Chapter 7. Designing Support for Browsing: A Research and Development Perspective

Many specific system features have been shown to invite and support browsing as an information-seeking strategy. We are beginning to acquire a set of techniques that define what is possible in designing such features. Determining what is optimal for different users, tasks, and settings requires systematic testing of techniques across the range of information-seeking factors. Because browsing requires users to coordinate physical and mental activities, systems that support browsing must solve technical and conceptual problems in a coordinated fashion. Technical challenges such as the computational power needed to manipulate huge vector spaces on-the-fly and display problems such as resolution limitations, refresh and scroll rates, window sizes and juxtapositions are difficult enough in isolation, but must be coordinated with other technical problems such as mechanisms for selection and control of information and conceptual problems such as what are the best representations of meaning for specific information items and what should be displayed at what time, in what form, and at what level of granularity. Programs of research are needed that address the technical problems of designing interaction styles for browsing, that determine the physiological and psychological boundaries on browsing activities, and that test various representations for browsable information. These are technical, user, and organizational areas respectively. Although different researchers and groups typically specialize in one of these problem areas, ultimately, support for browsing will depend on integrating results from all three.

Developments in physical interfaces, especially in integrating multiple devices (e.g., Jacob, et al, 1993) and progress in designing screen displays, metaphors, and dialogues will be informed by answers to questions about human capabilities and needs. Consider the following questions, for example. What are the perceptual and cognitive limits of recognition for symbols and images? What temporal-accuracy tradeoffs exist? What are the implications of these tradeoffs for immediate use of information versus retention? How many windows can be concurrently scanned for different types of targets? What are the temporal and spatial limits on concurrent scanning and monitoring? Answers to these and similar questions about physiological-psychological interactions will help designers to develop or select better physical and conceptual interfaces. Although we have good measures of perceptual abilities in reading (e.g., Rayner, 1978), text editing (e.g., Card, Moran & Newell, 1983), and other cognitive tasks, our understanding of how perception is used in different browsing strategies is quite limited. Additionally, there are mixed results regarding ways that humans can rapidly recognize and understand information; research in cognitive processing for all types of information activity remains a long-term problem.

Of particular interest here are conceptual interfaces in general and the ways that information is represented and managed by browsers in particular. Since browsing is more data-driven and less goal-driven than analytical search, it is critical that data be well-represented and easily managed. By *well-represented*, we mean that appropriate data types are provided, information is provided at appropriate levels of granularity, and information is displayed succinctly and aesthetically using organizational structures that are appropriate to the task and user. Also, because the same information can be useful for different types of problems and by different users, these representations should be flexible and adaptable. By *easily managed*, we mean that usable and effective mechanisms are provided for access (specifying representations), display (organizing representations), and manipulation (linking, filtering, aggregating, and analyzing information representations). In addition, it is essential that systems provide instant response as users apply these mechanisms and that tools for cutting, pasting, saving, and tracing search history are available.

The representation problem is fundamental to information science and addresses how information is organized and articulated mentally and in the world. In particular, we are concerned with providing a closer mapping between information seekers' mental representations of an information problem and representations for information objects in a search system. Since cognitive science offers a variety of theories about how humans represent knowledge, it is prudent to take a design approach that provides alternative representations and allows users to select from them according to their personal style, experience, and the problem at hand. Thus, the first challenge of design is to identify and create useful alternative representations for information. The cost of such an approach is that users must remain active and make choices as information seeking progresses. This tradeoff leads immediately to the second design challenge--supporting the user in understanding and manipulating the alternative representations. This chapter considers the problems of representation and mechanisms in systems that support browsing strategies and then calls for interfaces that support highly interactive browsing strategies as well as analytical strategies.

REPRESENTATIONS

Ideas find many forms of expression and the problems associated with representation of knowledge are central to research in information science. Whereas artificial intelligence is concerned with creating machine manipulable structures and mapping concepts onto those structures, information seeking is most often concerned with locating and interpreting existing representations that are manipulable by humans. Information seeking includes asking people for information since seeking information from an expert requires locating the expert and interpreting representations generated through conversation. These representations include objects such as books, articles, photographs, discrete utterances, and films. These human-manipulable representations vary greatly in scope, time required for human processing, and physical space required for storage. To support the information-seeking process, information systems are challenged to:

- provide mechanisms for access, display, and manipulation of information,
- display these representations to facilitate interpretation, and
- support extraction and manipulation of information from them.

Systems that allow information seekers to use browsing strategies provide tightly-coupled representations and mechanisms so that the information subprocesses proceed in parallel and problem definition is maintained as a continuously active subprocess.

Simple systems store representations in their entirety arranged in some specific order. For example, a file of incoming correspondence arranged chronologically in a file cabinet provides access by date, displays the entire document a page at a time, and depends fully on the information seeker for extraction and manipulation. Some systems store only surrogates, abbreviations, or summaries instead of full representations. For example, a computerized transcript file for grades summarizes many tangible representations of performance. Such a system may support access by some key (e.g., last name or identification number), display records one student at a time, and may support standard printing formats and possibly some computational functions across parts of the file.

More sophisticated systems separate the full representation from surrogates for these representations. For example, a card catalog provides access to books through author, title, and a few subject headings; displays a full bibliographic record on one or more cards (screens in the case of OPACs), and provides little extraction or manipulation help beyond screen dumps. Current archival practice provides at least three main levels of representation ranging from directories (e.g., the Directory of Archives and Manuscript Repositories in the United States), holdings for a specific archive (e.g., record groups, series and subseries in an archive), and specific record groups (e.g., boxes, folders, and items). Still other systems store representations in their entirety, provide access through multiple indexes, and store pointers or entire representations that are related. For example, consider a system that stores the full text of a book together with reviews for the book and a citation file for other works that cite the book. Access may be provided at many levels, display may be simply scrollable or highly controllable by the user, and extraction may go beyond cut and paste to grammatical analysis.

