Chapter 6. Browsing Strategies

Marco Polo had the opportunity of acquiring a knowledge, either by his own observation or what he collected from others, of so many things, until his time unknown." (The Travels of Marco Polo)

It is one thing to remember, another to know. Remembering is merely safeguarding something entrusted to the memory; knowing means making everything your own. (Seneca, Epistle ad Lucilium).

In contrast to the formal, analytical strategies developed by professional intermediaries, information seekers also use a variety of informal, heuristic strategies. These informal, interactive strategies are clustered together under the term *browsing strategie*s. In general, *browsing* is an approach to information seeking that is informal and opportunistic, and depends heavily on the information environment. Four specific browsing strategies are distinguished in this chapter: scanning, observing, navigating, and monitoring. The term browsing reflects the general behavior that people exhibit as they seek information by using one of these strategies.

Browsing is a natural and effective approach to many types of information-seeking problems. It is natural because it coordinates human physical, emotive, and cognitive resources in the same way that humans monitor the physical world and search for physical objects. It can be effective because the environment and particularly human-created environments are generally organized and highly redundant--especially information environments that are designed according to organizational principles. Browsing is particularly effective for information problems that are ill-defined or interdisciplinary, and when the goal of information seeking is to gather overview information about a topic, or to keep abreast of developments in a field.

The term "browsing" finds use in a variety of settings, including: forestry, zoology, architectural design, marketing, as well as information science (Chang & Rice, 1993). The computer and information science literature uses the term browsing simply to mean navigating, scanning, and scrolling, and many interfaces include "browsers" that allow users to move across or within screens, windows, records and databases. In traditional library literature, browsing was most often associated with card catalogs or book shelves. More recently, Bates (1989) has argued for the importance of browsing in conjunction with other online search strategies. We term this type of browsing, *across-document browsing*, since the information seeker browses across records or books to find items to examine more closely. The education and reading literature also have been concerned with perusing a document to extract gist or to locate a relevant passage, and terms such as skimming, scanning or browsing are used to distinguish this activity from reading or studying[1]. We term this type of browsing *within-document browsing* and believe it to be a high-priority research problem for electronic environments. Most types of across-document browsing aims to identify specific documents or objects that may then be browsed or studied according to the information seekers's needs. One exception is across-document browsing in a collection to gain a sense of scope or form. Thus, browsing is a term whose meaning must be taken in context.

A significant characteristic of browsing in electronic environments is the blurring of boundaries between document collections or databases and discrete documents or records. The implication for design is to provide systems that support seamless browsing across and within documents. A hypertext collection of articles may be considered a single document when linked to other hypertexts or as a collection of documents (nodes) in its own right. In either case, a single node or a collection of nodes is represented on the same display device, and the display is controlled by the same paging, scrolling, and windowing actions. In present electronic environments, there are no clear delineations for the user within either the physical display device or the physical actions taken to execute browsing activity. This situation is exacerbated in networked environments where documents from different databases appear in the same window or the same document is represented in different databases.

With respect to the information seeking framework presented in Chapter 3, browsing strategies are more dependent than analytical strategies on interactions between the information seeker and the system; the system has more influence on search progress during browsing. Also, during browsing the information-seeking subprocesses proceed in somewhat more parallel fashion, more time is spent in the examination subprocess, and there are a wider variety of cycles within the overall process.

Browsing offers significant challenges to information seekers and system designers. First, an entry point must be identified. Determining an entry point combines problem definition and query formulation. Entry points can be determined by random selection, by use of indexes or statistical analyses that return a set of items for browsing, and by opportunistic, iterative examination. The challenge to the information seeker is to relate personal knowledge about the topic to what the system represents and how its representations are organized. The challenge to designers is to make clear the system's scope and organization, and to suggest entry points to the searcher.

Second, information must be examined and assessed during browsing. The information seeker's challenge is to make rapid judgments of relevance, and the designer's challenge is to provide flexible display facilities to facilitate examination and assessment. These challenges illustrate how browsing blurs the information-seeking subprocesses. Problem definition, query formulation, execution, examination, and iteration are especially intertwined during browsing. This fluidity and integration is due to the parallel engagement of physical, perceptual, and cognitive processes during browsing; to the close coupling between the information seeker and the organizations and representations provided by the information environment; and to the open-ended information problems that lend themselves to browsing.

This chapter is rooted in the notion that the natural human inclination to browse and the organizations inherent in information systems are strongly reinforced and harmonized in highly interactive electronic environments that compliment and invite browsing. Based on this perspective, we argue for a design philosophy that amplifies human abilities and inclinations rather than one that aims to optimize organization and task analysis. We consider the variety of reasons why people browse, different types of browsing behavior, limitations of browsing, and a variety of ways that manual and electronic information systems can support browsing.

WHY BROWSE?

Browsing is used for a variety of purposes and manifests itself in different ways. Table 6.1 illustrates the main reasons people browse. In contrast to analytical strategies where goal definition is important for success and efficiency, browsing strategies may be applied to more informal or general goals and depend more heavily on the information encountered during search.

[Table 1 about here]

People browse to gain an overview of a physical or conceptual space. By scanning a scene or a document, key landmarks and characteristics are identified and used to form impressions of the scene or document and to make analogies to known scenes or concepts. For example, by scanning the title page, table of contents, section headings, index, and reference list of a book, we gain a sense of content scope, depth of coverage, the organizational perspective of the author, and are able to make rapid decisions about whether to invest time reading it. It is important to note that in the case of books, these attributes that we browse first are well-established standards to aid browsing. Hildreth (1982) points out that these features of books both encourage and enhance browsing. Electronic environments are only beginning to build in systematic support for browsing. We can browse software applications by spending a few minutes traversing menus, executing generic commands, or manipulating system objects, and examining the form and scope of help and instructions. At this point in the development of the science of interface design, it is unclear how a program or a database should "reveal itself" to users to encourage and enhance browsing. Full-text CD-ROM databases allow users to browse texts by jumping to highlighted query terms, but these systems are constrained by the limits of display technology. Browsing Internet newsgroups gives a sense of breadth for various groups and details for specific postings, but few organizational cues to gain a quick sense of scope and organization of all the groups.

Another reason people browse is to monitor a process. When driving, we regularly check how close we are to the road shoulder, look for landmarks and signs, and glance in the rear-view mirror to check the status of cars near to us. When systematically examining books on a shelf, we often jump ahead to gain a sense of how much time will be needed to complete exhaustive browsing. Some electronic search systems provide history or path tools so that users can quickly go back and examine the progress of their search. Readers often scan forward in a book or article to see what remains and to establish an anticipatory frame of reference, or scan backwards to quickly review what has been read and regain or revise their sense of context. These actions are indicative of the text reinspection reading strategy used by skilled readers in comprehending text (Garner, 1987). This type of browsing is particularly important for ongoing, accretional information seeking where the objective is to stay abreast of a field. Many information environments provide aids for monitoring such as sidebars, pull quotes, abstracts, review articles and book reviews. Developing interfaces that support and amplify browsing through the mass of information created in a domain is one of the most critical problems in information science and human-computer interaction.

People also browse because it requires less cognitive load than analytical search strategies. Browsing is highly dependent on human perceptual abilities to recognize relevant information. Analytical query formulations require us to apply cognitive resources to recall from memory specific terms that represent the concepts related to our problem. Taken to extreme, the advantage of lower cognitive load may be abused. For example, through cognitive laziness, information seekers may apply minimal cognitive effort and simply depend on perceptual and motor processes to stumble upon relevant information. In general, however, applying recognition allows the information seeker to concentrate more cognitive resources on the problem at hand. For tasks that are ill-defined or complex, browsing allows information seekers to devote full cognitive resources to problem definition and system manipulation and to involve the perceptual system in filtering words or images.

