emphasis on our own work which has been influenced by Shneiderman's passion for user control and empirical studies of innovative designs that give people control over information resources.

2.1 Recognize, accept, and formulate the Problem.

In the pre electronic era, most information needs arose in distinctly different physical settings than where search took place. For example, needs arose in face-to-face conversations, work places, and classrooms, and people turned to printed resources at hand, other people, or libraries to satisfy those needs. Clearly, needs still arise in these physical settings today, however, in today's homes, workplaces, and schools large amounts of time is spent in front of computer screens that themselves are both the stimuli for information needs and the sources for information. These work settings offers new opportunities for systems to provide support for recognition, acceptance, and formulation. Although recognition is ultimately up to human perception and cognition, electronic systems support alerting mechanisms and communication tools that identify needs or recommend new information sources.

Acceptance is strongly dependent on constraints of time. Thus, having the sources at hand rather than physically separate from the need stimulus can itself diminish one barrier to accepting the information need. If an information need arises while one works online it is much more likely that it will be accepted if a search system is at hand rather than if one must travel to a library or seek out a teacher or mentor. Easy to use and effective search systems also help people gain confidence to accept more information problems. Norman (2002) has argued that good design aesthetics improve system effectiveness and well-designed search systems at hand may also influence people to more readily accept information problems.

Problem formulation determines the effectiveness of a search and strongly determines the efficiency of a search. In intermediated search settings, problem formulation takes the form of a reference interview where the reference librarian asks questions about what is already known about the information need and what kinds of results would be useful to meet the need. Most searching does not have the benefit of professional intermediation and thus information seekers must give voice to their need alone. This mainly reduces to identifying words or phrases to use and selecting a search system, with issues like what kind of result (e.g., formal academic paper, raw data set, short article, photograph) and veracity of source kept implicit. Key success will ultimately depend on how well the information. Tools such as note pads, calendars, thesauri, encyclopedias, dictionaries, blogs and wikis can be consulted to help in problem formulation and good search systems will either integrate them or make them easily accessible on the WWW or in software applications and operating systems. In cases where problems are not well-formulated, browsing in generally pertinent resources can help to sharpen the focus, and in some cases solve the information problem at the cost of efficiency.

HCI and IR researchers are working to find ways to leverage information seeker and system context to support problem formulation through such tools as implicit and explicit recommender systems (e.g., Herlocker *et al.*, 2004) and user interfaces that take advantage of usage histories (Komlodi *et al.*, 2006). A recent trend is to find ways to leverage information technology to support and enhance creativity. Once again, Shneiderman is among the champions of this difficult and provocative research direction. His book, Leonardo's Laptop (2005) presents a vision of tools that support human creativity, including the creativity required to formulate information needs.

2.2. Problem expression.

Once the problem has been formulated in the mind of the searcher, it is necessary for them to perform a number of activities: they must select the collection they are going to search and the search system they are going to use, specify their problem in a way that is understood by the search system, and submit their query for system processing. These activities comprise the problem expression phase of the information-seeking process. Given the perceived coverage and efficiency of commercial Web search engines such as Google[1], Yahoo![2], or Windows Live Search[3], the issues of system/source selection and query execution are trivial for most of today's users. Although, when users know where the relevant material is located, they generally prefer to limit their searches to that library, collection, or range of documents (Shneiderman, Byrd & Croft, 1998). In this section we focus mainly on query formulation since this is an area where improvements can have a significant impact on information-seeking effectiveness (Marchionini, 1992).

Since the quality of queries directly affects the quality of search results (Croft & Thompson, 1987), considerable attention has been paid to eliciting complete and accurate problem descriptions from information seekers. Query formulation requires two types of mappings: a *semantic mapping* of the information seeker's vocabulary used to articulate the task onto the system's vocabulary used to gain access to the content, and an *action mapping* of the strategies and tactics that the information seeker deems best to forward the task to the rules and features that the system interface allows (Marchionini, 1995). The search queries that emerge from query formulation are only an approximate, or "compromised" information need (Taylor, 1968), and may fall short of the description necessary to retrieve relevant documents simply because the vocabularies of the user and the system differ too greatly (Furnas *et al.*, 1987).

To address differences in the semantic mapping between the information seeker's task vocabulary and the system vocabulary it is necessary to augment either or both representations to bring them into alignment. One way to do this is to dynamically expand the vocabulary of the system based on users' querying behavior and the observed frequency with which terms are used to retrieve documents. This technique is known as *adaptive indexing* (Furnas, 1985), and makes use of the keywords typed by users to perform commands or retrieve documents, to assign additional indexing terms to the documents in the collection. In the search domain, if many users type a query, then visit a particular document, the terms in that query can be assigned as additional indexing terms for that document, potentially improving future retrieval performance. The query completion and spelling correction facilities offered by search engines such as Google or Windows Live Search – either as a toolbar add-in or on their results page – represent ways in which systems can refine a user's description of their information needs rather than refining a system's descriptions of its documents. Query completion is often offered as a drop-down list below a query entry text box that is populated with popular query statements containing the same prefix as the query currently being typed. Spelling correction is generally offered on the results page as a clickable hyperlink after a query has been submitted containing a potentially misspelled word. Both of these techniques leverage the query logs generated during the search activity of many millions of Web users to help predict what the intended query formulation should be. The difference between query *completion* and – the better known – query *expansion* (c.f. Efthimiadis,

---

[1] http://www.google.com

[2] http://www.yahoo.com

[3] http://www.live.com

1996) lies in *when* the recommendations are offered during the search. Query completion is offered as the user types their query statement, and query expansion is offered after they execute their search. As we will describe in more detail in this section's example, query completion has the potential to positively impact query quality for the initial formulation (i.e., before the user has seen any search results), and query expansion can positively impact query quality in subsequent iterations (White & Marchionini, 2006).

Action mappings take possible sets of actions to the inputs that a search system can recognize, and therefore limit how queries may be expressed. The success of simple interface design adopted in commercial Web search engines has meant that many users are unfamiliar with anything other than the most basic of query forms, supporting simple keyword entry. Given that users typically now only visit one system (e.g., Google) and one collection (e.g., the Web), to conduct most of their searching, this presents an opportunity for such systems to act as portals and offer a broader range of services, including different ways to articulate queries than are currently available through the simple textual query input forms.

Search systems typically allow Boolean expressions and offer advanced search operators such as quotation marks that can improve the precision of search results, but must be learned and included in query statements. However, most users are unaware of these operators, mainly because their use is not publicized and the interface to compose queries with them is hidden from initial view. As a result, most users lack the additional skills required to formulate well-defined query statements. One approach that has proven effective is to train searchers to pose better queries by using thesauri (e.g., Sihvonen & Vakkari, 2004), or learning systematic search strategies (e.g., Bates, 1979). Although this is a good way of empowering users, it can be difficult to do on a large scale, and many users are generally more concerned with solving their information problems than learning how to search. Since teaching users querying skills may not be a viable option, systems must provide alternative, user-friendly ways to rapidly specify and refine queries.

It is possible to provide a wide range of interaction styles to support the information-seeker in expressing their problem. Such techniques include expert system intermediaries (Croft & Thompson, 1987; Fox & France, 1987), query-by-example (Zloof, 1975) dynamic thesauri (Suomela & Kekäläinen, 2005), or through eliciting details of the context that lies behind the problem (Kelly, Dollu & Fu, 2005). The Query-by-Image-Content (QBIC) system (Flickner *et al.*, 1995) allows users to query based on the visual properties of images such as color percentages, color layout, and textures occurring in the images. Such queries use the visual properties of images, so you can match colors, textures and their positions without describing them in words. Content based queries are often combined with text and keyword predicates to get powerful retrieval methods for image and multimedia databases. Embedded menus (Koved & Shneiderman, 1986), as described earlier, can be applied to hypertext links to consider the context in which they reside when menu items are selected. This could allow systems to better disambiguate user intentions. In addition, search systems can also provide hierarchical or faceted stimuli that surface the underlying organizational structure of the collection being searched and the attributes of documents that could be retrieved. Systems such as Flamenco (Hearst, 2006) provide hierarchical faceted categories of labels that are reflective of relevant concepts in a domain, and allow users to select category labels to express their problem and refine subsequent searches. Phlat (Cutrell *et al.*, 2006) and Stuff I've Seen (Dumais *et al.*, 2003) exploit the wide and varied associative and contextual cues that people remember about their own information to help them formulate queries and browse results. An interesting feature of these systems is that they allow users to assign metadata to documents to support future information-seeking episodes

Once the information need has been specified to a level that is agreeable to the information seeker, the search is then executed. The submission of a query to a search system typically marks the end of an iteration of the problem formulation phase of the information-seeking process. The separation of query creation, submission, and result examination (which will be addressed in the next section) may mean that users have to iterate many times to express their problem correctly. However, *dynamic queries* (Ahlberg & Shneiderman, 1992), allow users to formulate query statements with graphical widgets, such as sliders. As these widgets are manipulated, the system adjusts visualizations of the underlying data in real-time to allow users to easily identify trends and exceptions. Although dynamic queries are generally only of use for structured domains, they are incredibly powerful at supporting information exploration activities, and provide users with a useful set of visual stimuli (in the form of the sliders) that constrain how their problem can be specified. Dynamic overviews and previews – such as those offered in the Relation Browser (Marchionini & Brunk, 2003) – give the user information on the predicted effect of issuing the query through mouse brushing and mouse hover operations without the cognitive interruption of waiting for the retrieval of search results.

Search systems also can use traces gathered from the interaction of other users either within a document or between documents, to suggest alternative courses of action to the current user. The notion of "wear" on parts of a document (Hill *et al.,* 1992) or "footprints" between documents (Wexelblat & Maes, 1999) give a clear indication to users about where other users have been that may be useful to them in making decisions about where they should spend their time. Research in the area of information scent has also tried to characterize and visualize the search behavior of users within particular Web domains (Chi, Pirolli & Pitkow, 2000). Systems offering mediated searching capabilities (Muresan & Harper, 2002), assume the role of the human mediator or intermediary searcher, and interact with the user to support her exploration of a relatively small source collection, chosen to be representative for the problem domain. Based on the user's selection of relevant "exemplary" documents and clusters from this source collection, the system builds a language model of her information need. This model is subsequently used to derive "mediated queries," which are expected to convey precisely and comprehensively the user's information need, and can be submitted by the user to search any large and heterogeneous "target collections." This familiarizes the user with the subject area, helps them conceptualize their problem internally, and assists them in creating potentially powerful queries before exploring the full collection.

