

Cat-a-Cone: An Interactive Interface for Specifying Searches and Viewing Retrieval Results using a Large Category Hierarchy

Marti A. Hearst
Xerox Palo Alto Research Center
3333 Coyote Hill Rd
Palo Alto, CA 94304
hearst@parc.xerox.com

Chandu Karadi
School of Medicine, M121
Stanford University
Stanford, CA 94305
karadi@leland.stanford.edu

Abstract

This paper introduces a novel user interface that integrates search and browsing of very large category hierarchies with their associated text collections. A key component is the separate but simultaneous display of the representations of the categories and the retrieved documents. Another key component is the display of *multiple* selected categories simultaneously, complete with their hierarchical context. The prototype implementation uses animation and a three-dimensional graphical workspace to accommodate the category hierarchy and to store intermediate search results. Query specification in this 3D environment is accomplished via a novel method for painting Boolean queries over a combination of category labels and free text. Examples are shown on a collection of medical text.

1 INTRODUCTION

There exist today many large online text collections to which category labels have been assigned. MEDLINE, a huge collection of biomedical articles, has associated with it Medical Subject Headings (MeSH) consisting of approximately 16,000 categories (Lowe & Barnett 1994). The Association for Computing Machinery (ACM) has developed a hierarchy of approximately 1200 category (keyword) labels¹. Yahoo!, one of the most popular search sites on the World Wide Web, organizes web pages into a hierarchy consisting of thousands of category labels². And traditional online bibliographic systems have for decades assigned subject headings to books and other documents (Svenonius 1986).

¹<http://www.acm.org/class/1991/cr91.html>

²<http://www.yahoo.com>

The meanings of the category labels differ somewhat among collections, but usually they are intended to help organize the documents and to aid in query specification. Unfortunately, as reported in a recent paper from the library sciences community, users of online bibliographic catalogs rarely use the available subject headings (Drabenstott & Weller 1996). These and other authors put much of the blame on poor (command line-based) user interfaces which provide little aid for selecting subject labels and force users to scroll through long alphabetic lists.

Although many researchers have investigated techniques for automatically augmenting word-based queries with category labels, there has been surprising little research on advanced user interfaces for browsing in and selecting from large category hierarchies for the purposes of information access. There has been still less on how to integrate category hierarchies with retrieval results, and work is especially lacking in the support of search when *multiple* categories have been assigned to each document.

This paper describes an interactive user interface that integrates search and browsing of very large category hierarchies with their associated text collections. The prototype system, called the Cat-a-Cone³, uses existing 3D+animation interface components, applied in a novel way, to support browsing and search of text collections and their category hierarchies.

A key component of the interface is the separation of the graphical representation of the category hierarchy from the graphical representation of the documents. This separation allows for a fluid, flexible interaction between browsing and search, and between categories and documents. It also provides a mechanism by which a *set* of categories associated with a document can be viewed along with their hierarchical context.

Another key component of the design is assignment of first-class status to the representation of text content.

³The Cat-a-Cone integrates *category* hierarchies into *ConeTrees*. The name is almost a homophone of *catacomb*, a word whose secondary definition is *a complex set of interrelated things*.

The retrieved documents are stored in a 3D+animation book representation (Card *et al.* 1996) that allows for compact display of moderate numbers of documents. Associated with each retrieved document is a page of links to the category hierarchy and a page of text showing the document contents. The user can “ruffle” the pages of the book of retrieval results and see corresponding changes in the category hierarchy, which is also represented in 3D+animation. All and only those parts of the category space that reflect the semantics of the retrieved document are shown with that document.

In summary, this interface is designed to exhibit the following features (see Figure 1):

- Make large hierarchical category sets easier to view and understand:
 - Allow for viewing of entire hierarchy in one window (occluded portions unveiled through animation)
 - Aid in the understanding of the meanings of unfamiliar terms
 - Help disambiguate terms that occur in several places
- Couple viewing of categories with search and display of retrieval results:
 - Separate category labels from documents, but link both together to allow for a flexible two-way interaction
 - Assign first-class status to the display of document content, rather than reducing a document to a small glyph or title alone
- Introduce a new, simple method for specifying complex Boolean queries consisting of free text plus category subtrees
- Be compatible with recent advances in ranking algorithms and with some other GUIs for information access.

Our interaction model is similar to that described in (Agosti *et al.* 1992). These authors define a two-level architecture for linking documents and their “auxiliary data”. However, the implementation and that used in a followup study (Belkin *et al.* 1993) use a text-based interface which does not provide most of the affordances listed above. (Ingwersen & Wormell 1986) also describes a text-based interface that allows the user to alternate between free text and thesaurus terms.

Before describing the interface in more detail, Section 2 discusses our view of the role of categories in information access and related work. Then Section 3 describes the Cat-a-Cone and examples of its use, Section 4 discusses other graphical information access interfaces, and Section 5 summarizes the paper.

2 CATEGORIES AND INFORMATION ACCESS

In this section we discuss what is meant by category labels in this work, in particular contrasting categories with other kinds of meta-information and with thesaurus terms. We also motivate our design decisions by first describing results of recent experiments in improving document ranking by automatically combining category labels with free text, and then discussing the special considerations that accompany the search and display of documents to which multiple categories have been assigned.

2.1 Categories vs. External Meta-Information

Most documents have some kinds of meta-information associated with them – that is, information that characterizes the external properties of the document, that help identify it and the circumstances surrounding its creation and use. These attributes include author(s), date of publication, length of document, publisher, and document genre. Several research groups are exploring issues associated with user interfaces for exploiting this kind of information, usually employing variations on a table format, e.g., (Fox *et al.* 1993; Baldonado & Winograd 1997).