In addition to bringing additional resources from the perceptual system, there may be symbiotic advantages in applying multiple subsystems that are mutually reinforcing--in effect the whole result is greater than the sum of its parts. This parallelism effect is perhaps even stronger because browsing requires high levels of direct physical engagement on the part of the information seeker; we must move, scan, scroll, jump, and page more quickly through the shelves, pages, or screens. Thus, browsing may be more of a direct manipulation strategy than formal methods because it involves cognitive, perceptual, and motor systems concurrently and continually.

Using both physical and mental subsystems in coordinated activity may reinforce experience and memory for events. Our interviews with and observations of students using the Perseus system suggest that there may be some possibly positive interactions between the physical actions of using a lexicon and the cognitive actions of translating

Greek to English. The popularity of video games and simulations that use high degrees of direct manipulation also illustrates the importance of engaging multiple subsystems in complex activity. Our understanding of how the satisfaction associated with direct manipulation affects performance and learning is very poor. There is no definitive evidence whether learners remember more as a result of coordinating physical and mental subsystems or how intellectual task performance such as information seeking is affected. If performance can be improved by engaging physical and mental subsystems more intimately, there must be design efforts to coordinate physical and intellectual actions at workstations. Furthermore, for some specific types of information seeking, such as locating intermediate facts in a chain of processes or repetitive lookups, experience and memory should be minimized and speed and accuracy maximized.

While seeking information, people also browse to clarify the information problem. Since many of our problems are due to anomalous states of knowledge, browsing a space of information in the problem domain can help us to clarify and expand our knowledge state. We may retrieve a record, article, or book and examine it quickly to look for ideas about entry terms, alternative or related concepts, or to learn more about the problem so that it can be clarified and stated with precise language. Likewise, we browse to find something interesting that suggests a more specific or exciting problem that can then be attacked more formally. For example, in our study of attorneys and law intermediaries, subjects noted that one reason to go online rather than use inexpensive printed sources was that the problem was ill-defined (Marchionini, Dwiggins, Katz, & Lin, 1993). The electronic systems allowed these information seekers to quickly locate sets of legal documents that provided feedback to clarify what issues were involved in the problem.

Related to problem clarification, people browse to stimulate or develop a plan or a formal strategy. An information seeker may use a simple query to retrieve a large number of generally related documents and browse the title words or subject headings to determine technical vocabulary, identify related facets, and generally get to know the intellectual neighborhood before formulating a specific set of formal queries. This strategy is particularly useful if the information seeker is not an expert in the domain or when controlled vocabulary lists or thesauri are unavailable.

People also browse to discover and learn. Regardless of how comprehensive our efforts to investigate an area, we all experience the pleasure of serendipitous discoveries that are highly relevant to our information problems but found in surprising places and situations. Browsing is an activity that may yield new insights in two ways. On one hand, we may locate a highly relevant item by browsing resources beyond those systematically coded and assigned as relevant. This has low probability but high payoff--like an intellectual lottery. This is the basis of many arguments for serendipity. For example, we may find a highly relevant paper that was not indexed with terms guiding the search or we may locate an alternative special interest group.

On the other hand, we may not locate any specific document that would otherwise have been missed, but rather gain new insights or interesting associations for our problem by browsing alternative sources that use different tools, techniques, and data structures. Scholars are highly conscious of such possibilities and support institutional policies to facilitate such serendipity, for example, open stacks in libraries and cross-disciplinary colloquia. By browsing in an unfamiliar field, we may discover methodologies that have not been applied in our own field or gain new perspectives or metaphors to apply in our own work. In interdisciplinary fields like information science, scholars must use results from the many allied fields. In many cases they work at the boundaries of several fields and depend on browsing to span the boundaries to gain synthetic insights and integrate concepts. For example, Auster and Choo (1993) have noted that executives depend on scanning information sources across traditional boundaries to stay abreast in their field.

Finally, we browse because the environment invites browsing--in fact, stores and museums are carefully designed to stimulate browsing. The organizations in nature and created by civilization invite interaction and exploration. The importance of how systems invite and support browsing is considered in a subsequent section.

Browsing has become much more important as more information resources migrate to electronic environments. The interactivity characteristic of these environments affects all the reasons for browsing above but has particularly interesting possibilities for discovery and learning. Humans make decisions about how to use their time and resources for this type of information seeking according to constraints on these resources, probabilities for finding new knowledge, and payoffs for new knowledge. For example, during a brainstorming phase of information seeking or at an impasse in the process, if 20 minutes are available to invest in an environment (e.g., a library), a min-max heuristic may be applied to examine low probability sources to gain high payoff information that may suggest directions for beginning or restarting work. Electronic environments provide good potential to lessen time constraints and raise the probabilities of locating high-payoff information. First, electronic environments are more accessible and may invite more frequent explorations. Second, electronic environments can improve the odds of likelihood of payoff by providing feedback that is controllable by the user directly or by a profile automatically.

BROWSING STRATEGIES AND TACTICS

The literature of library and information science contains many attempts to distinguish different types of browsing. Review of browsing (Chang & Rice, 1993; Hildreth, 1982) cite a number of typologies developed by different investigators. Although there are differences in how the types are named, there seems to be agreement on three general types of browsing that may be differentiated by object of search (the information needed) and by the systematicity of tactics used. At one extreme, users perform directed (Herner, 1970) or specific (Apted, 1971) browsing. This form of browsing is very systematic and focused and is often driven by a specific object or target. For example, scanning a list for a known item is a highly directed type of browsing since the objective is highly specific in nature and known by the information seeker. Tasks such as verifying information, or retrieving dates or other facts allow repetitive and highly systematic browsing behavior. A second type of browsing is semidirected and predictive (Herner, 1970), or general purposive (Apted, 1971). This type of browsing has a less defined target and proceeds in a less systematic fashion. For example, entering a single, general term in a full-text electronic database and casually examining the retrieved records. This type of browsing allows various degrees of target definition but requires multiple probes of the system and thus is time consuming. A third type of browsing is undirected (Herner, 1970) or general (Apted, 1971) browsing that is characterized as having no real goal and very little focus--more like recreation than information seeking. For example, repeatedly changing television channels to find something of interest to watch.

In addressing browsing in electronic systems, Cove and Walsh (1988) and others have described a similar trio of browsing types based on the not-quite orthogonal dimensions of goal definition and behavioral systematicity. Thus, types of browsing have been characterized coarsely according to the specificity of the object sought, which in turn influences the systematicity of actions the information seeker takes while browsing. These three different degrees of systematicity will be referred to from most to least systematic as *systematic*, *opportunistic*, and *casual* browsing. Systematic browsing exemplifies more linear sequencing and regular iterations of the information-seeking subprocesses; opportunistic browsing tends to jump from examination to many different subprocesses; and casual browsing exemplifies almost random sequences of information-seeking subprocesses. Such behavioral characterizations are useful for describing browsing as an information-seeking strategy, but a more detailed and cognitively-oriented analysis is required to understand browsing, especially in regard to the role of the object of information sought.

One step in this direction is Kwasnik's study of how individuals browsed in paper, a command-driven electronic system, and a hypertext to identify the behavioral components of browsing (in press). In preliminary studies of people in a marketplace, she noted the difficulty of identifying a unit of focus (e.g., the object of browsing could be an entire menu or a menu item), that people tend to impose structure on the environment regardless of how unstructured it may be, and that even casual browsing quickly becomes purposeful as people relate the activity to their own interests. She identified six functions her subjects used while browsing. 1) Orientation refers to their need to understand the structure and content of the environment and continues to evolve as browsing continues. 2) Place marking allows users to remember salient objects, and 3) identification refers to their decisions about the relevance of specific objects. 4) Resolution of anomalies refers to browsers' attempts to understand objects that were not clearly identified. 5) Comparison and 6) transition are functions that respectively support orientation and identification, and allow browsing to progress. This work provides a more user-centered and cognitive approach to browsing than previous work.