These examples illustrate the range of possibilities for representation and designers of systems that support browsing strategies must consider both the type of collection and the types of surrogates in meeting the challenges of access, display, extraction and manipulation. Collections may contain full representations (primary information) or only surrogate information (secondary information), or various levels of each. Electronic systems have increased the likelihood of including both secondary and primary information but designers have yet to fully grapple with the distinctions of across-document and within-document browsing. Surrogates may be descriptive or semantic and either type may possibly be extracted automatically or crafted by hand.

A system can represent an item in full or in part. A book can be represented simply by its title or in its entirety depending on the needs and resources of users. Clearly, some levels of representation are useful for locating a specific item and others for studying or browsing within that item. A system can also present alternative representations for the same item. For example, software comes in different versions for different platforms, textual documents can be represented as simple ASCII files or as formatted text with different fonts and styles, and data sets may be displayed in different tabular or graphical forms. Similarly, a text can be represented by its words or by SGML codes or page layout pictures, depending on what is to be done with the document. Information objects can also be represented with added value--with extensions and links to objects beyond themselves. A fundamental difference between paper and electronic systems is how electronic systems blur the distinctions among these different types of representations. Since it was earlier argued that browsing blurs distinctions between information-seeking subprocesses, the combination of electronic environments and browsing makes for highly fluid and amorphous information seeking.

Systems aid browsing across documents primarily by providing partial representations or surrogates. In across-document browsing, salient surrogates for documents aid the user in rapidly deciding whether to examine the document more closely or not. There are many such representations used in manual libraries and databases, and electronic systems offer new classes of surrogates to aid users in browsing collections. These representations can be classified as descriptive or semantic.

Descriptive surrogates identify the attributes of items. Consider the following attributes. Creator attributes such as author, translator, producer, director, animator, editor, conductor, performer, programmer, designer, or artist can be useful in assessing relevance of items. Creator attributes are subsets of information objects; they are directly extractable. Additionally, attributes associated with creators, such as institution, cultural background, historical environment, philosophical perspective, or methodological approach may be useful in judging whether to examine objects more carefully. Notes on creators are especially useful when the information seeker is not an expert in the domain and does not know the people active in the field. These attributes may be part of the object or added value for the purposes of retrieval. Meta information such as medium of representation (e.g., text, video, computer graphic, sound, multimedia, etc.) may be a useful attribute for quickly identifying or eliminating large portions of a collection. Physical attributes such as size, number of images, running time, and provenance may also offer indicators of potential relevance and are especially useful in retrieving physical objects. Other descriptive attributes include copyright date, version or edition, and vendor or publisher.

Semantic attributes aim to represent the meaning of objects in a collection. For text collections, titles and subject headings are perhaps the most valuable indicators of possible relevance and many systems list documents by title as a default key. For image collections, thumbnail miniatures, low-resolution surrogates, or video abstracts may serve as browsable surrogates for objects in the collection. The main problem with semantic representations is capturing the "aboutness" of an information object (Saracevic, 1976). A major area of information science research is devoted to discovering methods for representing meaning of objects in a collection and in determining to what extent these representations may be automatically extracted. It seems clear that for complex objects such as books, articles, and films the best descriptors are tied to the information problem rather than some generic classification (Soergel, 1985). This is why systems that invite browsing must represent information according to needs expressed by individual users. An ongoing research problem is how much preprocessing can be done and how much must be done on-the-fly during user-system interaction. Systems that support analytical search strategies will be able to use more preprocessing and those that support browsing strategies require more on-the-fly processing.

Descriptive and semantic representations provide important cues to information seekers as they browse across document collections. Both types of representation can also be useful for within-document browsing but current systems do little beyond paging or scrolling to support within-document browsing. Systems that support rich descriptive and semantic indexing should allow users to select from many documents and to move into, around, and back out of those documents in seamless ways. Present systems require users to execute a discrete set of actions using different input/output mechanisms at each step. Queries are typically specified according to a query language or form screen and a list of titles or other surrogates are then displayed on a different screen/window. Different menus or commands are then used for browsing the retrieved list and selecting an item for display. This results in yet another screen/window with possibly another set of commands or menus to browse within the item and return to a previous step. In electronic environments, representations for collections and specific objects should be represented and controlled in common and compatible ways.

Representations are by definition structured to support processing, and these structures can be used to support within-document browsing. Some of these structures are provided by the creator and some are added by publishers or editors to produce a final, "published" document. Textual databases such as books or articles may have abstracts, tables or contents, section headings, paragraphs, sentences, quotations, footnotes, indexes, sidebars, tables and figures, appendices, and reference lists that browsers may find useful for extracting gist or focusing in on specific details. Electronic environments allow users to display any or all of these structural components in a variety of orders and juxtapositions. Users, should be able, for example, to quickly display and scan only footnotes or tables according to their needs.

Software also provides many structural elements that could be useful to those seeking information in the code. Global and local variable declarations, functions, procedures, run time ranges, print ranges, and comments all help users understand the program. Likewise, documentation, tutorials and user support services assist end users in understanding and using programs. Modular programming techniques and the integration of tutorials and online help have complicated examination of software code, making a single application program more like a collection of documents and thus blurring within and across document searching.

Artifacts of the electronic format can also provide structural enhancements to aide users even more. Standard general markup language codes (SGML) used for typesetting purposes may be useful representations in electronic texts and hypertexts, allowing users to search or scan according to code (Raymond, 1992). For example, searching for all quoted passages or for text that is bold-faced may assist users in locating specific types of information according to their particular needs. The visual layout of documents may also aid users in identifying parts of documents for browsing. One prototype document image analysis system uses page layout formats for specific journals to segment text for recognition and browsing (Nagy, Seth, & Viswanathan, 1992). The system allows users to select documents based on pictures of page structures. Thus, information seekers can use their visual memory for journal layouts to aid in browsing. Although the prototype requires hand-crafted descriptions for hundreds of production rules for each journal layout grammar, the system illustrates how spatial structure cues can be added to electronic search systems to assist information seekers in identifying and selecting potentially relevant documents during across-document browsing and then to move to particular sections of the text itself. These examples illustrate how structural aspects of documents can enhance within-document browsing and may be extended to across-document browsing as well. Additionally, they illustrate how electronic technology not only allows new levels of representation for authors and their collaborators, but allows compilers and individual users to manipulate those representations and add new ones.