By contrast, this paper focuses on category sets that have been assigned in an attempt to characterize the content or meanings found within the text of documents. Content-oriented category sets are difficult to depict graphically because they are abstract and there are potentially more meaningful combinations of content categories than external meta-information.

2.2 Thesaurus Term Associations

Many researchers have investigated automatic thesaurus creation, most often using word co-occurrences, e.g., (Crouch 1990; Ruge 1991; Evans *et al.* 1991; Grefenstette 1994). Thesaurus terms are related to category labels in information access systems, in that both kinds of information are used to improve recall. However, there are several important differences between the two.

Category labels are used to classify documents’ contents according to general subject areas and other semantic attributes. Categories are instantiated by the set of documents they are assigned to, and are represented by their labels. These labels are used as a kind of meta-information, and typically category labels are matched against other labels when used in search.

By contrast, thesaurus terms are usually used as alternative ways to express a concept. Thesaurus terms can compensate for the occasions when the actual words used by the user in the query do not match the way the concepts are expressed in the document. When a user adds terms to the query from a thesaurus (or

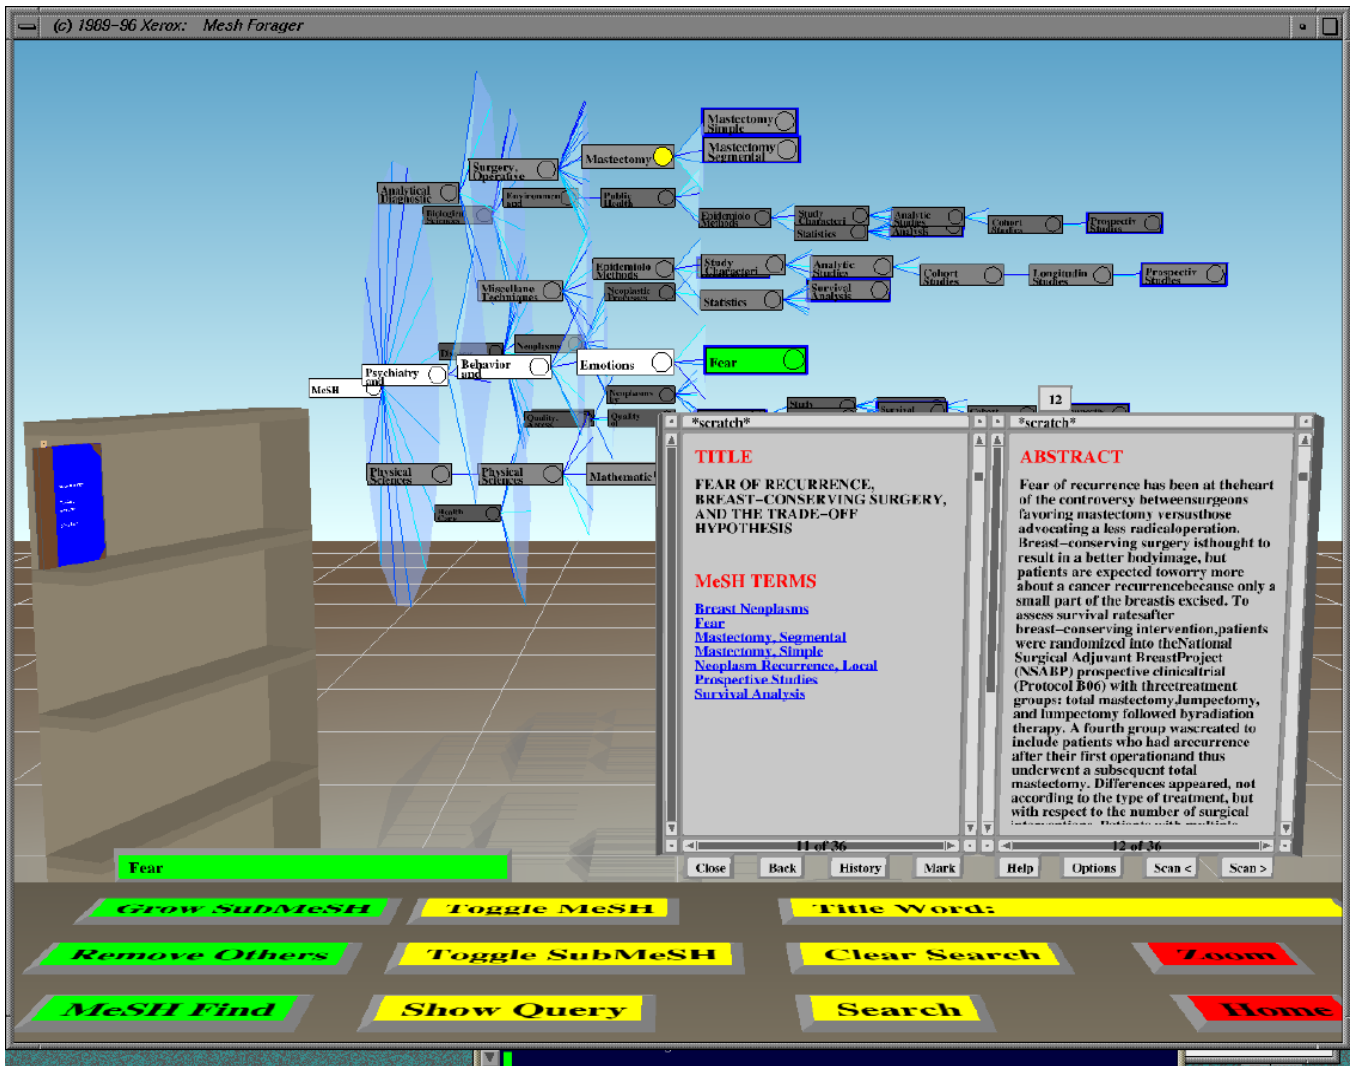


Figure 1: The Cat-a-Cone interface. Shown are the results of a search on category labels *Mastectomy* and *Radiation Therapy* in conjunction with the text word “lumpectomy” on a breast cancer subset of the MEDLINE collection. A ConeTree displays category labels and a WebBook shows retrieval results. The lefthand page shows the title and the category labels associated with the document. The righthand page shows the abstract associated with the document. Books that are the results of previous searches are stored in the workspace on the bookshelf, thus acting as a memory aid.