A user-centered model of browsing based on the general information-seeking framework that distinguishes different browsing strategies follows. In practice, information seekers use multiple strategies and often use default strategies to initiate search, but this model provides a basis for predicting rational initial strategy selections.

Like all information-seeking activity, browsing is dependent on interactions among the information-seeking factors: problem, domain, setting, user characteristics and experience, and system content and interface. Although the higher order interactions of all the factors influence the specific browsing tactics and behaviors information seekers select and exhibit, five sets of interactions determine dimensions that will be used to characterize browsing strategies. Two of these dimensions relate to the external and mental representations of the object sought, one is specific to the environment, and two relate to the physical and mental activity of the information seeker. The interactions are:

- The domain of the problem and the system mainly determine the external representation of the **object sought**,
- the problem and the seeker's knowledge of the domain and the system mainly determine the mental representation of the **object sough**t,
- the system and the setting determine the degree and type of organization of information in the **environment**,
- the system, the setting, and the seeker's mental models for them mainly determine which browsing strategies are available for application and the **degree of interactivity**, and
- the setting, and the information seeker's past experience and immediate mental and emotional states mainly determine the **cognitive effort**, motivation and resources applied during browsing.

The object sought during browsing must be considered from two perspectives: the characteristics of the object in the world and the representation of the object in the information seeker's mind[2]. The object may be a well-specified, discrete target (e.g., a book) that is identifiable by a single attribute (e.g., a call number). Such objects are well-managed by database management systems or string search techniques. If the single attribute is known by the information seeker (i.e., there is high definition in the user’s mental representation), such objects may be immediately recognizable upon perception (e.g., known item searches that seek to locate a specific fact or item). The object may be more general and complex (e.g., a concept or a document), and identifiable by an aggregation of multiple attributes. These objects are typically represented and managed with value-added techniques such as indexing and abstracting. Browsing for these types of objects requires that the information seeker perceive multiple attributes and execute multiple cognitive comparisons and inferences to ascertain pertinence. Objects may also be highly abstract and active (e.g., threads of argumentation in a longitudinal collection of group discussions) and identifiable by a variety of clusters of overlapping attributes. These types of objects may be saved and displayed by various techniques, but are not easily manageable. Browsing for such objects demands substantial cognitive effort and reflection on the part of the information seeker.

The objects of information seeking must also be represented in the mind of the information seeker. There are several factors that influence these representations. As discussed in Chapter 3, the information problem may be driven by the need to retrieve, to analyze, or to learn and discover; the information seeker will have a set amount of domain knowledge, information-seeking experience, personal characteristics, and immediate states of affect; and various setting constraints will be active. The object of information seeking may be represented in the seeker's mind as a discrete object with a single attribute (e.g., a date). Such objects allow the browser to simply compare the objects in the world on this attribute--a low cognitive load task. In fact, this sort of problem lends itself to more straightforward retrieval and browsing may demand considerably more user effort. Alternatively, the object can be represented in the browser's mind as a complex concept with multiple attributes. This situation will require more systematic comparisons, rankings of attribute salience, and thus more cognitive load. If the objects are represented as anomalous states of knowledge with many possible attributes which are themselves fuzzy, then the browser must not only execute many comparisons and inferences, but also make many interpretations about what may serve as attribute values.

The more broadly we decide to browse, the less constrained are the attributes and corresponding objects that we are willing to admit for careful examination. Furthermore, the attribute values may change as information seeking progresses and the information seeker's state of knowledge becomes less fuzzy. Thus, browsing for ill-defined or anomalous objects requires increased perceptual and cognitive load, and more time. It is also possible that information seekers may browse according to a conscious decision to engage in divergent thinking (no object is allowed), or by efforts to explore a space to select objects with similar benefits or preferences (choosing one object from many desirable objects). In these cases, as potential objects are perceived during browsing, various levels of cognitive effort may be applied to assess pertinence and determine what action to take next. Thus, browsing proceeds by comparing the mental representation for the object (an image of the object together with a set of attributes believed to be salient to the search) with the system's representation and altering expectations according to the outcome of each pattern-matching event.

The environment is determined by interactions among setting, search system, domain and task. The environment supports or inhibits browsing according to the level of organization for the objects represented and by the manner in which feedback is provided. A subsequent section is devoted to the important role that systems play in browsing, for the purpose of defining browsing tactics, degree of organization is used as a key dimension. Electronic environments may be characterized by the degree of feedback provided to users. This is the system's perspective of the amount of interactivity, a dimension considered from the user perspective here.

The information seeker is central to all information-seeking activity and since browsing is a specific case of this activity, the information seeker's personal information infrastructure and the situated moment influence browsing. The personal information infrastructure determines whether to browse or use an analytical strategy, what entry point to use, and which tactics to apply as browsing progresses. The user's emotional state, level of need, and resources such as time, money, and access privilege influence these decisions, determine whether systematic, opportunistic, or casual approaches are taken, and affect persistence and termination. For the purposes of defining browsing strategies, two dimensions are used here: interactivity and cognitive effort. The interactivity dimension is defined by the number and rate of choices and actions the user makes and takes during information seeking. Low interactivity is indicated by few choices for types of actions and infrequent turn-taking. Cognitive effort refers to the amount of reflection, analysis, integration, and decision making required during browsing. Low cognitive load is indicated by primarily perceptual processing and simple recognition and matching.

The types of actions information seekers exhibit when browsing are dependent on interactions among the information-seeking factors. From a user-centered, cognitive perspective, tactics and moves are mainly determined by the information seeker's internal representations of the sought object and by expectations about how the process of browsing will proceed. Browsers may expect that comparisons will be based on a simple match of a single attribute or that matching must be made on multiple attributes. In the case of more open-ended browsing, the experienced browser will expect to draw analogies and make inferences about whether information items found in the system are suitable as attributes for the fuzzy objects of interest. The internal representation of the object of browsing and the expectations about the kinds of matching that will be needed influence the tactics and moves the information seeker applies during browsing. For example, a single, well-defined attribute that a browser expects to occur in a section heading of an article will likely invoke a scanning strategy across section headings.

It is important to note that both types of mental representations change as browsing progresses. The object becomes more defined based upon attribute specification and frequency of occurrence and expectations about further progress are informed by reflection on the quality and quantity of attributes represented in the system. For example, consider a novice stock market investor who seeks information on how stock options work and has accessed an electronic discussion archive that provides scrollable subject headings posted by contributors to the discussion. Suppose that the terms "stock" and "option" are attributes the information seeker expects to match while scrolling. As browsing continues, additional terms may be added as attributes if those terms occur frequently with "stock" and "option" (e.g., the term, "contract") or understanding may grow based on frequent co-occurrences of other concepts such as "option listing," "put," "call," etc. These changes are dependent on feedback from the system and the degree of change is a simple indicator for distinguishing systematic, opportunistic, and casual browsing, respectively. An interaction between an information seeker and a system that exhibits occasional large changes in the mental representations for the object and the expectations about the process is more clearly opportunistic or casual than an interaction in which the object and expectations about how matching takes place vary incrementally as search progresses.