It is important to note that all of these techniques involve the user as an active participant in the specification and manipulation of problem descriptions. Techniques such as query completion and dynamic queries go a step further in that query specification is coupled closely with result examination, facilitating a fluid dialog between user and system that is vital for effective information access.

2.2.1. Example: A study of query completion.

In this example we describe and evaluate a query completion technique to support the rapid formulation and refinement of query statements for Web search. As a searcher enters their query in a text box at the interface, the interface provides a list of suggested additional query terms generated from the intermediate retrieved results, in effect offering query expansion options *while the query is formulated*. The terms are shown in a "Recommended words" list situated between the text box and the submit button used to execute the query. In Figure 1 we show a screenshot of the query completion component.

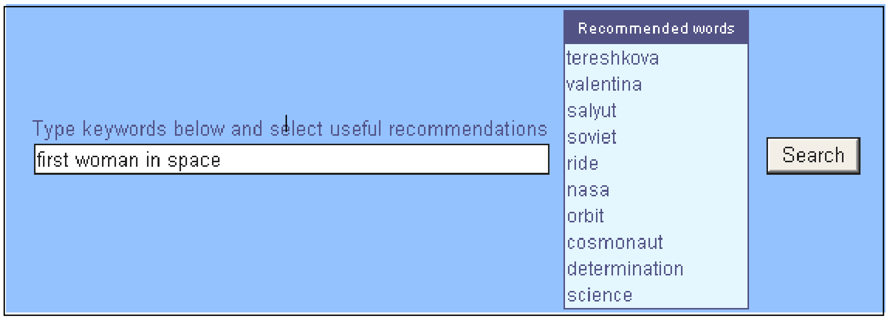

first woman in space

tereshkova
valentina
salyut
soviet
ride
nasa
orbit
cosmonaut
determination
science

Figure 1. Term suggestions in real-time at the interface. The list of "Recommended words" updates after each query word is typed in the text box. In this example the searcher has just pressed the spacebar.[4]

The additional query terms in the list act as a form of dynamic result preview, simultaneously removing the need for users to submit their query before seeing the key terms in the top search results, and supporting query formulation by suggesting alternative terms at a time where searchers may benefit most from this support (i.e., during initial query formulation). This implementation is different from that offered by Google or Windows Live Search described earlier. Those implementations use query logs to auto-complete the query rather than extracting important terms from the search results. As a result, they are more representative of what others are searching for than of what would be found if the query were to be executed. Thus, the query completion technique gives users a brief look ahead to results.

In order to determine how query completion is used – and when it may be useful – we conducted a user study involving 36 subjects in which we compared three search interfaces: a baseline interface with no query formulation support; the query completion interface (*QueryCompletion*), and a third interface that provides options after queries have been submitted to a search system (*QueryExpansion*). In particular, we used the data derived from the study to assess the quality of the queries generated across known-item and exploratory search tasks. Query quality is a complex construct that is dependent on many factors such as the searcher's knowledge about the need, search experience, system experience, and the mapping between the need and the information source. As an estimate of query quality we employed a panel of two judges who independently assessed the quality of every query expressed for all subjects using a 5-point scale. The judges met with one of the experimenters and discussed ways to assign values. The basic agreement was to examine the task, conduct a search, and then identify the key concepts in the task to use as basis for judging the subject queries. The judges then coded queries for one task together to establish a common rating scheme.

An analysis of query quality showed that offering query completion improved the quality of initial queries for both known-item and exploratory tasks, making it potentially useful during the initiation of a search, when searchers may be in most need of support. If query completion techniques are capable of enhancing the quality of some queries, and do not have a detrimental effect on other aspects of search performance, then there is a case for them to be implemented as a feature of all search systems. A promising characteristic of query completion is that it does not force searchers to use it, or indeed do anything radically beyond the scope of their normal search activities. Additional analysis of the findings, presented in (White & Marchionini, 2007), shows that compared to post-retrieval query expansion, query completion lowers task completion times,

---

[4] First woman in space: Soviet cosmonaut Valentina Tereshkova.

increases searcher engagement, and increases the uptake of system suggestions (44% of queries used suggestions in *QueryCompletion* versus 28% of queries used suggestions in *QueryExpansion*). In addition, our findings suggest that query completion made searchers more involved in their search and led to higher user satisfaction. However, the time at which query recommendations were offered did not affect the number of query terms, or the number of query iterations.

The *QueryExpansion* system offered query recommendations next to search results and led on average to the highest query quality across all queries. This may be because the system provided two types of support: searchers were shown the query recommendations, and they were shown the titles, abstracts, and URLs of the documents from which those terms were derived. The presence of this information may provide an additional source from which to choose terms, but perhaps more importantly, gives practiced, motivated searchers a sense of the type of documents that their query retrieved, and a sense for the context within which query modification terms occur in the collection.

An important finding from our study is that despite the effectiveness of query completion, it has the potential to introduce query skew if any of the recommendations are ambiguous[5]. If the technique is to be implemented for large-scale use, then care must be taken to implement it in such a way as to offer searchers some information about the predicted effect of their query formulation decisions. This study gave us insight into the circumstances under which query completion performs well, how searchers use it, and potential enhancements for the approach.

2.2.2. Other Support for Query Formulation.

Modeling the contextual factors that influence information-seeking can allow for the development of more robust information-seeking support (Ingwersen & Järvelin, 2005). Factors that need to be considered include relevance, uncertainty, user preferences, goals and motivation, task, and historic/societal/organizational contexts and traditions. The ability of systems to support information seeking may be enhanced if models can be developed that also incorporate the context within which the systems operate. It is conceivable that the systems' representation of the external context could be tuned through user input in a custom user interface. An alternative could be that systems take input from external "sensors" that report periodically on the state of the user, their search experiences, and environmental and situational factors likely to affect them. For example, Microsoft Exchange Server already monitors incoming communications in many forms, including email, telephone calls, and Instant Messaging (IM), and is aware of when a user is in meetings, where they are meeting, and who they are meeting with. With enhancements systems can also be aware of when user attention is diverted away from their personal computer (Horvitz *et al.*, 2003) or their active task (Horvitz, Jacobs & Hovel, 1999). All of these factors could be used to provide additional contextual information to a search system.

It is very seldom that a user is the first to encounter a particular information problem. Earlier in this section we described the use of techniques that leverage previous users' interactions to support the current searcher. The general focus in this area has been on passive collaboration through the use of the interaction behavior of many users to make recommendations to the current user population (Joachims, 2002; Agichtein, Brill & Dumais, 2006). However, an additional way to find relevant information is through questions that have been previously posed by other searchers. Frequently asked questions (FAQs) and Answer Gardens (Ackerman, 1998) are

---

[5] In the example shown in Figure 1, the term "ride" may seem appropriate for a journey into outer space. However, this term was recommended since Sally Kristen Ride (the first American woman in space) appears many times in the top-ranked documents. If this term was added to the query, it could certainly lead the user down incorrect search paths.

examples of applications whereby solutions to problems previously encountered are made publicly accessible for the benefit of others. However, such forums do not provide a means through which new questions can be posed and solutions sought. Social tagging (e.g,, tag clouds) in systems such as Flickr[6] and del.icio.us[7] support some of these same kinds of direct access without literally specifying a query.

The requirement for a search system to be a medium through which the user accesses a knowledge base was necessitated by the amount of information that must be searched and the requirement that an answer be furnished almost instantly. However, a beneficial side effect of the growth in the size and diversity of the Web has been a growth in size of user population. These users bring with them a diverse range of interests, expertise, and experience on a scale unimaginable two decades ago. Online question-answering communities such as Yahoo! Answers and Windows Live QnA leverage this user population to provide answers to user questions in close to real-time. Questions and answers are posed and offered as a temporally delayed dialogue written to a remote Web page visible to all who visit the site. As well as helping individual users to find the answers to their questions, these services can support the formulation and refinement of problem statements since they have to be presented in a way that is understandable by others, and can be used to create a repository of questions and answers for future reference.

As well as the passive collaboration techniques described earlier involving the use of interaction logs of many users, research has also considered providing more active collaborative experiences among groups of users who know each other. Collaborative search (Chi & Pirolli, 2006) is an emerging area of interest whereby multiple users can become involved in the pursuit of a single task. However, systems based on these principles tend to be designed for very specialized domains and/or devices. TeamSearch (Morris, Paepcke & Winograd, 2006) is a system that enables co-located groups of up to four people to simultaneously search collections of digital photographs, using a visual query language designed for a multi-user interactive tabletop. Maekawa *et al.* (2006) describe their system for groups of co-present people who each have a small, Web-enabled mobile device – to improve the efficiency of searching for information within a Web page (since scrolling through long Web pages on small screens is time-consuming), they allow a page to be divided into several parts, each of which is displayed on a different user's device to facilitate parallelization of visual search.

A few commercial products also offer support for collaboration during the performance of search tasks. For example, the search engine ChaCha[8] pairs searchers up with another person – supposedly skilled at searching and knowledgeable about the domain of interest – who assists them in formulating their query and suggesting interesting Web sites. The Windows Live Messenger IM client provides a "shared search" feature whereby conducting a Web search through the client allows the list of returned URLs to be displayed to both the searcher and his/her IM partner. Google Notebook allows a user to store clippings from several Web sites in one document; the tool provides a facility for allowing multiple users to add content to a single notebook. Tools that allow users to collaborate in the formulation and refinement of queries during an information seeking episode, and could potentially benefit users in terms of coverage, confidence, exposure, and productivity (Morris, 2007).

---

[6] http://flickr.com/

[7] http://del.icio.us/

[8] http://www.chacha.com

The advances described at the end of this section suggest that search climate is expanding from searching in isolation to encompass search as a social activity within a defined community of interest. The community may comprise work colleagues, academic peers, friends and family, or remote users with whom searchers have had no previous relationship. Interactive support for problem expression in these communities must consider social issues such as parsing and translating natural language questions, addressing cultural conventions, and directing questions to those with domain knowledge. We are all familiar with human-human interaction and potentially capable of expressing our problems in a way that is understandable to others more easily than to systems. Involving other users in problem expression has the potential to address some of the semantic mapping and action mapping constraints described earlier.