All and only those categories (and their ancestors) present in the current document are displayed in the category hierarchy. When the user clicks on a category label on the lefthand page, the corresponding category label in the ConeTree is rotated to the foreground and the labels of its ancestors are highlighted. When the user “turns” the pages of the book, the subtrees of the category hierarchy rotate, expand, and contract appropriately (categories that were present on the previous page and are present on the new page remain constant, categories that were on the previous page but are not on the new page are pruned away, and new categories that were not on the previous page but are on the new page are expanded out).

from relevance feedback for that matter), this usually means adding the actual words from the thesaurus into the query in hopes that they will appropriately match words in the documents of interest. Thesaurus terms are typically represented by the words themselves,

This distinction does not hold in all cases. Techniques such as Latent Semantic Indexing (Deerwester *et al.* 1990; Schütze 1993) map the meaning of one set of words to another by means of a vector representation. Thus it attempts to find general themes which could be thought of as categories although the representation is nonstandard. However, thinking of the terms as thesaurus terms rather than category labels tends to influence how they are used in the user interface.

2.3 Combining Category Labels and Free Text Search

Library catalog systems have long provided categorization information in the form of subject headings. Researchers have reported that these kinds of headings often mismatch user expectations (Svenonius 1986; Lancaster 1986). However, there is also evidence that when such subject heading information is combined with free text search, results are improved over using either categories or free text alone, although usually these improvements are small (Markey *et al.* 1982; Henzler 1978; Lancaster 1986). Most of this work was done in the bibliographic context and did not employ modern user interface technology.

Studies on the biomedical category set, MeSH, usually do not achieve strong improvements on MEDLINE searches with automatic addition of category labels over non-MeSH searches (Srinivasan 1996a). Only very small improvements were found in two studies (Hersh *et al.* 1994; Aronson *et al.* 1994) on a larger collection (using the Metathesaurus, a much larger hierarchy than MeSH (Schuyler *et al.* 1993)). Larger improvements were found in another study (Yang & Chute 1994), but under the assumption that a new query will match against a query that had already been seen and analyzed extensively with training examples.

However, recent careful studies (Srinivasan 1996b; Srinivasan 1996c) have shown that if done appropriately, adding in MeSH terms can lead to significant improvements in precision and recall simultaneously, even over an initial high baseline. Importantly, these improvements were not found by using the standard technique of automatically mapping natural language queries into semantically equivalent MeSH terms, as is done with most attempts to automatically improve ranking with category labels.

Rather, improvements were found by retrieving documents based on the free-text version of the query, taking the top-ranked documents for relevance feedback, and adding to the query the MeSH category labels that appeared in the top-ranked documents. This suggests

that first finding some good example documents, and then examining the category labels assigned to them, and using these to revise the query is a more effective way to improve rankings via category labels. This is precisely the kind of interaction that our framework is designed to support.

2.4 The Case for Multiple Category Representation

Many graphical user interfaces for text collections place documents within one point in semantic space, usually based on some measure of inter-document similarity (see Section 4). We assert there is a problem with assuming that documents can be placed into a single category or a single point in semantic space. Although real-life objects can (arguably) be assigned one place in a taxonomy (a truck is a kind of vehicle), the content of text is usually not so simply classified.

Consider for example a biomedical journal article entitled *Fear of Recurrence, Breast-Conserving Surgery and the Trade-Off Hypothesis*. This article has been manually assigned the MeSH category labels:⁴

Breast Neoplasms
Fear
Mastectomy (Segmental)
Mastectomy (Simple)
Neoplasm Recurrence (Local)
Prospective Studies
Survival Analysis

The article discusses a statistical study of the effects of a patient's fear of recurrence of breast cancer after a partial mastectomy versus the improved self image the patient has from retaining part of the breast. Thus, at a high level the human categorizers have placed the article into the semantic space at the intersection of Surgery, Statistics, and Psychology, since all three areas help characterize the complex subject matter of the document. At a more detailed level, the document content rests along the axes of a contrast between particular surgical procedures, a particular kind of statistical study and analysis, and the measurement of a particular psychological attribute.

When documents are clustered or grouped according to overall similarity, the distinctions about which axes they are similar on are not visible to the user. Placing the example document within a cluster of others might be done because any subset of the topic areas discussed were held in common. A cluster-based representation can be useful for some purposes, such as getting an overview of a collection's contents, but a

⁴Items in parentheses represent subheading modifiers for the main category labels. Our subset of MEDLINE documents have on average eight category labels assigned to them.

user interface that shows the inter-relations among documents according to their category labels (or in general, according to orthogonal semantic attributes) allows the user an alternative view of document similarity.