The behavioral moves exhibited by browsers are manifestations of browsing tactics. The physical actions should be good indicators of cognitive processing because browsing depends on close coupling of physical, perceptual, and cognitive activity. Four types of browsing strategies are defined: scan, observe, navigate, and monitor. Figure 6.1 categorizes these strategies by the five primary dimensions critical to browsing. Two dimensions are used for object: degree of external definition (high means very well-defined objects, e.g., simple, concrete data) and degree of definition in the information seeker's mind. External definition is mainly determined by interaction between the domain and search system. The degree of environmental organization for object representations is the dimension used for scaling environment. Environmental organization is mainly determined by interactions among the search system and the setting, with problem and domain to a lesser extent. The degree of interactivity (high means multiple, rapid turn-taking actions) and amount of cognitive effort (high means considerable reflection, decision making, and analysis while examining objects) are used as dimensions for the information seeker element. Interactivity is determined by interactions among all the information-seeking factors, with domain and outcomes contributing the least effects. Cognitive effort is primarily determined by information seeker characteristics and the problem with other factors contributing smaller effects. The values given in the figure are considered to be threshold values where the strategies become useful and typical. Clearly, any strategy can be applied for any browsing situation, but efficiency and effectiveness are determined by appropriate applications. Thus, any of the strategies may be used during systematic, opportunistic, or casual browsing.

[Place Figure 6.1 about here]

Scanning is the most basic browsing strategy. It is fundamentally a perceptual recognition activity that compares sets of well-defined objects to an object that is clearly represented in the information seeker's mind. Scanning benefits from highly organized environments that provide clear and concise representations. It can proceed in sequential fashion, according to some structural feature of the content, or through some sampling method. Two such tactical approaches are linear and selective scanning.

Linear scanning applies perceptual attention in continual and sequential fashion. For example, information seekers may scan title lists to identify potentially relevant documents to examine, or may fast-forward through television channels to locate an interesting program. Linear scanning is most applicable to sequential arrangements of similar objects with precise attributes that are recognizable with a single glance, and when the collection is reasonably small or the information seeker is confident about being in the general neighborhood that contains the object. Linear scanning is effective since the eye can recognize simple patterns in as little as 50 milliseconds, although eye movement plus recognition times vary from 125 to 500 ms. (Potter & Levy, 1969). Simon (1979) reports 250 ms fixations for people scanning squares on a chess board. Latencies are shorter when people know where a target will appear (Rayner, 1978), as is the case in linear scanning or reading.

A special type of linear scanning is monitoring of gauges or controls and there is evidence that linear sequences of stimuli in a stable position (e.g., reading by fixing on a screen position and having text move through the position) can dramatically decrease response latency. Payne and Lang (1991) termed such techniques "rapid communication visual display" and found response latencies with them to be two to three times faster than spatial displays; however, these speed increases brought severe penalties in error rates that were twice as high as rates with the spatial displays. How these speed and accuracy tradeoffs affect scanning tactics for different tasks and how they may be applied to interface design remain to be determined. For example, there is no evidence for how errors of omission or commission affect browsing for textual information and it may be that increased speed benefits may be worth such error penalties in some types of browsing. Errors of omission, for example, may be acceptable when scanning journals or databases to gain a sense of gist. Linear scanning is fatiguing, but the physiological and psychological bounds of scanning in electronic environments have yet to be determined.

An alternative to linear scanning is *selective scanning*. This tactic applies perceptual attention according to either an inherent or imposed stratification of the search space. Most search systems offer some inherent partitions that support selective sampling. For example, information seekers may scan the section headings within a book or the reference lists of different journal articles. This tactic is particularly useful to gain an overview of content or to identify a neighborhood for linear scanning. A variation of this tactic is to scan selectively by sampling the search space according to a partitioning rule. Rather than depending on some features or organizations in the search space, random or purposive samples of the search space are identified and scanned for relevance. For example, information seekers may fast-forward a videotape 100 feet at a time to locate a particular sequence, or may examine every first sentence of each screen of text to get a sense of content. The partitioning rule may take advantage of spatial memory and may be particularly useful in locating objects previously seen in the system (e.g., a statement that we believe appears about two-thirds of the way down a left-facing page about one-third of the way through the second chapter of a book). This tactic could be applied to gain an unbiased overview of database coverage or to make determinations of accuracy or integrity. Sample scanning may also be appropriate in discovering promising information neighborhoods in extremely large, unstructured databases.

Scanning tactics are most applicable to organized search spaces of reasonable size. What constitutes "reasonable" is often underestimated, especially in electronic environments, nonetheless one of the most obvious ways that electronic search systems can amplify human information-seeking abilities is by emphasizing and facilitating scanning tactics. Wiberley and Daugherty (1988) summarize literature related to library patron across-document scanning behaviors and suggest that users are willing to scan longer lists of citations from online searches than from OPAC searches and longer lists from electronic than manual systems. They report ranges of 7 to 50 items scanned with few above 20 for manual systems, but occasionally up to 200 for electronic systems. They also note that poor interfaces in electronic systems minimize scanning persistence. It seems plausible that if systems are specifically designed to support scanning rather than simply permit it, substantial lists of references will be scanned by users. There is little evidence on how users scan within electronic documents, although simple affordances such as highlighted query terms and indicators of progress through the document clearly assist within document scanning.

Scanning tactics are most useful for tasks where recognizable and discrete attributes are available. Scanning tactics are most often applied during systematic browsing strategies or for intermediate examinations during opportunistic or casual browsing. One measure of the costs of scanning tactics is attention time and research on scanning in electronic environments will be well-served by analyses of these costs as influenced by different information-seeking factors.

Observational strategies are the most general of all browsing strategies in that they have minimal thresholds for all the browsing dimensions except for cognitive effort. Browsers who apply these strategies assume they are in a promising neighborhood and react to stimuli from that neighborhood. Observational strategies depend on a great deal of parallel input. In busy street scenes, a browser may attend to a variety of sounds, sights, and motions rather than systematically scanning the scene. Advertisements in newspapers and magazines and provocative titles of books on shelves attract attention as we passively apply observational tactics. Like scanning, observational strategies are rooted in our physiological and psychological survival instincts and thus are naturally and easily applied as defaults. Observation does require interpretation and reflection to make sense of what is observed and to relate it to information-seeking objectives. Observations may lead to interesting discoveries but yield most initiation control to the environment. For this reason, it is most important for the information seeker to be in a relevant neighborhood.

Observations can be executed in systematic fashion but for ill-defined objects and fuzzy purposes, for example, regularly reading the morning newspaper or watching a news broadcast. Observation is the primary strategy used in opportunistic and casual browsing since it admits the widest range of objects and unorganized environments. Serendipitous observations are lauded as a benefit of browsing, but they occur as a result of reflection and association rather than simple perception and pattern matching.

Interfaces that aim to support observational strategies should provide alternative views of information. Because observational strategies yield significant control to the environment, information should be clearly represented and demarcated. Alternative representations should be available for user to transform so that pattern recognition may lead to association and reflection.

The *navigation* strategy balances the influence of the user and the environment. The environment constrains browsing by providing possible routes and the user exercises some control by selecting which routes to follow. Navigation is a term that is used broadly and differently in the literature. In much of the human-computer interaction literature, navigation is considered as synonymous with browsing rather than one of several specific browsing strategies. One model that does distinguish browsing and navigation was proposed by Waterworth & Chignell (1991). In their model of information exploration, they distinguish querying from browsing, object specification (command interaction styles) from object recognition (menu interaction styles), and navigation from mediation. In their model, navigation is taken to mean a high degree of user control and mediation depends more on system control.

As considered here, navigation is defined by relatively high thresholds for all the browsing dimensions. Objects must be specifable; information seekers must know what they are seeking, take an active role in interacting with the environment, and regularly reflect and make decisions about progress; and although navigation may occur in unstructured environments, high levels of organization greatly aid efficiency and effectiveness.