2.3. Results examination.

People spend most of their search time examining results returned by the search system. Results are often presented in lists that in turn lead to the primary object. Greene *et al.* (2000) distinguished overviews that display collections of results and previews that display abbreviated views of individual objects. The overviews can be simple lists or hierarchical lists, or visualizations, and the previews can be snippets or metadata records that stand as surrogates for the primary object. The HCIL and other groups have created and tested various systems to support easy movement between overviews and previews and the primary object. Some principles of display are well established. For example, Egan *et al.* (1989) demonstrated the benefits of highlighting query terms in the full text of results, and Norman and Chin (1988) and others demonstrated that hierarchical lists should be broad rather than deep. Card *et al.* (1999) assembled a book of readings on information visualization that includes examples from the pioneering visualization work at Xerox PARC and HCIL. We illustrate aims at improving results examination and integrating them more seamlessly with the other information seeking subactivities with two examples from the authors' recent work.

2.3.1. Example: Video surrogates

The large volume and range of digital video available on WWW demands good search tools that allow people to browse and query easily and to quickly make sense of the videos behind the result sets. Surrogates, such as the textual 'snippets' provided in the results lists of most search engines, are essential components of good user interfaces for all search systems but are even more crucial for video collections that offer information in multiple sensory channels and consume substantial human and system effort to transfer and consume. We distinguish surrogates from most metadata in that surrogates are designed to assist people to make sense of information objects without fully engaging the primary object, whereas metadata can serve this purpose but more often is meant to support retrieval and often is meant to be used by machines rather than people.

The mainstays of surrogation for all media are textual surrogates such as keywords and abstracts. However, there is considerable work devoted to video classification, segmentation, and keyframe extraction. Most of this work uses signal processing techniques focused on specific features such as color (e.g., Jain & Vailaya, 1998), texture (e.g., Carson et al., 1997), shape/objects (e.g., Smith & Kanade, 1998), faces (e.g., Senior, 1999), motion (e.g., Sim & Park, 1997; Teodosio & Bender, 1993) and higher level events (e.g., Qian et al, 2002; Smith, 2006), and various audio characteristics (e.g., Foote, 2000; Li et al., 2001; Witbrock & Hauptman, 1998). The most integrated surrogates emanating from this work are the Informedia skims (Christel et al., 1998; 1999).

Our work has focused on creating simple video surrogates, embedding them in video retrieval systems, and conducting user studies on their effectiveness for a range of information seeking subtasks. More than a dozen studies were conducted over a five year period for visual surrogates such as poster frames, storyboards, slide shows, and fast-forwards and these results were incorporated into an operational system (Open Video Project[9]). The results of these studies are reported in more than a dozen papers in the HCI and IR literature (see Marchionini et al., 2006 for a summary). Assessments were made for different variations of single keyframes (poster frames), storyboards (arrays of keyframes), slide shows, and fast forwards for a series of recognition and gist determination tasks, using task accuracy, time, and a battery of affective measures. Based on these results, the current system provides storyboard previews for all videos and fast forward previews for most of the videos.[10] The fast forwards are played at 64X speed based on empirical evidence from a comparison of a range of rates (Wildemuth *et al.*, 2003). Over the past four years the storyboards have been consistently used about twice as much as the fast forwards. This illustrates people's desire for control over surrogates as well as the additional requirement to launch a video player to view the fast forwards. Figure 2 shows a screen display for a preview of a video from the HCIL symposium series.

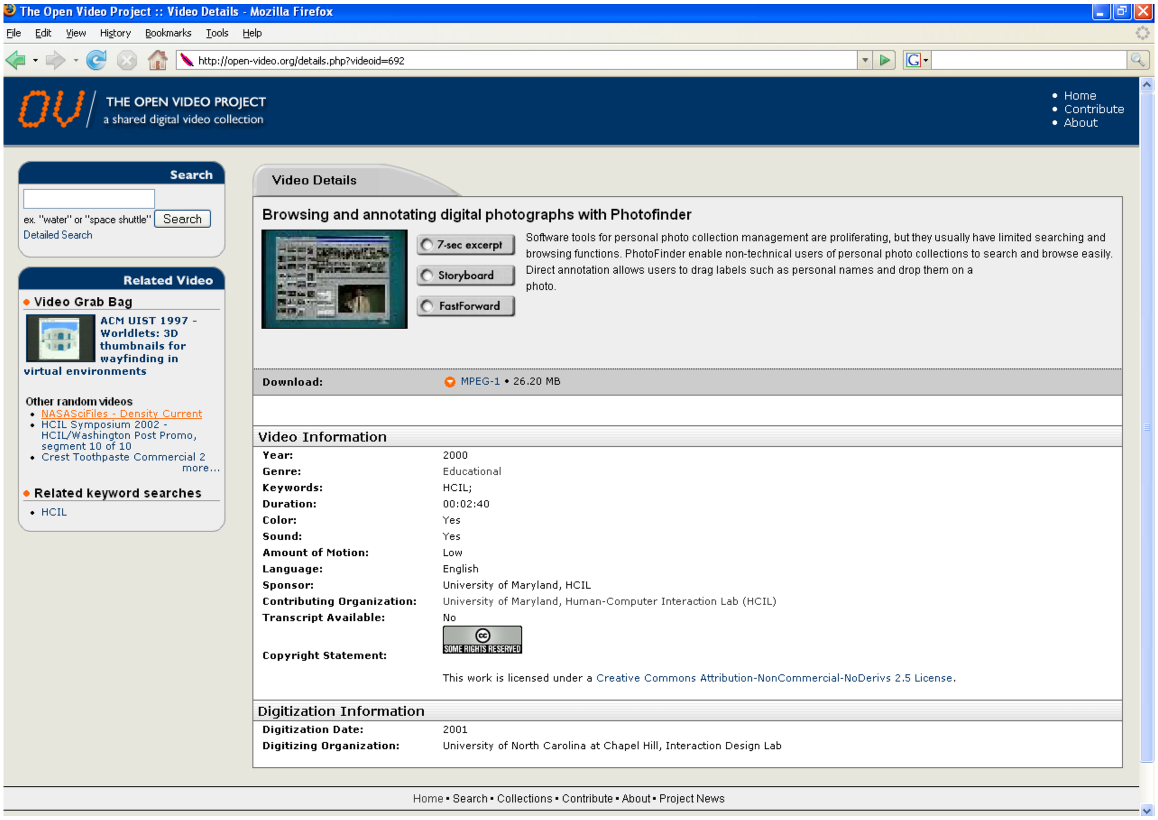

Browsing and annotating digital photographs with Photofinder

Software tools for personal photo collection management are proliferating, but they usually have limited searching and browsing functions. PhotoFinder enable non-technical users of personal photo collections to search and browse easily. Direct annotation allows users to drag labels such as personal names and drop them on a photo.

Download: MPEG-1 • 26.20 MB

Video Information

Year:	2000
Genre:	Educational
Keywords:	HCIL;
Duration:	00:02:40
Color:	Yes
Sound:	Yes
Amount of Motion:	Low
Language:	English
Sponsor:	University of Maryland, HCIL
Contributing Organization:	University of Maryland, Human-Computer Interaction Lab (HCIL)
Transcript Available:	No
Copyright Statement:	This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 2.5 License.

Digitization Information

Digitization Date:	2001
Digitizing Organization:	University of North Carolina at Chapel Hill, Interaction Design Lab

Figure 2. Open Video Screen Display showing three kinds of visual surrogate options.

We currently are working to incorporate audio surrogates into the search process. A recent study (Song & Marchionini, in press) was conducted to determine the tradeoffs between visual and audio information in video surrogates. A within subjects study with 36 participants was conducted that compared audio-only (spoken descriptions), visual-only (storyboards), and combined surrogates effects on five kinds of recognition and sense-making tasks (write gist, select pertinent keywords from lists, select best title, select pertinent keyframes from arrays, select best description from list). Dependent measures included performance accuracy on the five

---

[9] http://open-video.org

[10] Textual metadata is also provided for all videos and many videos also have short excerpts.

tasks, time to view surrogate, time to complete tasks and a suite of affective measures that included confidence, and judgments of usefulness, usability, engagement, and enjoyment.

As expected from the psychological literature on dual coding (Paivio, 1986) and learning effects of multimedia (Mayer & Gallini, 1990), the combined conditions were statistically reliably better on most performance tasks and preferred by participants. An important result found was that audio surrogates alone were almost as good as the combined surrogates on the performance task. Although the visual surrogates alone were significantly reliably faster to consume, there were no time penalties for audio and combined surrogates on task completion time. This study raised a here-to-for unasked question about synchronization of different information channels in multimedia surrogates. The evidence in favor of synchronized channels in the primary video is well established. Based on our observations that people were able to easily integrate the independent channels in the combined condition, even though they were not synchronized, it appears that multimedia surrogates need not synchronize information from different channels if it is clear to people that they are not coordinated. This work suggests that audio surrogates should be incorporated into video retrieval systems, that synchronizing different channels in surrogates may not be necessary, and that information seekers will be asked to control tradeoffs in time, satisfaction, and performance during results examination.

2.3.2. Example: Results in Context.

In this example we describe the use of content-rich search interfaces that extract and present the contents of the top-ranked retrieved documents, use them to promote exploration of the search results, and use this exploration as implicit feedback to support query refinement and retrieval strategy selection (White, Jose & Ruthven, 2005). In Figure 3 we show an example of a content-rich interface. Through applying sentence extraction techniques adopted from the summarization community, content-rich interfaces create a polyrepresentative search environment comprising multiple representations (or views) on each of the most highly-ranked Web documents.

As well as being represented by their full-text, documents are also represented by a number of smaller, query-relevant representations, created at retrieval time. These comprise the title (2)[11] and a query-biased summary of the document (3) (White, Jose & Ruthven, 2003) A list of sentences extracted from the top thirty documents retrieved scored in relation to the query, called *Top-Ranking Sentences* (TRS), include sentences from each document (1). Each sentence included in the top-ranking sentence list is a representation of the document, as is each sentence in the summary (4). Finally, for each summary sentence there is an associated sentence in the context it occurs in the document (i.e., with the preceding and following sentence from the full-text) (5).