2.5 Existing Category-based Interfaces

Most interfaces that depict category hierarchies graphically do so by associating a document directly with the node of the category hierarchy to which it has been assigned. For example, clicking on a category link in Yahoo! brings up a list of documents that have been assigned that category label. Conceptually, the document is stored within the category label. When navigating the results of a search in Yahoo!, the user must look through a list of category labels and guess which one is most likely to contain references to the topic of interest. A wrong path requires backing up and trying again, and remembering which pages contain which information. If the desired information is deep in the hierarchy, or not available at all, this can be a time-consuming and frustrating process.

The MeSHBROWSE system (Korn & Shneiderman 1995) allows users to interactively browse a subset of semantically associated links in the MeSH hierarchy. From a given starting point, clicking on a category causes the associated categories to be displayed in a 2D tree representation. The interface has the space limitations inherent in a 2D hierarchy display and does not provide mechanisms for search over an underlying document collection.

Internet Grateful Med⁵ is a World Wide Web-based service that allows an integration of search with display and selection of MeSH category labels. The site is designed to support many simultaneous users and is constrained by the limitations of HTML. After the user types in the name of a potential category label, a long list of choices is shown in a page. To see more information about a given label, the user selects a link (e.g., Radiation Injuries). The current context is lost, and a new web page appears showing the ancestors of the term and its immediate descendants. If the user attempts to see the siblings of the parent term (Wounds and Injuries) then the context is changed again. Radiation Injuries appears as one of many siblings and the context of its children is lost. To go back to the previous list of choices, this page is obliterated.

One recent interface focuses on displaying faceted category sets (Allen 1995). In this work, hierarchical category labels corresponding to the Dewey Decimal system are shown indented in a scrollable window, in a focus-plus-context manner similar to that used in the Superbook table of contents (Egan *et al.* 1989). All documents that have been assigned a selected category

are listed in another scrollable window. When the user issues a search, all categories that have a title with a hit in them are displayed in a scrollable window, with a number showing how many hits fall into each category. Thus this interface does not support display of multiple labels per document.

Our earlier work (Hearst 1994) emphasized many of the same points as in the present work. However, the interface for choosing and displaying categories could not support hierarchies of categories and did not provide a good mechanism for viewing the retrieved documents. The Cat-a-Cone interface represents a dramatic improvement.

3 THE CAT-A-CONE

We assert that separating the documents from the category hierarchy can open the door to more powerful search and display strategies, especially for collections in which documents have multiple category assignments.

Our approach consists of three main components. The first is a better representation of the category space, the second is a separate, compact representation of retrieved documents, and the third is a model of interaction that makes use of these visualizations in novel ways. The instantiation of these ideas is in a prototype interface that makes use of particular visualization technologies, but does not exclude using alternative technologies that support the same or similar functionalities.

Our prototype implementation (written in Common LISP, running on Silicon Graphics IRIX machines), makes use of the ConeTree 3D+animation visualization from the Xerox PARC Information Visualizer (IV) (Robertson *et al.* 1993). (Processing speeds and monitor quality are increasing rapidly enough to make general availability realistic within the next few years.) The ConeTree allows for the display of a very large category hierarchy all in one window. Those categories that are farther away and less legible can be rotated to the foreground with a simple click on the leftmost (highest ancestor) category label. The interface allows users to choose from menus, or gesture-based “gardening” commands for growing and pruning subportions of the hierarchy, as well as keyboard accelerators and traditional buttons on the “desktop” portion of the workspace for specifying searches.

If a label is terminated with a triangle, this signals the existence of a subtree. The user can expand the category one level down by a right-drag gesture, or can select the node with the left mouse button and then click on the “Grow SubMeSH” button on the workspace to expand the entire subtree (this button’s label and functionality toggles depending on the state of the selected label).

The examples in the following subsections are drawn

⁵<http://igm.nlm.nih.gov:80/>

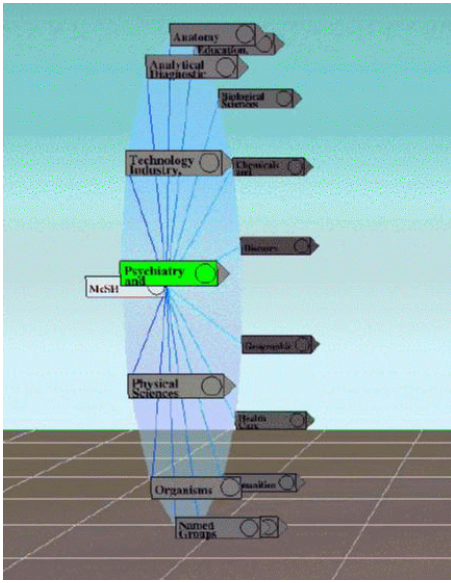


Figure 2: The top-level representation of the MeSH hierarchy as a search starting point.

from a subcollection of MEDLINE abstracts on breast cancer. This subcollection was further narrowed to include only those 403 documents that contained the word *lumpectomy* in their title or abstract (the equivalent MeSH term is *Mastectomy, Segmental*).

3.1 Starting a Search and Discovering Categories

Search interfaces must provide users with good ways to get started. (Shneiderman 1996) advocates a model of: first overview, then zoom and filter, then details, and repeat. In terms of the overview, an empty screen does not provide a good starting point. On the other hand, the contents of the entire hierarchy can be overwhelming, even if it can be made to fit in one window. The model of interaction used here allows for several different starting points and interaction sequences.

One starting point shows all of the top level categories initially, and allows the user to control the subsequent expansions (see Figure 2). The user can select one label to expand in detail, revealing all of its descendants rather than having to navigate laboriously through many pages to see all of the subtrees. If this is still too overwhelming, the user can select a few nodes and issue the command “prune others” to view a smaller space. Or the system can be programmed to automatically expand subtrees to a depth that has been determined by user studies to be most comprehensible.