Physical navigation is often used as a metaphor for traversal of a hypertext or a database. The activity users engage in as they follow links in a hypertext is the navigation that we consider to be a browsing strategy. In this strategy, the process and the information found along the way are what is important. A better metaphor for this strategy may be “grazing” or “berry-picking” (Bates, 1989), but the activity is commonly called navigation in the human-computer interaction and hypertext literature and so is adopted here. This type of navigation is in contrast to the goal-oriented physical navigation that uses cues from the environment only to check progress toward the goal.

In physical navigation, the object sought is most often a predetermined destination and the process involves adjusting course according to attributes (conditions) provided by the environment. Navigation connotes goal-oriented behavior where plans and subplans respond to the environment. This dynamic interplay between navigator and environment is what makes navigation a popular metaphor for human-computer interaction. Navigation is a useful name for a browsing strategy in so far as it denotes observations of the environment and adjustments in behavior based on these observations. It is a weak metaphor for browsing in an information space in two respects. Most importantly, the destination for information seeking is seldom predetermined--intellectual space is highly amorphous and much of what is called navigation involves problem definition and clarification. Second, information-seeking activity draws information to the information seeker rather than transporting the information seeker to the information. It is especially true in electronic environments that we remain physically stationary and gather information to our screen rather than travel to some location and idea. Using navigation as a global metaphor for browsing admits the side effects of becoming lost. In electronic systems this effect has been most often referred to as lost in hyperspace (Nielsen, 1990a). It is useful to distinguish being lost and confused as a result of the environment (e.g., poor maps or bad road signs in physical space, and poor interfaces in information space) and being lost or confused as a result of the information problem (i.e., being lost in one's own mind). Although it is likely that improved interfaces (e.g., graphic maps, place markers, organizational labels and cues) will minimize disorientation due to the system, being lost in a collection of thoughts will remain a human problem regardless of system advances.

This point of view that understanding content (mental orientation) is more critical than system orientation was reinforced by two investigations we conducted to determine the effects of blatant metaphors on learning. Two versions of a hypertext explanation of the information seeking framework were prepared using the GUIDE system. One version, called the "jump" implementation, reinforced the navigation metaphor. The screen consisted of a single text window and link anchors were immediately followed by the word "jump" in bold upper case letters. Clicking on the link anchor replaced the text with the text in the linked node. The other implementation, called the "bring" version, reinforced the metaphor of bringing information to the screen rather than jumping to it. Link anchors were immediately followed by the word "bring" in bold upper case letters. Clicking on the link anchor opened a new window that overlapped the active window immediately below the anchor.

Forty-eight library and information science students were randomly assigned to one of the two conditions. They received mouse training and training in their assigned linking mechanism during which time the metaphor was verbally reinforced. They were then were given 30 minutes to study the database. After their study period, subjects were given a written test to determine their understanding, asked to draw a diagram of the organizational structure of the database, complete a questionnaire, and participated in a focus group discussion. No statistically reliable differences were found between the groups and there was wide variance on all indicators within the groups. The study was replicated the following semester with a third condition of a linear electronic version of the database added. Results in that study were similar to the first study. In both cases, reports of disorientation were not correlated with other results or any of the main treatments. Apparently, the metaphor was not as important to performance as user abilities and the complexity of the content.

Navigation as a browsing strategy refers to the ongoing observation of environmental attributes, adjustments to the mental problem representation based on these observations, and resulting behavioral actions. Navigation is thus information seeking that proceeds incrementally based on feedback from the information system. The ways the system provides feedback is a critical factor in navigation.

This strategy is most clearly supported by the structural layouts of museums. Navigation has become an important (but weak) metaphor for electronic systems and is the primary inspiration for static hypertexts. In one navigation variation, users are invited to follow paths through the database by selecting one of possibly many links from a current node to other nodes. Thus, the tactics used are simple selections from the choices active on the path. Navigation strategies can be applied in casual or systematic fashion, although they depend on taking advantage of existing links or the user's ability to create new links. Navigation is an attractive compromise between user and system responsibility since the system invites or suggests links to follow but the user is free to choose from among many links or continue in linear fashion. Different systems provide distinct levels of navigational freedom, ranging from highly systematic, mandatory paths to absolutely no guidance whatsoever. Successful systems fall somewhere in between, providing some choice to users but suggesting directions and providing informative cues about progress. Systems that provide explicit hypertextual links support more systematic browsing strategies and those that provide implicit links support more exploratory or casual browsing.

A *monitor* strategy is often applied in conjunction with systematic browsing or other primary activities such as reading. Monitoring is most similar to scanning except it tolerates poorly structured environments. For example, while reading text related to a specific topic, a monitor browsing strategy "listens" for concepts related to another topic of interest. Professionals in any field often apply the monitor strategy automatically by spontaneously relating what they read in an unrelated area to their own field. It is this strategy that is perhaps most important for discovery and creative connections among disparate ideas, and is the reason scholars revere serendipity. Although this strategy may be partly subconscious, it depends on the user making associations among concepts in the mind and representations in the information space. Monitoring is enabled by perspectives or views of the database, which in turn are made up of various cues such as words or phrases, movements, or visual characteristics, depending on the information environment. Monitor strategies are focused on attributes of interest to the information seeker and are less dependent on stimuli in the environment than observational and navigation strategies.

These various strategies are applied by information seekers as part of the information-seeking process. Scanning and simple path following navigation are most often used in systematic browsing where objects are well-defined in the information seeker's mind and in the environment. These strategies are indicated by typically linear sequences of information-seeking subprocesses. Observation, navigation by cues in the environment, and active monitoring strategies are most often used in opportunistic browsing when exploration, learning, or accretion of knowledge are goals, objects are complex or ill-defined, or systems are unfamiliar or unstructured. These strategies require more cognitive load than systematic scanning or navigation since they demand reflection and frequent adjustments in objects and their attributes; the information-seeking subprocesses proceed in parallel and with varied branching and looping. In most situations, information seekers apply different strategies at different stages of their information seeking. For example, an information seeker may systematically scan a list of titles but opportunistically scan or navigate documents from the title list.

LIMITATIONS OF BROWSING

Although browsing is a natural and often effective information-seeking strategy, it has limitations and costs. These limitations are summarized in Table 2. and discussed below. Browsing requires the information seeker to perceive all potentially pertinent information items and each perception requires time. This "on-the-fly" processing has several consequences. Humans fatigue quickly, especially when doing repetitive tasks. Most studies of information-seeking behavior use subjects in a laboratory setting where performance motivation is influenced by study conditions. The Wiberley and Daugherty (1988) work discussed above suggests an upper limit of 200 items as the number of citations users are willing to scan in electronic environments, although these results were based on surveys of the literature and systems that were designed to discourage examination of long lists of citations. Lee & Whalen (1993) studied an automated "mug shot" system for helping crime witnesses identify suspects. They review consistent results from multiple sources that indicate that performance accuracy falls off rapidly between 100 and 200 image examinations. There is no evidence for how persistent users are when browsing electronic texts, and what factors influence persistence and fatigue. In fact, there is a need to determine a variety of physiological-psychological limits for browsing in various environments.

[Insert Table 2 about here]

Browsing offers severe temporal penalties when the search task is well-defined and a structured search system is available. If indexes or electronic search tools are available and the information problem can be represented in terms of the indexing language or parameters of the tool, then a formal query is likely to be more efficient. Furthermore, if the information space is ordered and a known unit is to be located, application of even a simple binary search strategy will be more efficient than scanning the list in order. The value of the information seeker's time and any time-sensitive charges may also mitigate against using browsing. In addition to the efficiency limitations of browsing in large information spaces, browsing may also limit the coverage that is possible. Thus, inefficiencies may lead to decreased effectiveness. It is mainly for these reasons that professional intermediaries are trained to avoid browsing to save connect charges and personal time.