---

[11] Numbers correspond to those in Figure 3.

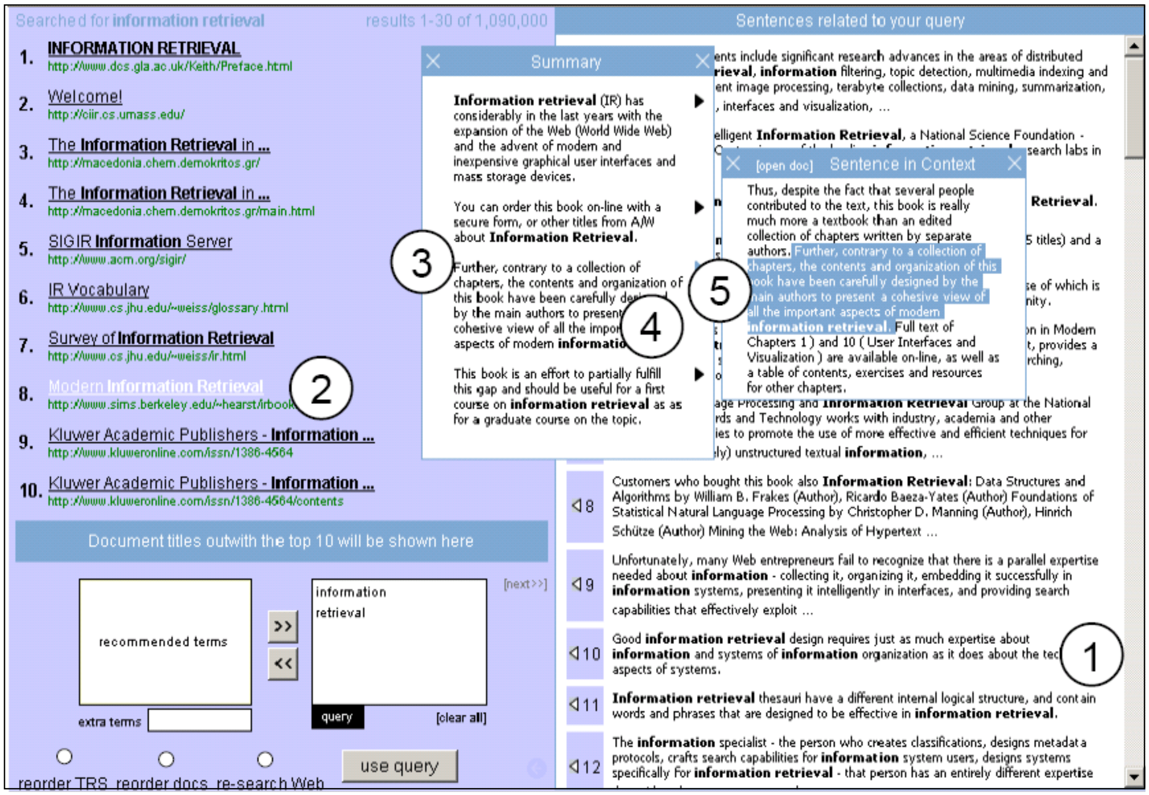

1. **INFORMATION RETRIEVAL**
http://www.dcs.gla.ac.uk/Keith/Preface.html
2. Welcome!
http://ciir.cs.umass.edu/
3. The **Information Retrieval** in ...
http://macedonia.chem.demokritos.gr/
4. The **Information Retrieval** in ...
http://macedonia.chem.demokritos.gr/main.html
5. SIGIR **Information** Server
http://www.acm.org/sigir/
6. IR Vocabulary
http://www.cs.jhu.edu/~weiss/glossary.html
7. Survey of **Information Retrieval**
http://www.cs.jhu.edu/~weiss/ir.html
8. Modern **Information Retrieval**
http://www.sims.berkeley.edu/~hearst/irbook
9. Kluwer Academic Publishers - **Information ...**
http://www.kluweronline.com/issn/1386-4564
10. Kluwer Academic Publishers - **Information ...**
http://www.kluweronline.com/issn/1386-4564/contents

Information retrieval (IR) has considerably in the last years with the expansion of the Web (World Wide Web) and the advent of modern and inexpensive graphical user interfaces and mass storage devices.

You can order this book on-line with a secure form, or other titles from A/W about **Information Retrieval**.

Further, contrary to a collection of chapters, the contents and organization of this book have been carefully designed by the main authors to present a cohesive view of all the important aspects of modern **information retrieval**.

This book is an effort to partially fulfill this gap and should be useful for a first course on **information retrieval** as as for a graduate course on the topic.

Thus, despite the fact that several people contributed to the text, this book is really much more a textbook than an edited collection of chapters written by separate authors. Further, contrary to a collection of chapters, the contents and organization of this book have been carefully designed by the main authors to present a cohesive view of all the important aspects of modern **information retrieval**. Full text of Chapters 1) and 10 (User Interfaces and Visualization) are available on-line, as well as a table of contents, exercises and resources for other chapters.

ents include significant research advances in the areas of distributed **rieval, information** filtering, topic detection, multimedia indexing and ent image processing, terabyte collections, data mining, summarization, , interfaces and visualization, ...

elligent **Information Retrieval**, a National Science Foundation - search labs in

age Processing and **Information Retrieval** Group at the National rds and Technology works with industry, academia and other ies to promote the use of more effective and efficient techniques for ly) unstructured textual **information**, ...

◁8 Customers who bought this book also **Information Retrieval**: Data Structures and Algorithms by William B. Frakes (Author), Ricardo Baeza-Yates (Author) Foundations of Statistical Natural Language Processing by Christopher D. Manning (Author), Hinrich Schütze (Author) Mining the Web: Analysis of Hypertext ...

◁9 Unfortunately, many Web entrepreneurs fail to recognize that there is a parallel expertise needed about **information** - collecting it, organizing it, embedding it successfully in **information** systems, presenting it intelligently in interfaces, and providing search capabilities that effectively exploit ...

◁10 Good **information retrieval** design requires just as much expertise about **information** and systems of **information** organization as it does about the tec aspects of systems.

◁11 **Information retrieval** thesauri have a different internal logical structure, and contain words and phrases that are designed to be effective in **information retrieval**.

◁12 The **information** specialist - the person who creates classifications, designs metadata protocols, crafts search capabilities for **information** system users, designs systems specifically for **information retrieval** - that person has an entirely different expertise

Figure 3. Content-rich search interface.

The document representations were arranged in interactive *relevance paths* (the order of which is denoted by the numbers in Figure 3), and encouraged interaction with the content of the retrieved document set. We call this approach *content-driven information seeking* (CDIS) since it is the content of the retrieved documents that drives the information-seeking process. This is in contrast to query-driven information-seeking, where searchers proactively seek information through the query they provide. Typically Web-search systems use lists of document surrogates to present their search results. This forces searchers to make two steps when assessing document relevance; first assess the surrogate, then perhaps peruse and assess the document (Paice, 1990). Such systems enforce a *pull* information seeking strategy, where searchers are proactive in locating potentially relevant information from within documents. In CDIS, it is the system that acts proactively, presenting the searcher with potentially relevant sentences taken from the document set at retrieval-time. The system uses a *push* approach, where potentially useful information is extracted from each document and proactively pushed to the searcher at the results interface. Searchers have to spend less time *locating* potentially useful information.

As the users explore the top-ranked search results through this interface, the system uses their interaction to make suggestions about additional query terms that may be appropriate to add to the original query, or retrieval strategies related to the estimated level of change in their information needs during the search session. Depending on the amount of divergence from the original request the system estimated, it would either take no action, recommend that the user reorder top-ranking sentences extracted from the top documents, reorder the top-ranked search results, or if the estimated change in need was sufficient, then re-search the Web.

We performed five user studies on variants of this interface, involving over 150 subjects over the course of three years. Each user study targeted a particular aspect of the interface, from the use of document representations to facilitate more effective information access (White, Ruthven & Jose, 2005), to different amounts of user control over aspects of the search process (relevance indication, query formulation, and action selection) (White & Ruthven, 2006). The findings of our research suggested that users found these content-rich interfaces useful for tasks that were exploratory in nature (i.e., where they needed to gather background information on a particular topic or gather sufficient information to enable them to make a decision about the best course of action). However, the interfaces were not as effective in known-item searches where users had to find a specific piece of information. In addition, users wanted to retain control over the strategic aspects of their search such as the decisions to conduct new searches, but were willing to delegate control for less severe interface actions to the system. A number of our studies compared this interface with the traditional interface offered by Google. The findings showed that searchers benefited from the additional information both in terms of subjective measures such as task success and more objective measures such as task completion time.

2.4. Problem reformulation.

The need to reformulate the problem, either as expressed or internally, is a common part of information seeking. The set of documents that are retrieved in response to a query often serves as feedback about the effectiveness of the query or the effectiveness of the system in interpreting the query. Deciding when and how to iterate requires an assessment of the information-seeking process itself, how it relates to accepting the problem, and the expected effort, and how well the extracted information maps onto the task (Marchionini, 1995).

Techniques such as Relevance Feedback (RF) (c.f. Salton & Buckley, 1990) have been proposed as a way in which IR systems can support the iterative development of a search query using examples of relevant information provided by the information seeker. RF is an effective technique in non-interactive experiments (Buckley *et al.*, 1994). However, few studies have investigated the use of RF (e.g., Koenemann & Belkin, 1996), and have highlighted problems in the use of RF by searchers at the interface. Typically RF systems require searchers to assess a number of documents at each feedback iteration. This activity includes the viewing of documents to assess their value and the marking of documents to indicate their relevance. There are a number of factors that can affect the use of RF in an interactive context. Relevance assessments are usually binary in nature (i.e., a document is either relevant or it is not) and no account is taken of partial relevance; where a document may not be completely relevant to the topic of the search or the searcher is uncertain about relevance. Previous studies have shown that the number of partially relevant documents in a retrieved set of documents is correlated with changes in the search topic or relevance criteria (Spink *et al.*,1998). Potentially relevant documents are therefore useful in driving the search forward or changing the scope of the search. The techniques used to represent the document at the interface are also important for the use of RF. Barry *et al.* (1998) demonstrated that the use of different document representations (e.g., title, abstract, full-text) can affect relevance assessments. The order in which relevance assessments are made also can affect searchers' feelings of satisfaction with the RF system (Tianmiyu & Ajiferuke, 1988).