In addition to taking advantage of structures within paper or electronic documents, representations may be enhanced, alternatives may be provided, and new representations may be added to support browsing. One type of enhancement is to provide new representations that augment abstracts and other compressed expressions of gist. Florance (1992) has proposed extracts for the medical literature to aide physicians in making clinical use of the medical research literature. Her extracts are frame-based representations for epidemiologic, diagnostic, treatment, and learning information in articles and allow a large community of information seekers to examine the literature from any of these specific perspectives depending on their immediate needs. These abstracts are similar to the concept of database views; they provide problem-oriented and personalized entry to a database.

Value may be added to documents or databases in a number of ways. Reviews, apparatus criticuses, commentaries, forwards or prefaces, movie out takes or montages, demo modes, and sample chords provide information seekers with additional information for decision making. In traditional environments these added values are carefully crafted as part of the production and marketing process or as part of the growth of a scholarly field. Other forms of added value include notes on allusions or metaphors, bibliometric features such as documents that cite the document of interest or sets of documents that are co-cited with it. Hypertext links and network gateways serve as obvious electronic connections to related materials that add value to specific information objects. In fact, the links or network nodes themselves provide useful interpretive value since they make relationships explicit. Special interest group lists provide structures for extended conversations among various participants and specific argument threads may be isolated and linked together according to the needs of users. Consider a video collection. Descriptive representations that occur as part of the entire representation include: title, scenes, visual effects, and credits. A form of video translation is illustrated if the scenes are shown as compressed (e.g., Quicktime) displays. Value may be added in many ways, including: descriptions, reviews, awards, promotional "rushes," and specially constructed montage abstracts. Although electronic representations are not necessary for most of these attributes, benefits of electronic technology accrue by making it easy or automatic to create additional and alternative representations.

Electronic environments offer end users possibilities for easily adding value. For example, in textual databases, users can form concordances, run grammar analyses, compute readability indices, or run clustering analyses to better examine documents according to their specific needs. Spatial management systems (geographic information systems) allow users to produce their own thematic maps rather than simply accessing and viewing existing maps (Laurini & Thompson, 1992). In hypermedia databases, users can produce simple node-link ratios, link plots that represent the degree of connectivity across all nodes (Bernstein, Bolter, Joyce, & Mylonas, 1991), or in-out indices that indicate which nodes serve more as indexes and which more as references (Botafogo & Shneiderman, 1991). Nielsen (1990a) has suggested combining word frequencies within nodes (intrinsic weights) and nodes directly linked to those nodes (extrinsic links) as indicators of relevance to guide hypertext navigation. These added values could be provided by the creators, but electronic systems could provide end users with tools to add these values themselves.

The possibilities for added values grow in shared environments. For example, in a library, borrowers could be asked to provide subject headings, abstracts, or commentaries for items and these could be made part of the total document. Books that borrowers underline or highlight offer benefits to some users but may distract and annoy others. Defacing books is moot in an electronic environment since comments and highlights can be saved separately and linked as desired by the user. However, the usefulness of other users' commentaries or markings still depends on who those users are and what their information problems and goals were as they entered their commentary. For this reason, user contributions to a document must be accompanied by profiles or contextual information about those users and their goals in using the materials--in turn, leading to privacy concerns. Electronic environments make these types of added values in shared environments much more feasible technically. If social and economic issues are solved, users can choose to display or hide other users' contributions or select them according to various sorting and filtering criteria.

This discussion illustrates the many levels, alternatives, and added types of representations that designers and users must consider. A book can be examined at a coarse level (e.g., the title alone) or fine level (e.g., the full text) depending on whether an information seeker is browsing across documents or within the document. Document structure (e.g., table of contents, headings, etc.) may be used for examination and additional related information (e.g., a book review) may be sought. These various representations may be useful to information seekers in locating a book in a large collection, determining whether to examine it more closely, and in executing careful examination for relevant information. In general, manual environments provide basic support for access, good display characteristics, and little support for extraction and manipulation. Electronic environments complicate this situation in several ways. Electronic environments support rich across-document browsing by providing multiple descriptive and semantic surrogates. They are presently quite weak in supporting within-document browsing but as more ways to take advantage of document structures are discovered, this may improve. Since the same physical display device is used to view the database, the text, and any related items, electronic environments tend to blur the physical distinctions of across and within-document browsing. Although it is advantageous to click a mouse and immediately display the full text or a cited article, the interface must provide clear demarcations between these representations as well as providing the selection mechanisms themselves. The trend toward interconnected systems offers opportunities and challenges for linking representations according to a variety of relational conditions and managing the representations themselves in addition to the content of interest. Most importantly, electronic systems can provide alternative representations and access mechanisms that empower users to manipulate, analyze, and re-represent information according to their particular needs. Since electronic systems also may allow documents to be changed by users, entirely new classes of problems, especially for representations that depend on preprocessing, are created. Presently, users find themselves in roles ranging from overwhelmed kids in a candy store to mystical information alchemists. Controlling the richness of representations requires interface mechanisms, and at present, we have little guidance as to how to best design such mechanisms.

MECHANISMS FOR MANAGING REPRESENTATIONS

Electronic environments offer alternative representations that can be controlled by users. Users can choose what levels of detail to display and what types of added value to include during information seeking. All of these structural elements are useful for specific types of information-seeking conditions, but having them all available at once may actually lead to overload and confusion. Consider, for example, a hymn book that provides a dozen indexes for various types of users and purposes and an electronic version where these preprocessed indexes are available and users can create their own on-the-fly. The possibilities increase for spending all one's time in pointer information rather than ever getting to the primary information. Clearly, a major problem is how best to organize and control these representations. Electronic environments on the one hand exacerbate the problem by providing rich sets of options, and on the other hand offer new mechanisms for managing the representations.

In addition to other mechanisms for storage, transmission, and help, systems must provide a range of mechanisms for managing representations. From the end users' point of view, mechanisms include widgets and tools for:

- specifying constraints on the universe of information available (e.g., queries) in order to access relevant information;
- managing screen and window displays; and
- linking, filtering, aggregating, and analyzing information representations.