An alternative starting point is to have the user type in a category label and see which parts of the hierarchy match or partially match that label. The user can then expand and explore other nearby regions.

As a final Cat-a-Cone starting point type, the user

can type in free text, causing the system to retrieve documents containing those words in their titles or abstracts. The user can then view the category labels associated with the retrieved documents. This kind of interaction should be useful in analogy with the findings of the experiments of (Srinivasan 1996b; Srinivasan 1996c) discussed above, in which good MeSH categories were associated with documents highly ranked against the free-text query.

For example, after a search on the category *Mastectomy* the system retrieves an article that has a link to *Survival Analysis*, a category which the user might not have known in advance. The user can then decide to delve more deeply into this topic by issuing a search on this category label, another category label, and one free-text word, yielding the book of Figure 1. This book in turn can aid the user in discovering that the MeSH hierarchy has a set of category labels pertaining to psychological issues, including *Emotions*.

When indexing new MEDLINE citations, indexers are instructed to use the most specific MeSH term available (Lowe & Barnett 1994). Often these low-level categories are meaningless to a patient or non-specialist. The context-preserving display of ancestor and sibling information provided by this representation can help the user see the general meaning of a term. For example, the category label Doxorubicin is the name of a chemical, and this is indicated by its ancestor labels.

3.2 Viewing Retrieved Documents

The Cat-a-Cone also makes use of a modification of WebBooks from the Web Forager project (Card *et al.* 1996), an extension of the Information Visualizer project. After a set of documents has been retrieved (in response to a query on free text and/or category labels), the documents are organized into a “book” of pages (see Figure 1). The lefthand page contains the document’s title and the list of category labels associated with that document. The righthand page shows the abstract and/or content of the documents (the pages are scrollable).

The cover of the book shows the query responsible for producing the retrieval results. When the book is closed and is the selected focus, the ConeTree can show a representation of all of the categories that have occurrences within the pages of the book. Note that since multiple categories have been assigned to each document, the book will in general contain many more category labels than were present in the original query. Only those parts of the hierarchy that contain categories that are in the book (and their ancestors) are shown automatically; the rest of the hierarchy is pruned away. The user is able to expand or contract any part of the hierarchy at any time.

When the user opens the book, the ConeTree is automatically modified to show only those parts of the

hierarchy whose categories appear in the document on the current page. The labels themselves are outlined in blue and the ancestors are shown without outline. This representation shows the space of concepts in which the document resides. When the user flips through the pages of the book, the representation of the tree adjusts accordingly. The use of animation helps retain the context of the category set. Often many category labels are shared among the book's retrieval results, so only a few offshoots of the hierarchy grow or are pruned as the user flips through the retrieved pages. The animation helps the user retain context, showing which parts of the category space differ from document to document.

Since the user can store the book away and reopen it at any time (making use of the workspace capabilities of IV) there is less reason to worry about getting "lost" or forgetting what happened in a previous session, because the retrieval results can be stored away and reused.

As currently described, the interaction between the category hierarchy and the retrieved documents is not entirely two-way: clicking on category labels does not influence the behavior of the retrieval results. This choice was made in part because there is a many-to-one mapping of category labels to book pages. However, the model could be altered as follows: clicking on a label causes the page of the book to turn to the next article containing that category label, if such an article exists.

In order to help further explain the retrieval results, other information access interface ideas can be incorporated into this representation. For example, we plan to place a TileBars (Hearst 1995) representation of the retrieval results into the first page of the book, which can be pulled out and viewed alongside the rest.

3.3 Query Specification

Research suggests that users often want to search on an "exploded" part of the MeSH hierarchy to improve recall (Lowe & Barnett 1994). In effect, they wish to specify a conjunction of disjunctions over category hierarchies, that is, require that at least one representative from each of a set of concepts be present in the retrieved documents. Our initial simple approach to query specification is to have the selection of the yellow circle on a category label indicate a disjunct of all of that category's descendants, inclusive. A conjunction is imposed over all selected subtrees. This approach is simple, but for queries such as (*hand OR foot*) AND *arthritis* it requires the user to select higher up in the anatomy subtree than desired.

For this reason, we have devised a novel scheme for specifying Boolean queries using the 3D environment. The user selects colors from a "palette" and "paints" subtrees with these colors, where each color representing members of a disjunction, and different colors indi-

cate different components of a conjunction. Additionally, free text search terms are typed into entry lines of corresponding colors. The users are instructed to think of the query in terms of colors rather than as a Boolean expression. For example, a user can specify a conjunction of disjunctions while thinking something like "I want documents that contain at least one green, at least one yellow, and at least one blue category or word." The retrieval results do not need to employ a strict Boolean filter; a quorum ranking strategy (Salton 1989) can be used instead. In this scheme, different colors receive different weight, and those documents represented by hits on more colors are ranked higher than those with fewer colors.

This query specification scheme thus stays within the object-centric paradigm and provides a simple way for users to link subsets of category labels and free text in complex Boolean expressions (with NOT represented by red).

4 OTHER GRAPHICAL APPROACHES

4.1 Graphical Concept Spaces

Several approaches map documents from their high dimensional representation in document space into a 2D representation in which each document is represented as a dot or other small glyph. The functions for transforming the data into the lower dimensional space differ, but the net effect is that documents are placed at one point in a scatter-plot-like representation of the space, and users are meant to detect themes or clusters in the arrangement of the glyphs. These systems include BEAD (Chalmers & Chitson 1992), ThemeScapes (Wise *et al.* 1995), and the Galaxy of News (Rennison 1994). Other systems, e.g. (Cutting *et al.* 1992; Maarek & Wecker 1994; Allen *et al.* 1993) display inter-document similarity hierarchically. The systems of (Fowler *et al.* 1991) and I³R (Thompson & Croft 1989) display retrieved documents in networks based on interdocument similarity.