RF is typically treated as a batch process where searchers provide feedback on the relevance of a number of documents and request support in query formulation. This may not be the best approach as in interactive environments searchers assess documents individually, not as a batch, and search is a sequential learning process (Bookstein, 1983). Incremental feedback (Aalbersberg, 1992) requires searchers to assess documents individually; they are asked about the relevance of a document before being shown the next document. Through this feedback process the query is iteratively modified. The method does not force searchers to use RF although it does force them to provide feedback and may hinder their abilities to make relative relevance assessments between documents (Florance & Marchionini, 1995). To resolve this problem, Campbell proposed an ostensive weighting technique (1999) that uses a "query-less" interface and browse paths between retrieved images to implicitly infer information needs. The paths followed through such *information spaces* are affected by the interests of the searcher. In Campbell's system, known as the *ostensive browser*, documents (images) are represented by nodes and the route traveled between documents by search paths. Clicking on a node is assumed to be an indication of relevance and the system performs an iteration of RF using the node clicked and all objects in the path followed to reach that node. The top-ranked images are presented at the interface and the searcher can select one of those shown, or return to a path followed previously. There is an implicit assumption that when choosing one image that this image is more relevant than the alternatives.

The process of retrieving relevant information is rich and complex. Bates (1990) suggested that there are situations where searchers may wish to control their own search and there are situations where they would like to make use of IR systems to automate parts of their search. As suggested by Fowkes and Beaulieu (2000) the level of interface support can be varied based on search complexity and associated cognitive load. Related empirical studies (e.g., Ellis, 1989) have shown that searchers are actively interested in their search and are keen to feel in control over what information is included or excluded and why. Other interaction metaphors (such as Rodden's use of a bookshelf to represent the current search context) have also been used to help searchers use RF systems (1998).

Web search systems such as Google offer RF by providing searchers with the opportunity to request "Similar Pages" and retrieve related documents. Jansen *et al.* (2000) showed that RF on the Web is used around half as much as in traditional IR searches. Therefore, the design of RF techniques for the Web needs to be more carefully approached than in other document domains as the searchers who use them are typically untrained in how to use search systems that implement them.

Systems such as Kartoo[12], the Hyperindex Browser (Bruza *et al.*, 2000), Paraphrase (Anick & Tipirneni, 1999) and Prisma (Anick, 2003) have all tried to incorporate feedback and term suggestion mechanisms into interactive Web search. Vivisimo[13] uses clustering technology to recommend additional query terms. These systems assume that Web searchers are mainly concerned with maximizing relevant results on the first page (Spink *et al.*, 2002), and rely on searchers to select the most appropriate terms (selected from the highest-ranked documents) to express their needs. These approaches typically assume top-ranked documents are relevant (i.e., use pseudo-relevance feedback) and give searchers control over which terms are added to the query. If the initial query is poorly conceived, irrelevant documents may be highly ranked, leading to erroneous term suggestions.

Interaction with feedback systems has an associated cost in terms of time and effort expended. Reading and rating a large number of documents is a costly activity that is not always justified by the results obtained. To be truly useful, searcher-system dialogue must have a perceived benefit to the searcher since they may depend on it directly. If this benefit cannot be guaranteed then feedback approaches based on passive observational evidence may be more appropriate since searchers have no pre-conceived expectations of their performance. Implicit RF (Kelly &

---

[12] http://www.kartoo.com

[13] http://www.vivisimo.com

Teevan, 2003) gathers relevance information unobtrusively from searchers' interaction, but with a reduced burden on them to provide relevance judgments.

One way RF can help is by suggesting additional *query expansion* terms for query modification (c.f. Efthimiadis, 1996). This modification can occur interactively with searcher participation i.e., interactive query expansion, or automatically without searcher involvement i.e., automatic query expansion. It is clear that the dynamism and action-oriented nature of the information-seeking process suggests that the user should be involved at all stages. Previous research in this area has shown that transparent query expansion interfaces (where system functionality is visible) are much preferred to opaque interfaces (where system functionality is hidden) (Koenneman & Belkin, 1996). Shneiderman and colleagues have advocated for user control and their active involvement in system activities (Shneiderman & Plaisant, 2004).

Today's search engines use the "wisdom of crowds" to suggest documents that may be worth investigating further based on the interaction decisions of many users. It is also possible to use the query formulation behavior of a large number of users to suggest query reformulations that others have entered (Anick, 2003). In a similar way, "search signposts" (White, Bilenko & Cucerzan, 2007) direct users to popular destinations that others have ended up following the submission of a query and subsequent traversal of a browse path. These are potentially useful approaches, but the most popular options for queries and documents may not always be the best options. Depending on the nature of the task being undertaken, users may want queries or documents that will provide them with new insights, unique perspectives, or have been visited or created by those with specialist domain knowledge. The challenge lies in being able to extract these queries and locations, given that they reside somewhere in the tail, and are not easily differentiable from non-relevant items. Sites such as StumbleUpon[14] use "collaborative opinions" from millions of Web surfers to help users discover new web pages that they probably would not discover through a search engine. Bringing relevant and previously unsurfaced documents to the attention of the information seeker will undoubtedly improve their ability to complete their tasks more effectively (and refine their problems if appropriate).

The provision of scratchpads or temporary bookmarks allows users to store information items as they are encountered during their search, and return to these later in the process to examine the contents or perhaps use them in query refinement. The Google Notes feature described in the previous section allows users to store documents and notes pertaining to documents as notes during their search, and the Scratchpad feature in Windows Live Image Search allows users to drag-and-drop images to form a collection during image browsing. However, once stored, these systems provide limited functionality on how to use the stored items. Possibilities include the use of them as RF to perform a new retrieval, visit multiple stored items simultaneously, publish these on Web or in a word processing document, or share these with other searchers. As we will describe in Section 2.5, information use is a vital part of the information-seeking process that is too often ignored. White, Song and Liu (2006) described a two-dimensional workspace used to support oral history search using concept maps created by middle-school teachers. During their search users can drag entities such as people and locations onto this map, form relationships between these entities, and use the concept map that emerges as the basis for conducting a new search. As an additional feature, the workspace also included functionality to create a movie presentation for their students automatically based on the concept map. This is an example of how the workspaces can be used to support the refinement of searches, but also the use of information following this refinement.

---

[14] http://www.stumbleupon.com

Earlier in this article we discussed "mediated searching" as a means through which a user's interaction with a restricted collection can help them create better representations of information needs before they interact with a heterogeneous collection. In situations where problem descriptions need to be reformulated, users also may benefit from the consultation of new sources of information or engaging in iterative dialog with systems (or users) with specific domain knowledge. Systems that use unobtrusive methods to infer interests are called attentive or adaptive systems. These observe the user (via their interaction), model the user (based on this interaction), and anticipate the user (based on the model they develop). Attentive *information* systems aim to support user's information needs and construct a model based on their interaction. In attentive systems, the responsibility for monitoring this interaction is usually assigned to an external *agent* or *assistant*. Examples of such agents include Lira (Balabanovic & Shoham, 1995), WebWatcher (Armstrong *et al.*, 1995), Suitor (Maglio *et al.*, 2000), Watson (Budzik and Hammond, 2000), PowerScout (Lieberman *et al.*, 2001), and Letizia (Lieberman, 1995).

Attentive systems accompany the user during their information seeking journey, and by observing search behavior (and other behaviors in inter-modal systems) they can model user interests. Such systems can typically operate on a restricted document domain or on the Web. The methods used to capture this interest and present system suggestions differ from system to system. Letizia (Lieberman, 1995), for example, learns user's current interests and by doing a lookahead search (i.e., predicting what searchers may be interested in the future, based on inference history) can recommend nearby pages. PowerScout (Lieberman *et al.*, 2001) uses a model of user interests to construct a new complex query and search the Web for documents semantically similar to the last relevant document. WebWatcher (Armstrong *et al.*, 1995), in a similar way, accompanies users as they browse, but as well as observing, WebWatcher also acts as a *learning apprentice* (Mitchell *et al.*, 1994). Over time the system learns to acquire greater expertise for the parts of the Web that it has visited in the past, and for the topics in which previous visitors have had an interest. Suitor (Maglio *et al.*, 2000), tracks computer users through multiple channels – gaze, Web browsing, application focus – to determine their interests. Watson (Budzik & Hammond, 2000), uses contextual information, in the form of text in the active document, and uses this information to proactively retrieve documents from distributed information repositories by devising a new query.

All of these systems can be classified as *behavior-based* interface agents (Maes, 1994), that develop and enhance their knowledge of the current domain incrementally from inferences made about user interaction. These systems work with the user's searching/browsing in a concurrent manner, finding and presenting documents to them during the search based on system inference of relevance/current interest. To predict what might be useful, an attentive information system must learn from a user's history of activity to improve both the relevance and timeliness of its suggestions. Attentive systems are personalized, developing and revising a user model throughout the whole search session. As the user model evolves, becoming a closer approximation to the user after each step, it should be able to recommend new documents should a significant change in need and/or user dissatisfaction be detected. Any new suggestions should be presented to users in an unobtrusive and timely way, either selecting opportune moments of prolonged inactivity or in the periphery of the current, active task. These concepts are embodied by systems with a *just-in-time* (JIT) information infrastructure, where information is brought to users just as they need it, without requiring explicit requests (Budzik & Hammond, 2000). Such systems automatically search information repositories on the user's behalf, as well as providing an explicit, query-entry interface. Attentive information systems can be distinguished by a few main characteristics. They are capable of gathering information on user behavior from a number of sources, even across multiple modalities. When only a single source is used, the probability of making incorrect inference of user intentions is high. In contrast, with multiple sources of evidence (e.g., many applications open concurrently) ambiguity can be removed and a more accurate user model can be constructed. Despite the potential effectiveness of such agents, to insure user satisfaction it is important that they provide levers and buttons through which their internal mechanisms can be controlled. In particular it should be users who initiate actions such as new searches, monitor search progress, and decide the order in which actions occur (Shneiderman, Byrd & Croft, 1997).