Users who are browsing primary materials not only have these monetary concerns, but the additional problem of staying on task and not being distracted by "interesting details" of the context. The nature of opportunistic browsing strategies in particular may decrease browsing effectiveness. These strategies may lead to disorientation or distraction, especially in systems that offer few navigational cues. Observational strategies are particularly prone to distraction and navigation tactics can lead to both disorientation and distraction. Garner, Gillingham and White (1989) have described how children and highly-skilled adult readers are subject to "seductive details" while reading non-fiction passages. Their work showed that interesting details about the topic of interest distracted both types of reader from recognizing and recalling the main points of the text. Browsing situations and highly interactive environments must certainly exacerbate this effect. As people browse, they may make discoveries, but they also may be distracted, confused, disoriented, or frustrated by peripheral and tangential information. This is especially problematic is situations where browsing is driven by needs to gain an overview of an area or to extract a sense of gist.

Related to distraction, which wastes time and may mislead information seekers, browsing in rich environments may also lead to information overload. This is especially true in highly interconnected electronic environments where the richness of relationships leads to confusion and overload rather than a deeper understanding of the concepts represented by the system. Clearly, tools and strategies for minimizing distraction and overload are needed.

In addition to affecting efficiency and effectiveness in direct ways, browsing strategies may have subtle influences on information-seeking performance. Browsing strategies are attention-dependent, relying heavily on recognition, and human attention is biased toward what we know and like. Therefore, what we recognize and attend to will naturally be biased. Our observations of users of various search systems illustrates human tendencies to "wishful thinking" as information seekers patch together bits of evidence rooted in availability rather than logical similarity. An artifact of electronic systems is what may be termed "cognitive laziness," a malady invited by systems that allow mindless search or browsing. Users flounder under the illusion of productivity as they enter queries or scan items without reflection or inference.

A related bias may be termed "cognitive inertia," a tendency to continue following paths or lines of evidence rather than examining contrary or alternative directions. Linear text reinforces such thinking and this may be one of the best information transfer qualities from author to reader. However, non-linear environments offered by electronics will mitigate against such information flow and may lead to breakdowns in information transfer.

Technology usually has not been specifically designed to support browsing, thus, there are hardware and software limitations on electronic browsing. The requirement of using a display device has so far limited how users can take advantage of the physical and spatial creativity of designers for retrieval purposes. Intellectual space must be broken across screens or windows which are temporal rather than static--all the pages of a book may look alike, but each page is a distinct physical entity. Since all the information objects used on a display are confined by the same static area and position, users are unable to use physical characteristics such as size, shape, thickness, or paper quality as access points for locating those objects. Since there is good evidence that human memory is episodic (e.g., Walker & Kintsch, 1985), the absence of physical cues to distinguish different intellectual experiences makes us even more dependent on symbolic retrieval mechanisms that are not directly manipulable.

Browsing strategies, like many other problem-solving methods, are subject to a law of diminishing returns. The limitations of browsing parallel the limits of full-text searching. Although there is evidence that adding additional terms to a query statement will improve recall in full text systems (Gomez, Lochbaum, & Landauer, 1990), such improvements must have limits because the salience of the additional words must eventually decrease as various facets of the concept are exhausted--eventually, all documents will be retrieved. Likewise, total reliance on browsing will yield a superficial perspective on a limited portion of an information space. Analytical search, browsing, and careful reading or study all have roles to play in most information-seeking problems.

Thus, the key to using browsing strategies is in selective application to appropriate problem situations and with those search systems that best support it. Users who simply wish to retrieve information in primary systems rather than learn about the context of the information may choose to use intermediaries to minimize costs or avoid the limits of the strategy.

HOW SYSTEMS SUPPORT BROWSING

Support for browsing across objects is built into many environments. Streets and buildings are organized to direct physical movement, and museums and stores invite users to explore personal interests while providing overall organizational constraints. Zoellick (1987) states "Browsing presumes a document collection is structured in some way." (p. 74) Although browsing can be used (and perhaps is the only approach to use) in random collections, organized environments make most browsing strategies much easier. Product catalogs group like items together for easy perusal and comparison and place related groups nearby to suggest other products that customers may purchase. Video controllers and automobile radios provide scanning mechanisms to quickly browse their respective information spectra. Libraries that provide open stacks to patrons shelve books according to a classification system that groups related books together, thus mapping conceptual similarities onto physical proximity on the shelves. A major limitation of this arrangement is that each book can only exist in one shelf location but may be related to many topics. Electronic libraries have enormous potential to overcome this limitation since the electronic document can be instantiated at many places in the classification system, i.e., electronic systems are able to manage many-to-many relationships. Although this potential is yet to be tapped, the probability of locating a specific item by browsing an electronic catalog that assigns items to many main topics is likely to be greater than the probability of locating the specific item by browsing only physical shelves.

Objects such as books also provide structural support for browsing. Tables of contents, chapter and section headings, and indexes are provided to make the organization of a book explicit and allow users to examine contents in non-linear fashion. Journal articles offer abstracts, and organizational structures such as headings and tables that support browsing and non-linear reading. Some journals require specific organizational structures and promote browsing in various ways. Medical journals often provide augmentations such as sidebars, use typographic techniques like placing methodology sections in type font sizes that are smaller than other sections, and require authors to write structured abstracts (Huth, 1987). Electronic documents can provide the same types of support for browsing and offer additional potentials such as hypertextual links that invite browsing.

most browsing is highly interactive and blurs distinctions between the information-seeking subprocesses. Determining an entry point combines search system selection and query formulation; examination, extraction, iteration, and problem definition progress concurrently as browsing continues. A primary challenge for electronic systems is to support this close coupling of the information-seeking sub-processes.

Close linking of input and output mechanisms is one way to parallel and reinforce highly dynamic browsing behavior. Users initiate queries or probes in the database and the system displays results that the user in turn rapidly manipulates. Thus, input mechanisms for query formulation and display control must be provided by the system seamlessly and rapidly so that user control and system display are in synchrony. Various types of electronic information systems support browsing (e.g., bibliographic, full-text, hypertext, image databases, hypermedia) and different techniques have been demonstrated (e.g., text processing, image processing, visual browsers, and information retrieval techniques) that support browsing. However, today's systems provide only one or two features or techniques added on to support browsing as a secondary strategy to analytical search (e.g., allowing users to page or scroll through document sets of titles returned by a query).

Electronic environments offer significant possibilities to support browsing as well as analytical search strategies. In fact, electronic systems that support combinations of browsing and analytical strategies provide the best promise to assist users in applying the most appropriate strategy for the problem. There is a rich collection of example systems that claim to support browsing, illustrating widespread recognition of the need for such systems and that specific techniques and mechanisms can be built. There is a dearth of evidence, however, related to how browsing with these environments affects information-seeking performance, which techniques and mechanisms are most applicable to different search tasks and with what user populations, and how these techniques and mechanisms may best be combined and integrated into general-purpose search systems. Table 3. provides a framework for various features and techniques that may be useful in supporting browsing. Representative examples of systems that illustrate one or more of these techniques and features follow. Examples are organized by generic type of database represented (bibliographic, full-text, and graphical) and by types of mechanisms (visual queries and browsers, information retrieval techniques). These examples are not meant to be exhaustive, but illustrate some of the ways browsing is supported in present systems.