The following discussion is meant to provoke questions and raise research issues as much as to provide guidance to designers. In the following discussion, a graphical user interface workstation connected to Internet resources over a high-speed channel is assumed. Thus, the universe of available information includes large, well-structured databases, data directories, semi-structured files, and specific objects such as programs and electronic mail messages. Three general classes of mechanisms are discussed that allow users to search and browse in such an environment. These classes are based on visual metaphors for different problem-solving tactics people apply within the information-seeking process. The three classes of mechanisms are not mutually exclusive and may be used together during information, but each offers specific advantages to the information seeker. Probing mechanisms for quickly locating specific items such as words or records provide advantages in well-defined information-seeking situations and for identifying general neighborhoods for browsing. Mechanisms for zooming and panning in, out and across levels of detail offer great potential for users to smoothly go across and within documents to meet their information needs and manage displays. Mechanisms for filtering and restricting information allow users to limit levels of detail or types of documents and provide good functionality for manipulating information.

Probes

Probes are specific investigations launched by information seekers; these investigations range from simple queries in a database to multiple-step examinations of segments of the information space. Probing is a form of hypothesis testing--hypotheses are formed, posed to the system, and results reflected upon. In the simplest case, hypotheses take the form of questions rather than predictions, for example, asking the system whether a word matches words in the database. String search has repeatedly proven its usefulness in information seeking and represents the first significant advantage of searching in electronic environments. String search allows users to probe a database or document for highly specific units--word fragments, words or phrases. Full-text word search mechanisms are similarly useful for testing matches in textual databases. Full-text search allows users to launch highly focused, complex analytical queries or highly interactive "guess and go" probes. Analytical probing of a database is the basis for most online systems today and highly focused, complex queries allow extremely precise predictions to be tested in huge databases. The guess and go approach was observed repeatedly in our studies of children and adults using different full-text environments. Information seekers use a word or phrase to rapidly find matches in document collections and then use browsing strategies to examine the context of those matches for relevance. This simple approach has proven remarkably effective in small and medium-scale collections and should certainly be provided in text-oriented systems.

We can imagine, however, more sophisticated types of probes for large collections and more complex problems. Systematic queries to a bibliographic database and subsequent merging of result lists is a common way of probing databases to determine coverage or for discovering neighborhoods or research fronts. Furthermore, relevance feedback may be used. For example, words or phrases from existing objects could be "picked up" using a pointing device and launched into a database. Queries may also be restricted to specific types of sources for similar purposes. For example, the MEDLINE system allows queries to be limited to review articles. This functionality is considered more fully in the section on filtering and illustrates how the different mechanisms overlap and work together in practice.

Since Boolean queries cause users so much difficulty, another direct manipulation approach may be used for established systems based on inverted indexes and exact match retrieval. For example, allow users to specify single term queries that return separate windows of document sets. These windows can then be dragged to merge or overlap. In the case of merging, one window can be dragged to the start or end of another and the records combined and sorted. In the case of overlap, one window can be placed on top of the other and only overlapping records displayed. In both cases, easy undo features must be provided so that users feel comfortable to quickly and directly merge (OR) and overlap (AND) document sets in exploratory ways.

Instead of focusing a probe on a single database, it can be broadcast. A single query can be sent to multiple databases, for example, the WAIS interface (Kahle & Medlar, 1991) allows queries to be sent to any number of databases. It may be as informative to study redundancies as well as unique results from the different databases. Queries posted to listservs are good examples of probes that use technology to broadcast to a specific community of people who may be able to assist the information seeker in solving their information problem.

More sophisticated probes may take the form of agents or "knowbots." Agents have simple inference capabilities and are provided with more than a word or phrase by the information seeker. An agent launched for a specific purpose may carry a form filled out by the information seeker which serves as a profile of the information need. When matches are found, inference rules are applied against weights and conditions in the form to determine whether to add the matching object to the report of results. More complex agents can be imagined that remain active and regularly report results. These types of probes may be considered as selective dissemination of information agents that keep people appraised of developments in specific areas of interest. Further speculation leads to anthropomorphic icons that engage the information seeker in dialogues to determine rule setting parameters, conceptual clusters of words and phrases identified and weighted as a kind of dynamic thesaurus, specific constraints or value assessments for information sources, and highly relevant "pearls" that act as exemplars. As is always the case with imagination and speculation, the engineering problems are significant. Building user models for such agents is a huge challenge since efforts to date have been stymied by defining what are the essential components of generic models, when they should be applied, what level of granularity should be included, how they should be evaluated, and human adaptability to situations in general (Allen, 1990; Daniels, 1986). Furthermore, management of agents will itself require time and effort on the part of users. Additionally, many of the serendipitous fruits of information seeking may be lost if information seekers limit their time to profile-specific documents.

Probes may be supported by many different interaction styles. Command languages are most obvious and exemplify the well-known tradeoffs between difficulty in learning and efficiency in use. Natural language processing seems obvious for asking questions but may promote imprecision and sloppiness in stating hypotheses. It is unlikely that query articulation will benefit much from verbal or typed language capabilities, but verbal dialogues that clarify and define the problem seem promising. Form fill-in and query-by-example interfaces provide rich context for designing probes and will continue to find wide application for well-defined problem situations. Menu systems severely limit expression but offer information seekers explicit categories or vocabulary for shaping a probe. These are particularly helpful if domain-dependent hierarchical sets are provided such as the object-search menus described in the Perseus system in the previous chapter. Advantages of direct manipulation for maps and other physical/spatial domains are obvious. Direct manipulation interfaces may also allow users to directly construct probes by "picking up" words or objects, placing the probe into sources, and possibly watching activity, but it is unclear what benefits other than physical ease of selection and placement accrue. It is even less clear how useful will be virtual reality mechanisms that allow users to launch themselves as probes into the information space.

Probes seem most appropriate for across-document searching and finding neighborhoods or entry points for browsing tactics, although some subjects in our studies of full-text searching said they would like to do string searches within a retrieved document. Probes are the main mechanism used in today's retrieval systems since in the simplest case, they are queries. They also lend themselves to the full range of interface styles. Probes are most useful in query formulation and execution processes and must be augmented with display mechanisms and manipulation mechanisms.