There have been attempts to create a medium of knowledge elicitation traditionally performed by human intermediaries. From this user models can be created that can be used to select retrieval strategies (Rich, 1983; Croft & Thompson, 1987; Brajnik *et al.*, 1996). Systems of this nature have focused on characterizing tasks, topic knowledge and document preferences to predict searcher responses, goals and search strategies. These systems typically make many assumptions about the search environment in which they operate and the searchers that use them. Search systems such as Grundy (Rich, 1983) tried to infer user preferences by characterizing search behavior, whereas systems such as FIRE (Brajnik *et al.*, 1996) have attempted to individuate the user modeling process. Systems like I^3R (Croft & Thompson, 1987) used different methods to improve query formulation and select appropriate retrieval strategies. I^3R used multiple retrieval techniques to form a better model of the searcher's information needs. Models were constructed in I^3R based on RF about what terms and concepts were of interest to searchers. This system required searchers to perform an active part in explicitly defining the model and their interests before using the system. This made users more in control of the system, and prevented the system's model of relevance from deviating too greatly from the searcher's (correct) model. Problem reformulation typically occurs because of unsatisfactory search results or a change in the knowledge state of the searcher during the search. Since it is the user that determines task completion, problem reformulation activities need to involve them as an active participant throughout. Visualizations can be used to provide information on the overlap between result sets and opportune areas of the information space yet to be explored. Scratchpads and temporary bookmarks that allow important facts and documents to be stored during the process should be offered. However, providing users with only the information they have stored may be insufficient to move their problem closer to resolution. Searchers should be able to use the information stored to find related information, allow users to cluster and form relationships between stored information, and support the exploration of previously uncharted regions of the information space using the experiences and opinions of others as a guide. Since users can only attend to a small number of items at any point in time, they should also be supported by systems that provide recommendations as the users search. These activities should attempt to maximize the novelty of information they provide by searching a broad range of locations that have not yet been visited by the searcher. However, such background system activities should only be brought into the foreground at the request of searchers, who should have ultimate control over system operations.

2.5. Information use.

Use depends on the information seeker understanding the results of search and making a decision that information is relevant, trustworthy, and as complete as necessary to meet the conditions of the information problem. Understanding results is dependent on a variety of searcher cognitive characteristics (e.g., knowledge about the search domain, inferential ability) and states (e.g., attentiveness), however, good system design can augment or amplify searchers' capabilities to understand results. Many of the query reformulation strategies discussed above aid understanding of intermediate results as search progresses. A collaboration of the first author and the HCIL aimed to help people find and understand government statistics in the WWW (Marchionini et al., 2006). This work focused on improving the vocabulary of websites (e.g., Hass et al., 2003), on-demand help (e.g., Plaisant et al., 2003), and exploratory search interfaces (e.g., Kules & Shneiderman, 2006; Zhang & Marchionini, 2004) and results were applied to websites at several US government websites.

The decision to stop searching and use some of the results is sometimes straightforward (e.g., a known item result that clearly answers a specific question) but in most cases the decision is satisfactory rather than optimal and searchers accept results that are “good enough” (e.g., most exploratory search problems). In either case, basic functionality at the operating system or application software levels like cut and paste and import/export are basic supports for taking results and incorporating them into other electronic documents. Many digital libraries provide bibliographic reference alternatives to make citations easy and also more sophisticated citation links that make finding related literature a simple mouseclick option. Likewise, browsers and some search systems provide tools for harvesting text, images, data, and other search results directly into work documents at hand.

The earlier discussion about the work environment and search environment converging at the network-connected desktop (or mobile device) motivates efforts to more fully unite information search and use of the retrieved information resources. Especially in cases where satisfactory rather than optimal results are found, designers can add extensions that continue search in the background and report back updates, build search histories that incorporate the products of use, and integrate the search facilities into general work applications.

3. Discussion.

Information seeking is a pervasive human activity that continues to gain importance in a massively connected world of digital information. Advances in hardware and networking have driven enormous progress in algorithms that leverage the amount of information in databases, the WWW, and the digital social interactions that hundreds of millions of people have each day. System designers have also made great progress by adopting human-centered approaches to design—conducting user needs assessments, adopting design guidelines rooted in human needs and universal access, doing both formative and summative user studies, and listening to feedback as people work with their systems. The confluence of these forces brings us to a renaissance in search research and development. Each day brings novel plug-ins or applications that add to our abilities to find, understand, and manage information in cyberspace. These advances also bring increasingly high expectations on the part of users that will continue to drive research and development. Three themes run throughout these innovations: interaction, representation, and integration.

User studies and reflection on systems that have succeeded in the marketplace demonstrate that people want to be in control over activities that are important to them, while happily acquiescing control over routine activities to systems that are trusted. The give and take between conscious control and automation is best managed by interactive systems that give people easy to exercise and change choices. Interactive systems have the added value of engaging attention and this combination makes interaction a desirable design goal. Inspired by the many dynamic query user interfaces at HCIL (e.g., Ahlberg *et al*.1992,; Kumar *et al*., 1997; Plaisant *et al*., 1998; Shneiderman *et al*., 2000; Williamson & Shneiderman, 1992), Marchionini and his colleagues defined an “AgileViews” design framework that aims to link multiple, rich representations with agile control mechanisms to integrate the query-results-reformulation cycle (Marchionini *et al.*, 2000). Five kinds of representations (views) were defined: overviews of information spaces; previews of information objects; reviews of past actions and results (histories); peripheral views of contextual information related to the view in active focus; and shared views that include collaborative or incidental representations of other people. Easy to manipulate control mechanisms such as hovering and brushing are mapped to actions such as quick collection partitioning, zooming and panning, and shifting focus across different views. This framework was empirically evaluated in Geisler's dissertation (Geisler, 2003) with a set of instances for video retrieval and provides a theoretical framework for design desiderata.

The trend in search has moved relentlessly toward richer digital representations. From the terse card catalogs and bibliographic databases that yielded pointers to documents, full text systems emerged and it is hard to find students who do not expect instant access to full text documents at their desktops. Beyond full text, today's systems include a variety of data in different forms ranging from genomic sequences to music to video to geospatial data to computer code. Consider systems such as the National Center for Biotechnology Information[15] that supports searching across data as diverse as genomic databases and bibliographic databases, or the Library of Congress that supports searching across laws, books, videos, sound recordings, manuscripts, and more from the same site. Even more challenging, multimedia combinations of these rich representations are increasingly common as mashups of retrieved data sets are integrated (e.g., statistical data retrieved and mapped onto real-time spatial displays). These rich representations challenge designers to help searchers distinguish different information forms as well as topics. Questions arise such as what level of aggregation to display in response to a query? How should user queries be disambiguated from a data type perspective? How might queries with multiple data types weight these different types? The most challenging issues emanate from query specification—how to support non-textual queries? Although there are examples of systems (e.g., hum a few bars to retrieve music; sketch a figure to retrieve an image), the state of the art is to either provide query-by-example interfaces or expect searchers to enter text.

Finally, there is a blurring between the search activities within the information seeking process. As we illustrated in the examples and discussions above, highly interactive systems closely couple search expression, results examination, and reformulation – and trends look toward even more integration in the years ahead. This integration is positive overall from a user perspective, but can lead to heavy-weight search systems. Alternatively, there is increased integration of search system capabilities into applications and operating systems. Today we have email applications with built-in search as well as cross application search tools that work across applications. It is likely that more search support will be built into all applications while specialized search systems will emerge with advanced capabilities and alternatives for integrating information seeking processes into daily workflows.

4. Conclusion

.

The HCI and IR communities have played a pivotal role in the emergence of search as an enabling technology for many computer users. Ben Shneiderman has strongly influenced this work by postulating clear principles about user control and building and evaluating a variety of user interfaces that illustrate these principles. In this article we have used an information-seeking framework to demonstrate the importance of this synergy in areas such as problem formulation and expression, result examination, and information use. Technological advances and the increased involvement of the user in information-seeking have made it easier for users to find what they need in the well-defined cases. Although some progress has also been made in helping people make sense of (understand) what is found, there are many opportunities to expand work in this area further. Incorporating search into frequently used applications, such as Web browsers, IM software, and office applications, can begin to realize our vision of a fluid and productive search experience. Multiple perspectives on search results and information spaces, and richer representations of documents and queries can facilitate more extensive exploration and the resolution of more complex information problems. Many of the systems we have described in

---

[15] http://www.ncbi.nlm.nih.gov/

this article have focused on making users more *involved* in information-seeking activities; it is vital that this continues. However, search systems of the future will also focus on making users more *informed* by providing explicit support for learning and investigation within a wider work task context.

5. References

Aalbersberg, I.J. (1992). Incremental relevance feedback. In *Proceedings of the 15th Annual ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 11-22.

Ackerman, M. (1998). Augmenting organizational memory: a field study of Answer Garden. *ACM Transactions on Information Systems,* 16(3), 203-224.

*Ahlberg, C., Williamson, C., and Shneiderman, B. (1992). Dynamic queries for information exploration: An implementation and evaluation. In *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems*, pp. 619-626.

Agichtein, E., Brill, E. and Dumais, S.T. (2006). Improving web search ranking by incorporating user behavior information. In *Proceedings of the 29th Annual ACM SIGIR Conference on Research and Development in Information Retrieval,* pp. 19-26.

Anick, P. and Tipirneni, S. (1999). The paraphrase search assistant: Terminological feedback for iterative information seeking. In *Proceedings of the 22nd Annual ACM SIGIR Conference on Research and Development in Information Retrieval,* pp. 153-159.

Anick, P. (2003). Using terminological feedback for web search refinement: a log-based study. In *Proceedings of the 26th Annual ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 88-95.

Armstrong, R., Freitag, D., Joachims, T. and Mitchell, T. (1995). WebWatcher: A learning apprentice for the world wide web. In *Proceedings of the AAAI Spring Symposium on Information Gathering from Heterogeneous, Distributed Environments,* pp. 6-12.