Systems such as VIBE (Korfhage 1991) and the InfoCrystal (Spoerri 1993) ask the user specify the query in terms of k words (although category labels could be used instead) where k is a small number. They then display, for each subset of the k categories, the number of documents that contain that subset of words. These systems show the categories in a graphical concept space, and do not provide a mechanism for choosing which of a large number of words or categories to choose from, nor do they suggest new words or categories, nor associations among category labels, and they do not introduce methods for associating the text of the documents with new words or categories.

The Lyberworld system (Hemmje *et al.* 1994) makes use of a ConeTree but uses it quite differently than the

Cat-a-Cone. Lyberworld uses the ConeTree to display the navigation path generated by a sequence of search steps, placing the documents associated with the results of the search at the first level of the hierarchy, expanding on the words of a given document's node at the next level, and repeating.

Some researchers, e.g., (Pedersen 1993; Carpineto & Romano 1996), have employed a graphical depiction of a lattice for query formulation, where the query consists of a set of constraints on a hierarchy of categories (actually, semantic attributes in these systems). This is one solution to the problem of displaying documents in terms of multiple attributes; a document containing terms A, B, C, and D could be placed at a point in the lattice with these four categories as parents. However, if such a representation were to be applied to retrieval results instead of query formulation, the lattice layout would in most cases be too complex to allow for readability.

In the AIR/SCALIR interface (Rose & Belew 1991) a connectionist network determines in advance a set of terms that characterize documents from a collection of bibliographic records. The term nodes are connected to the document nodes via edge links, so the user can see which documents are associated with each important term. If there are a large number of links between associated terms and documents, or if the links are not neatly organized, the relationships will be difficult to discern.

Finally, Kohonen's feature map algorithm has been used to create maps that graphically characterize the overall content of a document collection or subcollection (Lin *et al.* 1991; Chen *et al.* 1997). The regions of the 2D map vary in size and shape corresponding to how frequently their corresponding themes occur in the collection. Regions are characterized by single words or phrases, and adjacency of regions is meant to reflect semantic relatedness of the themes within the collection. If a document is strongly associated with the region according to the training of the feature map, its title can be viewed via a pop-up window over that region; documents can be associated with more than one region.

4.2 A Comparison

A recent evaluation (Chen *et al.* 1997) compared the Kohonen feature map representation on a browsing task to that of Yahoo!. The results found that some users expressed a desire for a visible hierarchical organization, others wanted an ability to zoom in on a subarea to get more detail, and some users disliked having to look through the entire map to find a theme, desiring an alphabetical ordering instead. Many found the single-term labels to be misleading, in part because they were ambiguous (one region called "BILL" was thought to correspond to a person's name rather than counting

money). The subjects like the ease of being able to jump from one area to another without having to back up as is required in Yahoo! and liked the fact that the maps have varying levels of granularity. (The authors concluded that this interface is more appropriate for casual browsing than for search.)

These results all support the design decisions made in the Cat-a-Cone. Hierarchical representation of term meanings is supported, so users can choose which level of description is meaningful to them. Furthermore, they can view different levels of description simultaneously, so more familiar concepts can be viewed in more detail, and less familiar at a more general level. An alphabetical ordering of the categories coupled with a regular expression search mechanism allows for straightforward location of category labels. Retrieved documents are represented as first-class objects, so full text is visible, but in a compact form. Category labels are disambiguated by their ancestor/descendant/sibling representation. Users can jump easily from one category to another and can in addition query on multiple categories simultaneously (something that is not a natural feature of the maps). The Cat-a-Cone has several additional advantages as well, such as allowing a document to be placed at the intersection of several categories, and explicitly linking document contents with the category representation.

5 CONCLUSIONS

The Cat-a-Cone interface allows for fluid two-way interaction between selection of category labels for search and display of multiple category labels within retrieval results. The ConeTree provides easy selection of subparts of the category hierarchy, to help users understand unfamiliar terms and usages of ambiguous terms by seeing their contexts. The book-based representation of the retrieval results with the corresponding display of multiple subparts of the category hierarchy helps show which categories are a part of the results, what these categories mean, and which new categories might be useful to search on.

Although this is a simple combination of simple ideas, it seems to produce a powerful, intuitive and original way for users to use large category hierarchies to aid them in query specification and understanding of retrieval results. In future work we plan to evaluate the interface using as subjects breast cancer patients and clinicians.

This kind of interaction should also be useful for other tasks, for example, helping authors of articles for ACM publications *choose* which category labels to assign to their documents. The user can tell the system to find documents whose textual contents are similar to their new document's, and then examine the resulting

category hierarchy.

The ConeTree as currently implemented does not satisfy fully the needs of the design. For example, the system should reformat subtrees when space becomes available as a consequence of pruning away other subtrees. Additionally, a recent study shows that fisheye, variable zoom algorithms work better than full zoom algorithms for navigating networks (Schaffer *et al.* 1996), and an improved facility for focus-plus-context of this kind could be used to help direct user attention. Subtrees should be easily “pulled off” and stored for later use. Finally, the display should reflect the frequency of the categories for a given retrieval result with a visual analogue such as “Read Wear” (Hill *et al.* 1992) or using a Kohonen feature map (Lin *et al.* 1991) as an overview of the highest-level categories.

We view information retrieval as a complex task that requires many different tools. Support for category browsing and search is one important capability, but by no means solves the entire problem. This interface must be combined with others to create an effective information seeking system.