Zooms and pans

A natural metaphor for displaying different levels of representation is the camera that allows users to zoom in and out of levels of detail and to pan across a scene. If representations of information objects are strictly nested, it makes sense to move up and down the levels of representation by mapping the screen to a lens that may be zoomed in and out at different levels of detail. A *zoom* is a mechanism for changing display by moving up and down levels of representation (e.g., from document list to a specific document). A *pan* is a complimentary mechanism for changing display by moving within a level of representation (e.g., form one part of a text to another).

The zoom metaphor may map the lens onto physical levels of detail or onto an alterative hand-crafted or computed representation. Zooming within physical levels is illustrated by the RightPages system that displays journal covers, allows users to zoom in to tables of contents, and then to full text (Story, et al, 1992). In their system, a “zoom in” and a “zoom out” button are provided to allow discrete moves among the three levels of representation. Many other application packages provide zoom mechanisms for moving in and out of physical levels of detail (e.g., Virtual Museum (Apple, 1992). Zooming has also been used as a metaphor for moving across alternative representations or views of a specific level. For example, the European Space Agency's Information Retrieval Service offers a zoom command that allows users to examine subsets of retrieved sets ordered according to frequency of occurrences for various attributes such as subject heading or date (Ingwersen, 1984). Many other systems (e.g., Microsoft Bookshelf) use zoom to move from hit list to specific entry and back out.

For representations that are not hierarchical but rather alternative or added value, the pan may be appropriate. Just as a camera can zoom in or out, it may be made to move laterally rather than in and out. Panning facilities would admit paging within a document or moving to a related document. Zooming and panning provide a powerful visual metaphor for many browsing activities. Systematic observation is a powerful problem-solving strategy and underlies this set of control mechanisms that allow information seekers to examine a field of view and move across different fields of view.

A particularly interesting interface that uses zooming as an interaction metaphor is the PAD system (Perlin, 1991) and the PAD++ descendant (Bederson, 1994). This system treats the screen as an infinitely scalable surface. Notes and addendums can be squeezed between any other objects in recursive fashion by zooming the screen in and out. This is especially useful for nested data such document citations, outlines, texts, and notes. The continuously zoomable interface allows the user to select different views easily. It also illustrates an amplification of human eye-hand coordination by allowing graphical tools to be applied at any level of magnification. Common paint programs allow two or a few levels of drawing detail (e.g., moving from full size, to 50% size to "fat bits" scales) but PAD supports a continuous and potentially infinite set of levels.

Zooming and panning lend themselves to direct manipulation interaction styles. Users apply perceptual feedback to determine where to focus and thus have direct control over the environment. This control is costly, however, in very large databases since so much information must be perceived, if only briefly. Zooming and panning may thus be best used in small or medium scale environments or after analytical or guess and go strategies have identified a promising neighborhood. Several interface design problems related to how zooms and pans are implemented must be solved. For example, should a zoom button work like a fast forward button and provide continuous movement through different levels or should it act in a more discrete fashion and step through the levels incrementally? The discrete implementation is technically easier, but what are the costs to the user? Similarly, how much of a level should be shown and how long must it be displayed for users to make decisions about zooming further?

Panning is particularly good for within document browsing and zooming provides natural opportunities for users to move easily across reasonable numbers of documents as well. As boundaries among specific documents, ancillary materials, and document collections continue to dissolve, effects on users must be studied to determine whether interfaces that support zooming and panning should provide clear distinctions as objects are displayed.

Filters and templates

Probes are particularly good mechanisms for accessing representations and zooming/panning are particularly good managing the display of representations. *Templates* are mechanisms that constrain what is probed or displayed and are particularly useful for filtering information. Since both filters and templates are used as metaphors for the function of filtering information, both will be considered here.

Information seeking for the purpose of learning or accretion of knowledge or for open-ended tasks is not so much a problem of finding information as it is of filtering out irrelevant information. Even rather narrow fields of interest generate huge volumes of information that may be relevant to the needs of experts in those fields. The challenge is to use one's knowledge of the field to ignore most of the information available and focus on that that is most pertinent to one's needs. Much of the work on information filters (e.g., Lai, Malone, & Yu, 1988; Malone, Lai, & Fry, 1992; Fischer & Stevens, 1991) has focused on ways to manage incoming streams of information such as electronic mail or Internet broadcast news postings. Belkin and Croft (1992) have argued that information filtering is conceptually congruent to information retrieval and that techniques for retrieval should be useful for filtering. They argue that principles for information retrieval apply to incoming streams of information as well as to the usual purpose of searching out information in different locations. It is natural to use filtering metaphors such as light and sound filters or chemical solution filters to represent the filtering process and the interface icons that initiate them. For example, Young and Shneiderman (1992) implemented a filter-flow interface for database queries that proved superior to a SQL query language for novice users. Users controlled the flow of water through a system by selecting attribute values as filters and arranging them in linear (Boolean AND) or parallel (Boolean OR) arrangements. The visual effects and attribute selection presumably combined to improve user understanding and performance. This system illustrates the visual use of the filter concept to augment probing mechanisms.

We can imagine other metaphors for filtering, for example, a recycling metaphor. Consider a toolbox organized around tools for sorting and organizing objects. The toolbox might contain sieves for separating coarse and fine objects; shaker boxes used to order items according to some weight (e.g., frequency of term occurrence, length of document, cost, etc.); a set of specialized magnets used to grab all items with a very specific property (e.g., a Boolean query, all photos taken before 1936, etc.); catalysts that interact with items to eliminate unwanted objects, highlight those with specific characteristics, or change representation (e.g., dissolve all items from a specific journal, change all tabular data to bar graphs, etc.); and conveyor belts to support hand sorting of various objects (e.g., rapid visual browsing of documents).

Just as simple inference capabilities could be added to probes, they could be added to filters as well. Instead of active agents, these filters may be more readily compared to sentinels. Sentinels could "listen" for specific classes of information, reject most and classify other information units according to its rule base. Although sentinels are preferred to more active and "intelligent" agents, the same problems of specifying profiles and managing one's sentinels become an issue. Clearly, intelligent probes and the sentinels used as filters differ only metaphorically; functionally they are identical and exemplify the similarities of selective dissemination of information and information filtering.