Balabanovic, M. and Shoham, Y. (1995). Learning information retrieval agents: Experiments with automated web browsing. In *Proceedings of the AAAI Spring Symposium on Information Gathering from Heterogeneous, Distributed Environments,* pp. 13-18.Bates, M. (1979) Information search tactics. *Journal of the American Society for Information Science*, 30: 205-213.

Bates, M. J. (1990). Where should the person stop and the information search interface start? *Information Processing and Management*, 25(5): 575-591.

Beaulieu, M. (1997). Experiments on interfaces to support query expansion. *Journal of Documentation*, 53(1): 8-19.

Bookstein, A. (1983). Information retrieval: A sequential learning process. *Journal of the American Society for Information Science,* 34(5): 331-345.

*Botafogo, R. and Shneiderman, B. (1991). Identifying aggregates in hypertext structures. In *Proceedings of Hypertext '91*, pp. 63-74.

Brajnik, G., Mizzaro, S. and Tasso, C. (1996). Evaluating user interfaces to information retrieval systems: A case study of user support. In *Proceedings of the 19th Annual ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 128-136.

Bruza, P., McArthur, R. and Dennis, S. (2000). Interactive internet search: Keyword, directory and query reformulation mechanisms compared. In *Proceedings of the 23rd Annual ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 280-287.

Buckley, C., Salton, G. and Allan, J. (1994). The effect of adding relevance information in a relevance feedback environment. In *Proceedings of the 17th Annual ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 292-300.

Bush, V. (1949). As we may think. *Atlantic Monthly*, 176, (July 1945), 101-108.

*Card, S., Mackinlay, J. and Shneiderman, B. (1999). Readings in information visualization: Using vision to think. San Francisco: Morgan Kaufmann.

Carson, C., Belongie, S., Greenspan, H., and Malik, J. (1997). Region-based image querying. In *Proceedings of IEEE Workshop on Content-based Access of Image and Video Libraries*, pp. 42-49.

Christel, A., Hauptmann, A. Warmack, A.S. , and Crosby, S. (1999). Adjustable filmstrips and skims as abstractions for a digital video library. In *Proceedings of the IEEE Advances in Digital Libraries Conference*, 98-104

Christel, M. Smith, C.R. Taylor, and D. Winkler, Evolving Video Skims into Useful Multimedia Abstractions. In *Proc. CHI '98, ACM* (1998), pp. 171-178.

Cutrell, E., Robbins, D.C., Dumais, S.T. and Sarin, R. (2006). Fast, flexible filtering with Phlat - Personal search and organization made easy. In *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems*, pp. 261-270

Chi, E.H. and Pirolli, P.L. (2006). Social information foraging and collaborative search. In *Proceedings of Human-Computer Interaction International Workshop*.

Chi, E.H., Pirolli, P., and Pitkow, J. (2000). The scent of a site: a system for analyzing and predicting information scent, usage, and usability of a Web site. In *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems*, pp. 161-168.

Croft, W. B. and Thompson, R. H. (1987). I^3R: A new approach to the design of document retrieval systems. *Journal of the American Society for Information Science*, 38 (6): 389-404.

Dumais, S.T., Cutrell, E., Cadiz, J.J., Jancke, G., Sarin, R. and Robbins, D.C. (2003). Stuff I've Seen: A system for personal information retrieval and re-use. In *Proceedings of the 26th ACM SIGIR Conference on Research and Development on Information Retrieval*, pp. 72-79.

Egan, D., Remde, J., Gomez, L., Landauer, T., Eberhardt, J., and Lochbaum, C. (1989). Formative design evaluation of superbook. *ACM Transactions on Information Systems*, 7(1): 30-57.

Efthimiadis, E.N. (1996). Query expansion. *Annual Review of Information Systems and Technology*, 31: 121-187.

Ellis, D. (1989). A behavioural approach to information retrieval system design. *Journal of Documentation*, 45, 171-212.

Flickner, M., Sawhney, H.S., Ashley, J., Huang, Q., Dom, B., Gorkani, M., Hafner, J., Lee, D., Petkovic, D., Steele, D., Yanker, P. (1995). Query by image and video content: The QBIC system. *IEEE Computer*, 28(9): 23-32.

Florance, V. and Marchionini, G. (1995). Information processing in the context of medical care. In *Proceedings of the 18th Annual ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 158-163.

Foote, J. (2000). Automatic Audio Segmentation using a Measure of Audio Novelty. In *Proceedings of IEEE International Conference on Multimedia and Expo*, pp. 452-455.

Fowkes, H. and Beaulieu, M. (2000). Interactive searching behaviour: Okapi experiment for TREC-8. *Proceedings of the 22nd BCS-IRSG European Colloquium on IR Research*.

Fox, E. & France, R. (1987). Architecture of an expert system for composite document analysis, representation and retrieval. *International Journal of Approximate Reasoning*, 1(2), 151-175.

Furnas, G. (1985). Experience with an adaptive indexing scheme. In *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems*, pp. 131-135.

Furnas, G.W., Landauer, T.K., Gomez, L. M. and Dumais, S. T. (1987). The vocabulary problem in human-system communication. *Communications of the ACM*, 30(11): 964-971.

Geisler, G. (2003). *AgileViews: A Framework for Creating More Effective Information Seeking Interfaces.* Unpublished doctoral dissertation, University of North Carolina, Chapel Hill. Available online at: http://www.ischool.utexas.edu/~geisler/info/geisler_dissertation.pdf

*Greene, S., Marchionini, G., Plaisant, C. and Shneiderman, B. (2000). Previews and overviews in digital libraries: Designing surrogates to support visual information seeking. *Journal of the American Society for Information Science*, 51(4): 380-393.

Haas, S.W., Pattuelli, M.C., and Brown, R.T. (2003). Understanding Statistical Concepts and Terms in Context: The GovStat Ontology and the Statistical Interactive Glossary. *Proceedings of the 66th Annual Meeting of the American Society of Information Science and Technology*, 193-199.

Hearst, M.A. (1995). TileBars: Visualization of term distribution information in full text information access. In *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems*, pp. 59-66.

Hearst, M (2006). Clustering versus faceted categories for information exploration, *Communications of the ACM*, 49(4): 59-61.

Herlocker, J., Konstan, J., Terveen, L., and Riedl, J. (2004). Evaluating collaborative filtering recommender systems. *ACM Transactions on Information Systems 22*(1): 5-53.

Hill, W.C., Hollan, J.D., Wroblewski, D. and McCandless, T. (1992). Edit wear and read wear. In *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems*, pp. 3-9.

Horvitz, E., Kadie, C.M., Paek, T. and Hovel, D. (2003). Models of Attention in Computing and Communication: From Principles to Applications. *Communications of the ACM*, 46(3): 52-59.

Horvitz, E., Jacobs, A. and Hovel, D. (1999). Attention-sensitive alerting. In *Proceedings of UAI '99, Conference on Uncertainty and Artificial Intelligence*, pp. 305-313.

Ingwersen, P. and Järvelin, K. (2005). The turn: Integration of information seeking and retrieval in context. New York: Springer-Verlag.

Jain, A. and Vailaya, A. (1996). Image retrieval using color and shape. *Pattern Recognition*, 29(8): 1233-1244.

Joachims, T. (2002). Optimizing search engines using clickthrough data. In *Proceedings of the 8th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 133-142.

Joachims, T., Freitag, D. and Mitchell, T. (1997). WebWatcher: A tour guide for the World Wide Web. In *Proceedings of the 16th Joint International Conference on Artificial Intelligence*, pp. 770-775.

Jansen, B. J., Spink, A. and Saracevic, T. (2000). Real life, real users, and real needs: A study and analysis of user queries on the web. *Information Processing and Management*, 36 (2): 207-227.

Kelly, D. and Teevan, J. (2003). Implicit feedback for inferring user preference. *SIGIR Forum*, 37 (2): 18-28.

Kerne, A., Koh, E., Dworaczyk, B., Mistrot, J., Choi, H., Smith, S., Graeber, R., Caruso, D., Webb, A., Hill, R. and Albea, J. (2006). combinFormation: a mixed-initiative system for representing collections as compositions of image and text surrogates. In *Proceedings of 6th ACM/IEEE Joint conference on Digital Libraries*, pp. 11-20.

Kolmlodi, A., Marchionini, G. and Soergel, D. (2006). Search history support for finding information: User interface design recommendations from a user study. *Information Processing and Management*, 43(1): 10-29.

*Koved, L. and Shneiderman, B. (1986). Embedded menus: Selecting items in context. *Communications of the ACM*, 29 (4): 312-318.

Koenemann, J. and Belkin, N. J. (1996). A case for interaction: A study of interactive information retrieval behavior and effectiveness. In *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems*, pp. 205-212.

Kuhlthau. C. (1991). Inside the Search Process: Information Seeking from the User's Perspective. Journal of the American Society for Information Science 42 (5) 1991, 361-371

Kules, B., & Shneiderman, B. (2003). Designing a metadata-driven visual information browser for federal statistics. In *Proceedings of the 2003 National Conference on Digital Government Research* (pp. 117-122).

*Kumar, H., Plaisant, C., & Shneiderman, B. (1997). Browsing hierarchical data with multi-level dynamic queries and pruning. *International Journal of Human-Computer Studies*, 46,103-124

Lieberman, H. (1995). Letizia: An agent that assists web browsing. *Proceedings of the 14th International Joint Conference on Artificial Intelligence*, pp. 475-480.

Lieberman, H., Fry, C. and Weitzman, L. (2001). Exploring the web with reconnaissance agents. *Communications of the ACM*, 44(7): 69-75.

Maekawa, T., Hara, T., and Nishio, S. (2006). A collaborative web browsing system for multiple mobile users. In *Proceedings of IEEE International Conference on Pervasive Computing and Communications*, pp. 22-35.

Maes, P. (1994). Agents that reduce work and information overload. *Communications of the ACM*, 37(7): 30-40.

Maglio, P. P., Barrett, R., Campbell, C. S. and Selker, T. (2000). SUITOR: An attentive information system. In *Proceedings of the Annual Conference on Intelligent User Interfaces*, pp. 169-176.

Marchionini, G. (1992). Interfaces for end-user information seeking. *Journal of the American Society for Information Science*, 29(3): 165-176.