Acknowledgements

Karadi was supported in part by a Medical Scholars Grant from Stanford Medical School. We would like to thank Larry Fagan, Bill York, and Stu Card for their help with and encouragement of this work. Additional figures can be found at <http://www.parc.xerox.com/ia>.

References

- AGOSTI, M., G. GRADENIGO, & P.G. MARCHETTI. 1992. A hypertext environment for interacting with large textual databases. *Information Processing & Management* 28.371–387.
- ALLEN, ROBERT B. 1995. Two digital library interfaces that exploit hierarchical structure. In *Proceedings of DAGS95: Electronic Publishing and the Information Superhighway*, Boston, MA.
- , PASCAL OBRY, & MICHAEL LITTMAN. 1993. An interface for navigating clustered document sets returned by queries. In *Proceedings of ACM COOCS: Conference on Organizational Computing Systems*, Milpitas, CA.
- ARONSON, A., T. RINDGLESCH, & A. BROWNE. 1994. Exploiting a large thesaurus for information retrieval. In *Proceedings of RIAO '94; Intelligent Multimedia Information Retrieval Systems and Management*, 197–216.
- BALDONADO, MICHELLE Q. WANG, & TERRY WINOGRAD. 1997. Sensemaker: An information-exploration interface supporting the contextual evolution of a user's interests. In *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems*. To appear.
- BELKIN, N., P. G. MARCHETTI, & C. COOL. 1993. Braque – design of an interface to support user interaction in information retrieval. *Information Processing and Management* 29.325–344.
- CARD, STUART K., GEORGE G. ROBERTSON, & WILLIAM YORK. 1996. The webbook and the web forager: An information workspace for the world-wide web. In *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems*, Vancouver, Canada.
- CARPINETO, CLAUDIO, & GIOVANNI ROMANO. 1996. Information retrieval through hybrid navigation of lattice representations. *International Journal of Human-Computer Studies* 45.553–578.
- CHALMERS, MATTHEW, & PAUL CHITSON. 1992. Bead: Exploration in information visualization. In *Proceedings of the 15th Annual International ACM/SIGIR Conference*, 330–337, Copenhagen, Denmark.
- CHEN, HSINCHEN, ANDREA L. HOUSTON, ROBIN R. SEWELL, & BRUCE R. SCHATZ. 1997. Internet browsing and searching: User evaluations of category map and concept space techniques. *Journal of the American Society for Information Sciences (JASIS)*. To appear.
- CROUCH, C. J. 1990. An approach to the automatic construction of global thesauri. *Information Processing and Management* 26.629–640.
- CUTTING, DOUGLASS R., JAN O. PEDERSEN, DAVID KARGER, & JOHN W. TUKEY. 1992. Scatter/Gather: A cluster-based approach to browsing large document collections. In *Proceedings of the 15th Annual International ACM/SIGIR Conference*, 318–329, Copenhagen, Denmark.
- DEERWESTER, SCOTT, SUSAN T. DUMAIS, GEORGE W. FURNAS, THOMAS K. LANDAUER, & RICHARD HARSHMAN. 1990. Indexing by latent semantic analysis. *Journal of the American Society for Information Science* 41.391–407.
- DRABENSTOTT, KAREN M., & MARJORIE S. WELLER. 1996. The exact-display approach for online catalog subject searching. *Information Processing and Management* 32.719–745.
- EGAN, DENNIS E., JOEL R. REMDE, LOUIS M. GOMEZ, THOMAS K. LANDAUER, JENNIFER EBERHARDT, & CAROL C. LOCHBAUM. 1989. Formative design evaluation of SuperBook. *Transaction on Information Systems* 7.
- EVANS, DAVID A., KIMBERLY GINTHER-WEBSTER, MARY HART, ROBERT G. LEFFERTS, & IRA A. MONARCH. 1991. Automatic indexing using selective NLP and first-order thesauri. In *Proceedings of the RIAO*, volume 2, 624–643.
- FOWLER, RICHARD H., WENDY A. L. FOWLER, & BRADLEY A. WILSON. 1991. Integrating query, thesaurus, and documents through a common visual representation. In *Proceedings of the 14th Annual International ACM/SIGIR Conference*, 142–151, Chicago.
- FOX, EDWARD A., DEBORAH HIX, LUCY T. NOWELL, DENNIS J. BRUENI, WILLIAM C. WAKE, LENWOD S. HEATH, & DURGESH RAO. 1993. Users, user interfaces, and objects: Envision, a digital library. *Journal of the American Society for Information Science* 44.480–491.
- GREFENSTETTE, GREGORY. 1994. *Explorations in automatic thesaurus discovery*. Kluwer international series in engineering and computer science. Kluwer Academic Publishers.
- HEARST, MARTI A. 1994. Using categories to provide context for full-text retrieval results. In *Proceedings of RIAO '94; Intelligent Multimedia Information Retrieval Systems and Management*, 115–130.
- . 1995. Tilebars: Visualization of term distribution information in full text information access. In *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems*, Denver, CO.