An alternative metaphor to filters is the notion of templates. A template for drawing guides the pencil by restricting where the instrument may go. From another perspective, a template only allows viable objects or relationships to show through. For example, a spreadsheet template constrains the worksheet by assigning labels, values, and relationships to each cell. We can imagine information seekers selecting or designing templates for specific databases and types of information problems and placing them over incoming streams of information or on databases accessed purposefully. For example, a medical specialist could design a template that only allows new MEDLINE records added since the previous day to pass through. Thus, templates can be user profiles that users can directly manipulate according to their personal needs. For well-defined databases such as MEDLINE, templates are easily designed according to the database fields and are simply stored queries for automatic execution. Each template "cell" could contain verbose sets of items acceptable for that particular field. For example, dozens of medical subject headings and free-text terms could be included in a subject cell and when that template was applied to the database alone or in combination with other templates, large numbers of records would be blocked out. The resulting set of records may be large, but could be simply scanned or further restricted with other templates. Additionally, computational tools could be applied to specific cells or collections of cells, for example, links to previously located documents, aggregations with related documents, and statistical or lexical analyses. The template mechanism developed with such capabilities represents what Kay (1984) described as the next generation of spreadsheets.

Another implementation of the template metaphor is exemplified by the Perseus system. Juxtaposed pull-down menus contain scrolling lists of descriptors for various object classes. These object classes were hand-crafted to the classical Greek world and the resources in Perseus. For example, to locate a vase image of domestic animals, a user might pull down the object menu and select vases, then pull down the classes menu and pick animals. Menu selections restrict what parts of the database shows through in the results window. Other menus can be used in combination to further restrict what remains, for example, another object menu could be used to restrict objects to the Munich and Tampa museums.

A different implementation might allow users to directly manipulate, drag, and place templates over the universe of information and immediately see what remains. The templates manipulated in this manner would need identifiers and control panels for setting or defining constraints. Using menus for categorical attributes and sliders for interval attributes seems a reasonable way to aid users in setting template constraints. Once one or more templates were in place, settings could be adjusted and the residual information immediately perceived. Thus, placing and setting templates serve as an implementation option for what Shneiderman (1992) calls dynamic queries. The part of the database that is displayed depends on the parameter settings. Changing a setting with sliders or menu selections instantly redisplays the data elements that “show through” the newly configured template.

It is unlikely that any of these classes of mechanisms are by themselves robust enough to support all types of browsing. Probes, filters, and templates are similar in that they require the user to specify some hypotheses or constraints to begin, whereas zooms and pans afford immediate direct manipulation. Just as information seekers apply different browsing strategies according to the task, setting, information system, and personal experience, they will want to apply different mechanisms in different situations. Balancing rich sets of alternatives against learnability costs will remain a design challenge. The pros and cons of each mechanism are summarized below.

Easily manageable mechanisms are desired that allow users to access representations, control display, and manipulate information. Probes give good access capabilities; they allow users to specify what is wanted. Probes give generally poor display control; users either have to specify the display format as part of the probe or accept default displays. Manipulations of information such as aggregations or analyses are possible through highly qualified queries or systematic queries that produce sets that are combined. Filtering manipulations are possible by adding field limitations to the query. Mechanisms for probing are the most common tools that information seekers have in today’s systems.

Zooms and pans provide excellent display capabilities; they allow users to display information at levels of granularity and in juxtapositions that meet their needs. Zooms and pans are generally poor for access in large systems since they rely on perception and the user viewing, if only briefly, information in order to make acceptance/rejection judgments. Manipulations such as links may be effectively controlled with zooms and pans, but filtering, aggregation, and analysis are not supported with these mechanisms.

Templates provide excellent filtering capabilities and if template cells are properly instrumented, new ranges of linking, aggregation, and analysis are possible. When used alone, access is coarse, allowing all information in the database to pass through except that explicitly filtered out. Thus, templates may be unacceptable in very large databases. Filters and templates give good display control if users can instantly see the results of parameter changes, as in the case of dynamic queries.

From the information seeker's point of view, browsing in electronic environments is a problem of determining what representations to use for objects in the information universe and then examining these representations to identify relevant information or to change their choice of representation. From the system's point of view, supporting browsing is a problem of providing the user with alternative representations and mechanisms for controlling them. Distinctions have been made among representations according to granularity (e.g., whole/part), form and context (e.g., different formats for the same data), and value added (e.g., addition of new representations, links to others). Generic metaphors that may be useful for organizing classes of control mechanisms have also been discussed. What is needed is a unifying metaphor for representations and mechanisms that admit both the user and system points of view.

A Geometric Metaphor

Our present conceptions of browsing are constrained by manual and paper-based systems and it will take some time before browsing mechanisms are modeled on ubiquitously known electronic system features. One way to press the search for new thinking is to start with a generic and robust mathematical model and ask how a specific problem maps onto the objects and operations of the model. Such an approach focuses search for new discoveries and has proven useful in various scientific fields (e.g., chemical elements, new planets, etc.). A first step in this process is to choose a model and apply it as a metaphor for information seeking.

One mathematical model that has been proposed is to consider algebraic mappings from one form of information to another. Heilprin (1985) has proposed that compression of information from one form to another is a fundamental phenomenon in information science. Rather than well-defined isomorphic or homomorphic mappings, he has argued that a flexible similarity relation is necessary and calls such mappings paramorphisms. Although this model is powerful for capturing compression and decompression, a geometric model is here proposed that may be more familiar and concrete. Since one reason to present such a model is to provoke thought, familiarity should admit thought by larger segments of the community.

Geometry provides a framework for defining specific types of objects and operations on those objects. In geometry, operations vary according to how they preserve attributes such as angle, distance, and shape. Isometries such as reflections, rotations and translations preserve all these attributes but change the context (e.g., orientation in a plane) of objects. Projections preserve shape attributes such as the number of edges, but allow attributes such as angles and distance to vary. A special type of projection, dilations, preserve angles but not distance and yield similar figures. These represent the most obvious and numerous information compressions. Topological transformations preserve neighborhoods but not shape, distance or angles. Mathematical systems are appealing because they offer a framework to test new ideas and to determine where components of the framework are missing. Using such a model, we may be able to define information objects and operations that behave predictably, and determine what types of objects and operations best deserve attention for development.