Marchionini, G. (2006). Exploratory search: From finding to understanding. *Communications of the ACM*, 49(4), p. 41-46.

Marchionini, G. (2006). Toward human-computer information retrieval. *Bulletin of the American Society for Information Science and Technology*. June/July.

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

*Marchionini, G. and Shneiderman (1988). Finding facts vs. browsing knowledge in hypertext systems. *IEEE Computer*, 21(1): 70-80.

Marchionini, G. and Brunk, B. (2003). Towards a general relation browser: A GUI for information architects. *Journal of Digital Information*, 4(1).

Marchionini, G., Geisler, G., & Brunk, B. (2000). Agileviews: A Human-centered framework for interfaces to information spaces. In *Proceedings of the Annual Meeting of the American Society for Information Science*, pp. 271-280.

*Marchionini, G., Haas, S., Plaisant, C., & Shneiderman, B. (2006). Integrating Data and Interfaces to Enhance Understanding of Government Statistics: Toward the National

Statistical Knowledge Network Project Briefing. *Proceedings of 7th Annual International Conference on Digital Libraries (DG06)*. San Diego, CA. May 21-24, 2006. ACM Press. 334-5.

Marchionini, G., Wildemuth, B., & Geisler, G. (2006). The Open Video Digital Library: A Mobius strip of theory and practice. *Journal of the American Society for Information Science and Technology*, 57(2): 1629-43

Mitchell, T., Caruana, R., Freitag, D., McDermott, J. and Zabowski, D. (1994). Experience with a learning personal assistant. *Communications of the ACM*, 37(7): 81-91.

Morris, M.R. (2007). Cooperative interfaces for exploratory web search: Motivating multi-user search UIs. In *Proceedings of the ACM SIGCHI Workshop on Exploratory Search Interaction* (in press).

Morris, M. R., Paepcke, A., Winograd, T. and Stamberger, J. (2006). TeamTag: Exploring centralized versus replicated controls for co-located tabletop groupware. In *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems*, pp. 97-104.

Muresan, G. and Harper, D.J. (2002). Topic modeling for mediated access to very large document collections. *Journal of the American Society for Information Science and Technology*, 55(10): 892-910.

Mayer, R., and Gallini, J. (1990). When is an illustration worth ten thousand words? *Journal of Educational Psychology, 82*(4), 715-726.

Mayer, R., & Moreno, R.Paice, C. D. (1990). Constructing literature abstracts by computer: Techniques and prospects. A split-attention effect in multimedia learning: Evidence for dual processing systems in working memory. *Journal of Educational Psychology, 90*(2), 312-320.

Nelson, T. (1983). *Literary machines*. Mindful Press, Distributed by Eastgate Systems, Watertown, MA.

Norman, D.A. (2002). Emotion and design: Attractive things work better. *Interactions Magazine*, 9(4): 36-42.

Norman, K. and Chin, J. (1988). The effect of tree structure on search in a hierarchical menu selection system. *Behaviour and Information Technology*, 7(11): 51-65.

Pavio, A. (1986). *Mental representations: A dual coding approach*. Oxford: Oxford U. Press.

Pirolli, P. and Card, S. (1999). *Information* foraging theory. *Psychological Review, 106*(4), 643-675.

Plaisant, C., Marchionini, G., Bruns, T., Komlodi, A., and Campbell, L. (1997). Bringing treasures to the surface: Iterative design for the Library of Congress National Digital Library Program. In *Proceedings of the ACM SIGCHI Conference on Research and Development in Information Retrieval*, pp. 518-525.

*Plaisant, C., Kang, H. and Shneiderman, B. (2003). Helping users get started with visual interfaces: multi-layered interfaces, integrated initial guidance and video demonstrations. In *Proceedings of 10th International Conference on Human-Computer Interaction*, pp. 790-794.

*Plaisant, C.,Shneiderman, B., Mushlin, R. (1998). An information architecture to support the visualization of personal histories. *Information Management and Processing*, **34** (5), 581-597

Qian, R., Haering, N. and Sezan, M.I. (2002). A computational approach to semantic event detection in video. In A. Bovik, C. Chen, and G. Dmitry (Eds.), *Advances in Image Processing and Understanding: A Festschrift for Thomas S. Huang*. Series in Machine Perception and Artificial Intelligence, NJ: World Scientific, 199-235.*Management*, 26 (1), 171-186.

Rich, E. (1983). Users are individuals: Individualizing user models. *International Journal of Human-Computer Studies*, 51, 323-338.

Rodden, K. (1998). About 23 million documents match your query... In *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems (Doctoral Consortium)*, pp. 64-65.

Senior, A. (1999). Recognizing faces in broadcast video. In *Proceedings of the International Workshop on Recognition, Analysis, and Tracking of Faces and Gestures in Real-Time Systems*, pp. 105-110.

Salton, G. and Buckley, C. (1990). Improving retrieval performance by relevance feedback. *Journal of the American Society for Information Science*, 41(4): 288-297.

Sihvonen, A. and Vakkari, P. (2004). Subject knowledge, thesaurus-assisted query expansion and search success. In *Proceedings of the RIAO Conference*, pp. 393-404.

*Shneiderman, B. (2002). *Leonardo's Laptop: Human needs and the new computing technologies*. Cambridge, MA: MIT Press.

*Shneiderman, B. (1983). Direct manipulation: A step beyond programming languages, *IEEE Computer* 16(8): 57-69.

*Shneiderman, B. (1994). Dynamic Queries for Visual Information Seeking, *IEEE Software*, 11(6): 70-77.

*Shneiderman, B. and Plaisant C. (2005). *Designing the user interface (Fourth Edition).* Reading, MA: Addison-Wesley.

*Shneiderman, B., Byrd, D. and Croft, W.B. (1997). Clarifying search: A user-Interface framework for text searches. *D-Lib Magazine* (January).

*Shneiderman, B., Byrd, D. and Croft, W.B. (1998). Sorting out search: A user-interface framework for text searches. *Communications of the ACM*, 41(4): 95-98.Shneiderman, B., Feldman, D., Rose, A., & Grau, X. F. (2000). Visualizing digital library search results with categorical and hierarchial axes. In *Proceedings of the Fifth ACM International Conference on Digital Libraries (San Antonio, TX, June 2-7, 2000)* (pp. 57-66). New York: ACM Press.

Sim, D-G and Park, R-H. (1997). A two-stage algorithm for motion discontinuity-preserving optical flow estimation. *Computer Vision and Image Understanding*, 65(1): 19-37.

Smith, M. and Kanade, T. (1998). Video skimming and characterization through the combination of image and language understanding. In *Proceedings of IEEE International Workshop on Content-based Access of Image and Video Database*, pp. 61-70.

Song, Y. and Marchionini, G. (2007). Effects of audio and visual Surrogates for making sense of digital video. To appear in *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems*, (in press).

Spink, A.,Griesdorf, H. and Bateman, J. (1998). From highly relevant to not relevant: Examining different regions of relevance. *Information Processing and Management*, 34(5): 599-621.

Spink, A., Jansen, B. J., Wolfram, D. and Saracevic, T. (2002). From E-Sex to E-Commerce: Web search changes. *IEEE Computer*, 35(3): 107-109.

Srinivasan, S., Petkovic, D., and Ponceleon, D.(1999). Toward robust features for classifying audio in the CueVideo system. In *Proceedings of ACM Multimedia*, 393-400.

Suomela, S. and Kekäläinen, J. (2005). Ontology as a search-tool: A study of real users' query formulation with and without conceptual support. In *Proceedings of the European Conference on Information Retrieval Research*, pp. 315-329.

Taylor, R.S. (1968). Question-negotiation and information seeking in libraries. *College and Research Libraries*, 29, 178-194.

Teodosio, L. and Bender, W. (1993). Salient stills from video. In *Proceedings of ACM Mulitmedia*, pp. 39-46.

Tianmiyu, M.A. and Ajiferuke, I.Y. (1988). A total relevance a document interaction effects model for the evaluation of information retrieval processes. *Information Processing and Management,* 24 (4): 391-404.

Vakkari, P. (2004). Subject knowledge, thesaurus-assisted query expansion and search success. In *Proceedings of the RIAO Conference*, 393-404.

Wexelblat, A. and Maes, P. (1999). Footprints: history-rich tools for information foraging. In *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems*, pp. 270-277.

White, R.W., Bilenko, M. and Cucerzan, S. (2007). Studying the use of popular destinations to support Web search interaction. *Manuscript under review.*

White, R.W. and Marchionini, G. (2007). Examining the effectiveness of real-time query expansion. *Information Processing and Management*, 43(3): 685-704.

White, R.W. and Ruthven, I. (2006). A study of interface support mechanisms for interactive information retrieval. *Journal of the American Society for Information Science and Technology*, 57(7): 933-948.

White, R.W., Jose, J.M. and Ruthven, I. (2005). Using top-ranking sentences to facilitate effective information access. *Journal of the American Society for Information Science and Technology*, 56(10): 1113-1125.

White, R.W., Song, H. and Liu, J. (2006). Concept maps for oral history search and use. In *Proceedings of the ACM Joint Conference on Digital Libraries*, pp. 192-193.

White, R.W., Jose, J.M. and Ruthven, I. (2003). The influencing effects of query-biased summarisation in web searching. *Information Processing and Management*, 39(5): 707-733.

Wildemuth, B., Marchionini, G., Yang. M., Geisler, G., Wilkens, T., Hughes, A. and Gruss, R. (2003). How fast is too fast? Evaluating fast forward surrogates for digital video. In *Proceedings of the ACM/IEEE Joint Conference on Research on Digital Libraries*, pp. 221-230.

Wilson, T. D. (1997). Information behaviour: an interdisciplinary perspective. *Information Processing and Management,* 33(4), 551-572

Witbrock, M. and Hauptmann, A. (1998). Artificial intelligence techniques in a digital video library. *Journal of the American Society for Information Science,* 49(7): 619-632.

Zhang, J., and Marchionini, G. (2005). Evaluation and evolution of a browse and search interface: Relation Browser++. *The National Conference on Digital Government Research.* (Atlanta, GA: May 15-18, 2005). ACM Press. 179-188.