[Insert Table 3. about here]

Bibliographic and online search systems

Early OPAC systems did very little to support subject searching at all, let alone providing support for browsing (Hildreth, 1982). Second-generation OPACs provided for subject access by allowing queries to be posed according to controlled subject headings (e.g., Library of Congress Subject Headings) and in many cases through keywords in the title. Several of these systems refer to one of these access methods as browsing, leading to more confusion about what browsing means. Beheshti (1992) noted that current OPACs lack display characteristics to allow browsing and proposed a graphic interface that represents MARC records visually rather than simply textually. Other OPACs have also used graphic interfaces to aide users in browsing either the bibliographic records or book surrogates (e.g., The Science Library Catalog of Borgman, et. al., 1990; The Book House system of Pejtersen, 1989) provide alternative strategies for search, including hypertext-like browsing capabilities. Figure 6.2 illustrates the graphical display used in the Science Library Catalog. Studies of patron use showed that users used picture browsing (various icons) almost as much as analytical search strategies, and that users were highly satisfied with the interface and their results. OPAC designers have begun to consider browsing as a legitimate search strategy but this is not generally the case for online systems since connect charges encourage off-line planning and analytical strategies. Since bibliographic databases on CD-ROM are not constrained by connect charges, some systems provide browsing support in the form of highlighted query terms, scrollable indexes, and jumps between query terms in text.

[Place Figure 6.2 about here]

Allen (1992) found that perceptual speed, and cognitive abilities such as logical reasoning, and verbal comprehension influenced information seeking in a CD-ROM bibliographic database. The system named free-text searching "browse" and this was the most-often selected option. Based on these results, he recommended that display of linear lists of hits be modified for users who lack perceptual speed, perhaps using hierarchical display of document lists. His results are related to those of Vicente, Hayes, & Williges (1987) who reported the importance of spatial ability to users of hierarchical file systems. These results are interesting because they remind designers to consider perceptual and spatial abilities of browsers, thus adding temper to assumptions about rapid display and graphical representations.

Full-text search systems

Given the quantity of text that professionals must deal with, techniques for browsing full-text systems are of great importance. Tenopir (1985) noted that full-text searching allows users to make immediate judgments about relevance and recommended that careful consideration be given to how text is displayed. She also noted that documents containing more than ten occurrences of a query word or the word occurring in four or more paragraphs had average precision three times higher than overall precision. These results demonstrate that statistical analyses of terms and documents should prove to be useful in helping information seekers select entry points for within-document browsing and in locating promising sets of documents when browsing across documents. Tenopir and Ro (1990) noted that users frequently browsed through articles for background information on topics and required good scrolling capabilities as well as multiple windows to allow multiple documents to be used at the same time.

Three techniques have provided clear advantages for browsing: string search capability, highlighting terms in text, and hypertext links. First, string search is the mainstay of most full-text search systems, serving as the primary way to query and probe the database for promising entry points. Liebscher's work demonstrated user preferences for string search even over hypertext links (1993). Second, displaying highlighted query terms in text is a simple but effective technique to support browsing and information seeking in general. Much of our work at the University of Maryland has illustrated that users made good use of highlighted query terms in electronic encyclopedias, computer science articles, and legal literature. Many users jumped from highlighted term to highlighted term rather than paging or scrolling. Egan, et al (1989) also found that highlighting query terms was a powerful tool for subjects using a SuperBook version of a statistics book. Third, the essence of hypertext is browsing by following links and the development of hypertext systems has led to vigorous interest in browsing techniques from both physical and conceptual interface perspectives. Hypertext systems have supported both explicit links among text fragments through highlighted terms or phrases (embedded menus, Shneiderman, 1992) and implicit links through string search or pull-down menus. Figure 6.3 illustrates how multiple representations are used together with string search and highlighted query terms in the SuperBook system. Figure 6.4 illustrates the embedded menus approach to hypertext link anchors in the HyperTies system.

[Place Figure 6.3 about here]

[Place Figure 6.4 about here]

Other developments have improved the browsability of electronic text. Larger screens, improved type font and size control, and better screen layouts have somewhat reduced the penalties of reading text on a screen rather than on paper (e.g., Hansen & Haas, 1989), although reading text on electronic displays takes longer than reading paper documents (Gould, Alfaro, Barnes, Finn, Grischkowsky & Minuto, 1987). Many full-text systems provide simple ways to scroll, page, jump to the next highlighted term, and jump to the next document in a list, and skilled users may be able to use such features to mitigate reading speed penalties. More importantly, these mechanisms for moving rapidly within texts should facilitate browsing and it remains to be determined how browsing speeds differ in paper and electronic settings.

A variety of techniques to structure and display hypertext links have been proposed and tested. Furuta and Stotts (1989) defined a browsing semantics based on a Petri net model to dynamically constrain how users navigate hyperdocuments. Their Trellis system (see Figure 6.5) can determine whether a link should be active or not depending on previous traversal events. It aids the user in making browsing decisions by displaying the current state of the network and four text windows related to nodes in the net. Thus, browsing can be supported by dynamically making links available according to context and by allowing users to visually see links and overall organization.

[Place Figure 6.5 about here]

Graphics search systems
Improved computing systems have led to the development of image databases. Slide collections, clip art, images of magazines or journals, and video sequences have been made accessible in a variety of ways. Most image collections are searchable only through words or phrases assigned to each image. For example, images in Perseus can be displayed by using an object keyword search that locates all catalog cards that use the target word somewhere in the description of the object (catalog cards have multi-paragraph textual descriptions in addition to basic identification and provenance information). Perseus images may also be located by using multiple pull-down categorical indexes for period, date, creator, and subject (Marchionini & Crane, 1994). Figure 6.6. illustrates the hierarchical indexes available in Perseus. Although these systems provide good access to either analog or digital images, browsing is often tedious due to the time it takes to retrieve and display images. More importantly, these approaches depend on access through textual descriptions rather than graphic properties.

[Place Figure 6.6 about here]

A technique that directly aims at supporting browsing of images is to provide what are called "thumbnails" or miniatures (Nielsen, 1990) of the images. Erickson & Salomon (1991) describe an interface to a large-scale news story database that includes browsing support through what they term "bird's eye views" of the database. These views are miniature displays of articles that show surrounding articles and graphics. Information seekers can scroll through these views and use visual cues to select and display the text or images in a large window. The Perseus Project has developed a thumbnail browser that displays 9 miniature views of graphics objects in the corpus. Users can rapidly scan through these miniatures before selecting and displaying detailed views to study. In Figure 6.7, nine thumbnail views for a vase that has multiple different images allows users to quickly select a view for detailed study or to display the next nine images.

[Place Figure 6.7 about here]

Another approach to supporting graphical search is to allow users to provide a graphical query. A system developed by Garber and Grunes (1992) allows art directors to search an image database by either providing an example image to act as a template or by selecting criteria from menus. This interface was designed according to careful study of how art directors actually search for images and illustrates how important are task analysis and customized interfaces for different problem areas. This approach is an analog of word queries and progress in pattern matching research may make such systems more viable.

A combined graphic and textual approach to chromosome mapping is provided by the ImageQuery interface at the Lawrence Berkeley Laboratory (U.S. DOE, 1992). As part of the Chromosome Information System, this interface allows researchers to specify queries or browse through miniature icons of the images. Buttons for zooming to different magnifications of the image and displaying textual descriptions and parameters are provided in a multiple window environment.

A different approach to accessing image collections is to compress image collections by only displaying the most significant elements or scenes. This is particularly appropriate for film or video. Rorvig, (1993) described a method for "abstracting" video. He proposed a technique for selecting frames from video sequences that results in 700:1 compression for scanning visual documents. Frames are selected according to a set of visual characteristics (e.g. hue, line frequency, angle frequency, etc.) and a statistical analysis that focuses on large deviations in characteristics.

These examples demonstrate that the current state of work on browsing graphics databases is focused on finding representations and mechanisms that are workable, i.e., technical developments. As more technical alternatives are developed, research can progress on understanding how users work in graphics databases and what principles should guide interface design.