- HEMMJE, MATTHIAS, CLEMENS KUNKEL, & ALEXANDER WILLETT. 1994. LyberWorld – a visualization user interface supporting fulltext retrieval. In *Proceedings of the 17th Annual International ACM/SIGIR Conference*, 249–259, Dublin, Ireland.
- HENZLER, ROLF G. 1978. Free or controlled vocabularies: Some statistical user-oriented evaluations of biomedical information systems. *International Classification* 5.21–26.
- HERSH, WILLIAM R., DAVID H. HICKMAN, BRIAN HAYNES, & K. ANN MCKIBBON. 1994. A performance and failure analysis of SAPHIRE with a MEDLINE test collection. *Journal of the American Medical Informatics Association* 1.51–60.
- HILL, WILLIAM C., JAMES D. HOLLAN, DAVE WROBLEWSKI, & TIM MCCANDLESS. 1992. Edit wear and read wear. In *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems*, 3–9.
- INGWERSEN, PETER, & IRENE WORMELL. 1986. Improved subject access, browsing, and scanning mechanisms in modern online IR. In *Proceedings of the 9th Annual International ACM/SIGIR Conference*, 68–76, Pisa, Italy.
- KORPHAGE, ROBERT R. 1991. To see or not to see – is that the query? In *Proceedings of the 14th Annual International ACM/SIGIR Conference*, 134–141, Chicago.
- KORN, FLIP, & BEN SHNEIDERMAN. 1995. Navigating terminology hierarchies to access a digital library of medical images. Technical Report HCIL-TR-94-03, University of Maryland.
- LANCASTER, F. 1986. *Vocabulary control for information retrieval, second edition*. Arlington, VA: Information Resources.
- LIN, XIA, DAGOBERT SOERGEL, & GARY MARCHIONINI. 1991. A self-organizing semantic map for information retrieval. In *Proceedings of the 14th Annual International ACM/SIGIR Conference*, 262–269, Chicago.
- LOWE, HENRY J., & G. OCTO BARNETT. 1994. Understanding and using the medical subject headings (MeSH) vocabulary to perform literature searches. *Journal of the American Medical Association (JAMA)* 271.1103–1108.
- MAAREK, Y. S., & A.J. WECKER. 1994. The librarian's assistant: Automatically assembling books into dynamic bookshelves. In *Proceedings of RIAO '94; Intelligent Multimedia Information Retrieval Systems and Management*.
- MARKEY, KAREN, PAULINE ATHERTON, & CLAUDIA NEWTON. 1982. An analysis of controlled vocabulary and free text search statements in online searches. *Online Review* 4.225–236.
- PEDERSEN, GERT SCHMELTZ. 1993. A browser for bibliographic information retrieval, based on an application of lattice theory. In *Proceedings of the 16th Annual International ACM/SIGIR Conference*, 270–279, Pittsburgh, PA.
- RENNISON, EARL. 1994. Galaxy of news: An approach to visualizing and understanding expansive news landscapes. In *Proceedings of UIST '94, ACM Symposium on User Interface Software and Technology*, 3–12, New York.
- ROBERTSON, GEORGE C., STUART K. CARD, & JOCK D. MACKINLAY. 1993. Information visualization using 3D interactive animation. *Communications of the ACM* 36.56–71.
- ROSE, DANIEL E., & RICHARD K. BELEW. 1991. Toward a direct-manipulation interface for conceptual information retrieval systems. In *Interfaces for information retrieval and online systems*, ed. by Martin Dillon, 39–54. New York, NY: Greenwood Press.
- RUGE, GERDA. 1991. Experiments on linguistically based term associations. In *Proceedings of the RIAO*, 528–545.
- SALTON, GERARD. 1989. *Automatic text processing: the transformation, analysis, and retrieval of information by computer*. Reading, MA: Addison-Wesley.
- SCHAPPER, DOUG, ZHENGPING ZUO, SAUL GREENBERG, LYN BARTRAM, JOHN DILL, SHELLI DUBS, & MARK ROSEMAN. 1996. Navigating hierarchically clustered networks through fisheye and full-zoom methods. *ACM Transactions on Computer-Human Interaction* 3.162–188.
- SCHÜTZE, HINRICH. 1993. Word space. In *Advances in neural information processing systems 5*, ed. by Stephen J. Hanson, Jack D. Cowan, & C. Lee Giles. San Mateo CA: Morgan Kaufmann.
- SCHUYLER, P. Y., W. T. HOLE, M. S. TUTTLE, & D. D. SHERERTZ. 1993. The UMLS metathesaurus: representing different views of biomedical concepts. *Bulletin of the Medical Library Association* 81.217–222.
- SHNEIDERMAN, BEN. 1996. The eyes have it: A task by data type taxonomy. In *Proceedings of Visual Languages '96*, Boulder, CO.
- SPOERRI, ANSELM. 1993. InfoCrystal: A visual tool for information retrieval & management. In *Proceedings of Information Knowledge and Management '93*, Washington, D.C.
- SRINIVASAN, PADMINI. 1996a. Optimal document-indexing vocabulary for medline. *Information Processing and Management* 32.503–514.
- . 1996b. Query expansion and medline. *Information Processing and Management* 32.431–443.
- . 1996c. Retrieval feedback in medline. *Journal of the American Medical Informatics Association (JAMA)* 3.157–167.
- SVENONIUS, ELAINE. 1986. Unanswered questions in the design of controlled vocabularies. *Journal of the American Society for Information Science* 37.331–340.
- THOMPSON, R. H., & B. W. CROFT. 1989. Support for browsing in an intelligent text retrieval system. *International Journal of Man [sic] -Machine Studies* 30.639–668.
- WISE, JAMES A., JAMES J. THOMAS, KELLY PENNOCK, DAVID LANTRIP, MARC POTTIER, & ANNE SCHUR. 1995. Visualizing the non-visual: Spatial analysis and interaction with information from text documents. In *Proceedings of the Information Visualization Symposium '95*, 51–58. IEEE Computer Society Press.
- YANG, YIMING, & CHRISTOPHER G. CHUTE. 1994. Expert network: Effective and efficient learning from human decisions in text categorization and retrieval. In *Proceedings of the 17th Annual International ACM/SIGIR Conference*, 13–22, Dublin, Ireland.