Consider transformational operations defined on information objects. The most obvious operation, similar to a geometric dilation, is to represent a document at different levels of granularity. For example, moving from a bibliographic record to a table of contents, then to a list of section headings, then to a list of first and last sentences for all paragraphs, and finally to the full text, illustrates applying four negative dilations to the document. Of course, what is preserved are not angles and shapes, but concept classes and the gist of the document. Document dilations aim to preserve primary meaning or "aboutness" for a document. The greater the dilation, the less defined are the intra-structural relationships and details of the document, but we gain overview perspective for making comparisons across documents (e.g., screen display real estate) and minimize time to scan the entire representation. Abstracts, thumbnail images, and most of the examples of alternative representations and added values listed above are types of document dilations. Radiation readings from different altitudes or area resolution of photographs taken by satellites define scales that also lend themselves to dilations.

Isometries for document operation are less obvious. What is required is to keep the same representation but change its orientation? Displaying the same document in different intermediate sets of retrieved documents or from different databases may aide users in viewing the many contexts of a document. Translations of a text by different translators or different editions of a work also illustrate such transforms. Changing images from PICT to TIFF formats or a word processing file from one platform to another are also translations. Isometric operations may have good potential for images since backgrounds, colors, or brightness may be varied for good effect. Three-dimensional renderings of organs based on CAT scans are examples of rotational transforms. Different maps of the same region for different criteria such as geography, political unit, topography, vegetation type, or population density also illustrate possible isometric transforms.

A large void exists with respect to topological transformations. What are they? These seem to be information alchemy. By posing the question, we may stimulate novel re-representations and mechanisms for manipulating them. Displaying tables of data as graphs is one example of such a re-representation. Simple reformatting is an isometry, but reformatting that provides new access may be more significant. For example, reformatting text for learners has proven effective for readers with visual perceptions disabilities (Stuben & Vockell, 1993). In addition to font size variations and adding more space between words, underlining the first words of each line assisted readers with such disabilities to track text more easily during reading. Another possible example of topological transformations may be morphological variations for words. Synonymic transformations may also be considered as preserving meaning but not word form. Perhaps if this book were in electronic form, the reader could apply a customized thesaurus to systematically augment (or replace?) words.

Just as geometric objects can be combined under various composition functions, information objects can also be combined. For example, documents can be merged, integrated, or linked. Although literal merging may not be promising for text, graphic overlays have been demonstrated by using a transparency film metaphor to aide information seeking and analysis (Belge, Lokuge & Rivers, 1993). For textual databases, composition functions may apply morphological comparisons or simple matching of concordances.

BROWSING AND HUMAN-COMPUTER INTERACTION

Because browsing is both highly interactive and parallel in nature, the interactions among the various components of the information-seeking framework are more fluid and concurrent than during more formal analytical search. The nature of browsing demands interfaces that support easy and flexible control, high-quality display, and rapid response time. The close coupling of user, domain, setting, task and system is reflected in the information-seeking process. The phases of the information-seeking process are more integrated during browsing, especially when browsing is driven by learning, exploring, and accretion of knowledge goals.

Problem acceptance and definition are especially well-integrated throughout browsing activity. Casual browsing may be initiated by recognition and acceptance of the opportunity for browsing rather than some information problem or need, although as Kwasnik (in press) points out, goals emerge as browsing progresses. While browsing, the information problem is continually redefined and in many cases it is the user's desire for the redefinition process that initiates and drives browsing. The human-computer interface most affects these subprocesses by virtue of the quality of display and system response speed. These are essential so that the flow of interaction is not interrupted as browsers continue to get a better view of the information problem.

Search system selection during browsing may be a matter of availability or situation (e.g., whatever environmental stimuli are at hand) but typically involves information seekers in making inferences about probabilities of success for specific systems or accepting default systems based on experience. Since the environment plays a significant role in browsing strategies, information seekers must develop strategies for monitoring the degree to which they are "pulled" by the environment rather than controlling it. Hypertext linking can lead to distraction, and trivial details in “edutainment” environments can easily distract those who have a propensity for casual browsing and observational tactics. Search system selection and query formulation blur in networked environments, since they both entail selection of entry points for browsing. The development of standard query languages and generic network interfaces are a first step in disassociating the interface from specific databases, but it remains dubious whether any single, generic "window" on the universe of information will be sufficient to support the views required for all purposes and users.

As environments become larger and more complex, determination of entry points becomes more crucial to finding promising conceptual neighborhoods for browsing.

This is especially true in online and networked environments where many databases may be available. One approach is to create intelligent guides to database selection. Wang (1990) designed an algorithm for a database guide to bibliographic databases in the field of business and management. His algorithm depended on characterization of the databases according to factors such as scope, coverage, and indexing. Another approach is to create mechanisms for browsing the various neighborhoods. Zoom and pan mechanisms are promising for this purpose. This approach requires high-level views of systems and reinforces the need for multiple levels of representation within and across search systems.

Query formulation, a subprocess that requires significant cognitive resources in analytical search, is manifested during browsing as identifying a neighborhood. This is a particularly crucial activity when browsing large search systems. In electronic environments, this means choosing a representation to examine. In libraries, the catalog is typically used to identify relevant items and those items are located as an entry point for browsing. This works where books are shelved according to a classification system and stacks are open. When using an analytical strategy, a query is carefully formulated and well-delimited from other steps in the process. Much of the interaction with the search system and significant amounts of time revolve around this step. When browsing, the information-seeking process is more fluid. The query formulation step identifies an entry point for extensive execution and examination that is guided by problem definition and reflection/iteration. The interface is critical here with respect to presenting representations and control mechanisms that allow easy selection, rapid feedback, and easy alternatives based on feedback. Whereas an interface designed to support analytical search may offer suggestions and help, an interface designed to maximize browsing will offer direct manipulation of the information space and tend to leave interpretation and decision making to the user.