Visual queries and browsers

Representing information graphically has long been recognized as a powerful way to aid understanding and dissemination (Tufte, 1990; Tukey, 1977) and there is substantial interface work underway to develop visualization tools for scientific phenomena. Graphical approaches to information retrieval have also be proposed and developed, especially for the case of database management systems (e.g., Herot, 1980; Fogg, 1984). Larson (1986) described a visual approach to browsing well-structured databases. His system supports four types of browsing operations: structuring objects for examination, filtering objects, panning to nearby objects, and zooming to different levels of detail. Users perform these operations by manipulating a graphical scheme that illustrates database objects and relationships. Travers (1989) developed a hypertext interface for the CYC database that uses boxes to represent hierarchical relationships among nodes. Color saturation is also used to reinforce semantic distances among outer and embedded boxes. A technique proposed by Furnas (1986) provides users with ways to focus attention on objects of interest while maintaining a sense of context through peripheral views of the database. An intuitively appealing approach to viewing an information space, his "fisheye" views have been implemented in several systems, including SuperBook (Egan et al).

Perhaps the most extensive collection of graphical techniques have been implemented by Card, Mackinlay, and Robertson at Xerox. Based on a general cognitive model of human performance and using 3D graphical techniques, their systems allow users to view and manipulate significantly more information per screen than other interfaces. Their perspective wall (see Figure 6.8) displays three sides of a rectangular solid in perspective so that the "front" wall shows a large rectangular array of cells and the two "side" walls angle back in perspective with diminishing arrays of cells (Mackinlay, Robertson, & Card, 1991). Users can easily slide walls to bring arrays of interest into front view. Another technique (see Figure 6.9.) uses "cone trees" to represent large hierarchical structures that users can rotate and manipulate to find information (Robertson, Mackinlay, & Card, 1991). In one example, 600 directories comprising 10,000 files were represented on a single display. These techniques offer promising directions for systems that support browsing through high-performance displays, although they require sophisticated workstations.

[Place Figure 6.8 about here]

[Place Figure 6.9 about here]

Another collaborative effort that involves multiple researchers is the work of Shneiderman and his colleagues at the University of Maryland Human-Computer Interaction Laboratory. This work applies direct manipulation techniques to information seeking under the name "dynamic queries." Systems that support dynamic queries provide sliders for setting interval data values and hierarchical menus to select categorical data values, and immediately display the results of settings. A museum system allowed users to quickly select types of archaeological sites, historical periods, dates of visit, and participation costs (Weiland & Shneiderman, in press). Another implementation (see Figure 6.10) allowed users to manipulate sliders for cost, location, number of bedrooms, and other house features to visualize what homes were on the market (Williamson & Shneiderman, 1992). Each change of a slider immediately resulted in an updated map with available homes plotted. This approach is intuitive and pleasing to use and users in their study consistently conducted faster searches with this system than with paper or natural language interfaces. This type of system offers good potential for problems that lend themselves to graphical representation and have ordinal criterion values. Another representation called tree maps uses horizontal and vertical slices and color to display large hierarchical structures such as directories of files (Shneiderman, 1990). Using this technique, directories of over 1000 files can be browsed for various parameters such as size and file type.

[Place Figure 6.10 about here]

Information retrieval techniques and mechanisms to support browsing

Although traditional information retrieval research aimed at optimizing analytical search, often eschewing informal browsing strategies, many of the fruits of information retrieval work may be applied to systems that invite and support browsing across documents. The most obvious information retrieval advance to support browsing is relevance feedback. As described in Chapter 2, relevance feedback adjusts user queries based on assessments of retrieved documents. Harmon (1992) in a review of relevance feedback, reported that variations of relevance feedback doubled the performance of systems she tested. Most potent were techniques that expanded queries rather than simply reweighing terms in the original query. She found that expanding with only twenty selected terms was superior to expanding with all terms. This seems to be another example of a law of diminishing returns in specifying terms in queries--more terms is generally better but after a point, too many terms either lead to intolerably low precision or do not yield any additional recall. Systems that include relevance feedback interfaces include Salton's SMART system (Salton & McGill, 1983), Frisse's Dynamic Medical Handbook (Frisse & Cousins, 1989), Williams' RABBIT system (Williams, 1984), Walker's OKAPI (Walker, 19), and Erickson and Salomon's Desktop Information System (Erickson & Salomon, 1991). The Wide Area Information Server (WAIS) interface has become popular for searching Internet sources and uses relevance feedback (Kahle & Medlar, 1991). Relevance feedback is like browsing in that it depends on feedback as the user explores and probes the system interactively. The underlying mechanisms of relevance feedback depend on statistical or probabilistic techniques, but these may be hidden from end users. For sophisticated browsers, however, future systems could allow users to manipulate the underlying relevance feedback parameters (e.g. how terms are re-weighted for query modification, what terms are used for query expansion).

Another collection of information retrieval techniques that can underlie browsing interfaces are clustering algorithms. Salton (1989) noted that browsing is promoted by the proximity of similar documents that results from clustering (see Rasmussen, 1992, for a discussion of different approaches to clustering). Cutting, Karger, Pedersen, & Tukey (1992) report a technique that applies clustering to allow browsing of large databases. Their "scatter/gather" technique is specifically designed to assist users in exploring open-ended or ill-defined problems and to work in conjunction with analytical approaches as needed. They contrast searching an index in a book which requires a term or query and examining a table of contents which gives a sense of what types of questions the book may answer, and base the browsing part of their scatter/gather approach on the latter. The basic approach is to use a clustering algorithm to scatter the collection into document clusters and report cluster summaries to the user. Based on which clusters the user chooses for further study, the system gathers these clusters and re-clusters them to produce a smaller scattering. The process continues until there is a small number of documents that can be examined individually.

The increasing computational power of computers has prompted some researchers to combine sophisticated information retrieval techniques and graphics display techniques to support browsing. Korfhage (1991) has developed a technique for reducing the high dimensionality of a document space to graphically display the space according to user queries and profiles. The queries and profiles are reference points from which all other documents are distanced according to similarity metrics. High-resolution display of the distance space allows users to view clusters of documents according to their own queries and profiles rather than some absolute clustering based on frequencies and co-occurrences alone. His GUIDO system illustrates a pioneering approach to personalized graphical views of semi-structured document collections (see Figure 6.11).

[Place Figure 6.11 about here]

Lin (1993) proposed a framework for presenting two-dimensional displays of document collections that is based on a neural network clustering of documents. He used Kohonen's feature map algorithm to cluster documents and prepare a graphical map that represents topic importance by area size and topic similarity by proximity. Instead of ranking documents in linear fashion based on some similarity measure, the algorithm maps the high-dimensional vector representation of titles in the collection to a two-dimensional map that preserves similarity through geometric distance. The resulting map defines areas for topical clusters that are juxtaposed to semantically similar clusters and whose areas are proportional to the frequency of occurrence in the corpus (see Figure 6.12). Lin compared his computer-generated maps to maps produced by humans and found that subjects could use the different maps equally effectively to locate titles. Both the computer-generated maps and those produced by humans were superior to a random map. Although generating such maps are computationally expensive at this time, such maps offer good possibilities for supporting the informal information-seeking strategies collated under the term browsing.

[Place Figure 6.12 about here]

In this chapter, browsing strategies were defined and contrasted to analytical strategies. Four different browsing strategies were defined according to how the various information-seeking factors interact during information seeking. Limitations of browsing were discussed and the importance of the environment was illustrated with examples of systems that support various aspects of browsing. Clearly, browsing strategies and analytical strategies complement one another and systems should support the preferences and abilities of diverse users. In the next chapter we examine how such systems can be developed.

Chapter 6 Notes.

1. Skimming is not considered here since it is particular to text, whereas scanning is independent of media.

2. Norman (1988) discusses the importance of matching "knowledge in the head" and "knowledge in the world" for successful human-computer interaction.