Whereas execution of a query and examination of results are discrete activities in analytical search, execution and examination are inextricably linked during browsing. In fact, the concurrent and mutually-reinforcing physiological and the cognitive activities are an important characteristic of browsing. Actions such as scanning, jumping, and navigating are much more common during the execution and examination subprocesses and likewise distinguish analytical and browsing searches. Since display of information is critical to examination, interfaces should provide organized representations and window arrangements, high resolutions, rapid response, and flexible tools for changing representations and marking traversal points. Because execution and examination require extensive selection actions, natural language styles may not be as useful as point and select styles.

Information extraction aids problem definition and reflection, iteration, and termination in similar fashion whether information seekers are searching formally or informally, but in both cases, search systems can provide better tools for assisting users with integrating extractions into their internal and external knowledge bases. Interface designers should aim to develop new tools beyond cut and paste and bookmarks to improve this subprocess. Since extraction of gist and ideas is usually the real aim of browsing, contextual information may be useful. For example, showing where in a list or array a record appears or displaying the data dictionary as well as the record may be especially helpful when customized views or indexes are used.

Reflection is an ongoing metaprocess necessary for any type of search to continue and eventually terminate. The speed and quality of feedback from the system are essential to the monitoring process. Too much delay or too little interpretable feedback cause frustration and early termination. Good display characteristics and the ability to easily change representation level will lead to more perseverance. The information-seeking framework and results from human factors research provide some guidance for system design to support browsing by informing the types of representations that may assist browsing and how mechanisms for manipulating those representations should be implemented.

Research on eye movements during reading, picture scanning, and visual search illustrate the importance of task and user expectations in predicting processing efficiency. Evidence that saccade latencies are significantly shorter when people know where the next target will appear (Rayner, 1978) encourage linear or systematic screen layouts and task-action grammars. Basic physiological-psychological research such at this and results from studies of motor control latencies (e.g., Fitts' law that predicts time to completion based on distance to target) provide the basis for interface design metrics such as layout appropriateness that minimize mouse movements and assure predictable and consistent placement of objects such as menu items or icons (Sears, 1992). Because browsing is dependent on physical-psychological interaction, such considerations are even more critical for interfaces designed to support browsing. Representations must not only be rapidly displayed and easily interpretable, but they must be predictably and consistently placed and organized to promote rapid eye and hand coordination. Likewise, selection and control mechanisms must be optimally placed and sequenced. These obvious transfers of results to browsing must be tested and extended by human factors researchers. For example, since left-to-right eye movements are expected by readers of English (and consequently lead to faster left-right saccades than right-left saccades during reading), do similar expectations hold for menus and icons? Likewise, since parafoveal vision is thought to be important during reading to prepare the eye for the next saccade and subsequent fixation, what similar value do adjacent screen objects offer during browsing?

Models for HCI such as the GOMS model of Card, Moran and Newell, have proven useful for predicting behavior and informing design for text editors. This model does not include metacognitive aspects of cognition and ignores learning effects, but Peck and John (1991) have extended the model to a well-structured information seeking domain. Their intensive study of a user of a help system suggested that GOMS may be extended to this particular type of browsing. Carmel, Crawford, & Chen (1992) have also applied the GOMS framework to explain the cognitive processing of browsers. The information seeking framework presented here is more global (considers setting and domain as well as user, task, and system) and less precise than a GOMS model of information seeking, but does reflect similar primitives. The GOMS goal relates to the information problem manifested as an information-seeking task; operators relate to the system tools such as probes, zooms, and filters for representing information; methods are the strategies and tactics known to users and possible to apply in a system; and selection rules are based on the information seeker's personal information infrastructure and how the system invites and supports various methods.

DESIGN PERSPECTIVE FOR SUPPORTING INFORMATION-SEEKING STRATEGIES

Ultimately, the many techniques that have been proposed and implemented in prototype systems must be tested, compared, and integrated into full-scale production-level systems that are used on a daily basis. It is unlikely that any information environment will be complete without supporting both analytical and browsing strategies. General-purpose systems must provide a variety of access methods and tools to accommodate disparate communities of users, and successful systems will allow users to customize different analytical and browsing strategies. At present, there are two approaches to bringing about this integration.

One approach is to create integrated systems that provide a broad range of retrieval functions and mechanisms for controlling those functions. One such system that uses expert system techniques to assist the user in applying powerful statistical retrieval techniques and browsable displays is Thompson and Croft's I^3R System (1989). They report good performance with the system but it requires significant computational resources. It provides a glimpse of the top-down approach to integration and depends on central intelligence built into the system to aid the local intelligence of users. Another example is the prototype multimedia browsing system of Gecsei and Martin (1989) that uses a fisheye display and weighted vectors to support browsing in a multimedia database.

A different approach to integration provides a variety of specific tools and depends on users using their local intelligence to apply the proper tools for complex information problems. The desktop metaphor illustrates this approach to integration in that rather than using large, integrated packages that handle many classes of problem, a variety of independent packages (e.g., word processor, painting program, file manager, communications program, etc.) are able to share data objects on the desktop. A similar bottom-up approach for information retrieval would provide a collection of analytical and browsing tools and allow users to apply them according to their particular needs.

In either case, users must understand the distinctions between browsing across documents and within documents, what different levels of representation are available, and how to manipulate and control these different representations.

Clearly, humans do not employ strictly analytical or strictly browsing strategies for exploring complex information, but rather various patterns, strategies, tactics, and moves associated with both. Search strategies are defined on a continuum with analytical and browsing extremes. The distinctions between any group of search strategies is in large part indicated by how parallel and tightly integrated are the information-seeking subprocesses. The most carefully planned analytical search illustrates discrete and sequential steps through the subprocesses and the most casual, observational browse illustrates examination of the environment which stimulates acceptance, definition, and reflection in parallel. Design practice has been driven first by simple performance parameters such as returning sets of objects in acceptable time and more recently by returning sets of objects that are useful for the task at hand. Task-oriented design has improved the usefulness of information retrieval systems for professional intermediaries but user-centered designs have begun to emerge. Designs that support browsing are well-advised to build on human capabilities and propensities first.