From Keyword Search to Exploration: How Result Visualization Aids Discovery on the Web

Bill Kules
School of Library and Information Science
The Catholic University of America,
Washington, DC 20064
kules@cua.edu

Max L. Wilson, m.c. schraefel
Intelligence, Agents, Multimedia (IAM) Group,
School of Electronics and Computer Science
University of Southampton
SO17 1BJ
{mlw05r, mc}@ecs.soton.ac.uk

Ben Shneiderman
Department of Computer Science and Human-Computer Interaction Laboratory,
Institute for Advanced Computer Studies
University of Maryland at College Park, College Park, MD 20742
ben@cs.umd.edu

Please send correspondence to Bill Kules (kules@cua.edu)
(202) 319-6278 (office)
(240) 599-7671 (fax)

Abstract

A key to the Web's success is the power of search. The elegant way in which search results are returned is usually remarkably effective. However, for exploratory search in which users need to learn, discover, and understand novel or complex topics, there is substantial room for improvement. Human computer interaction researchers and web browser designers have developed novel strategies to improve Web search by enabling users to conveniently visualize, manipulate, and organize their Web search results.

This monograph offers fresh ways to think about search-related cognitive processes and describes innovative design approaches to browsers and related tools. For instance, while key word search presents users with results for specific information (e.g., what is the capitol of Peru), other methods may let users see and explore the *contexts* of their requests for information (related or previous work, conflicting information), or the properties that associate groups of information assets (group legal decisions by lead attorney).

We also consider the both traditional and novel ways in which these strategies have been evaluated. From our review of cognitive processes, browser design, and

evaluations, we reflect on the future opportunities and new paradigms for exploring and interacting with Web search results.

Table of Contents

Table of Figures

1. Introduction

A key to the Web's success is the power of search. The elegant way in which search results are returned is usually remarkably effective. However, for exploratory search in which users need to learn, discover, and understand novel or complex topics, there is substantial room for improvement. Exploratory search methods are also valuable when users have poorly formed goals and when the data indexing methods do not match the users queries. Human computer interaction researchers and web browser designers have developed novel strategies to improve Web search by enabling users to conveniently visualize, manipulate, and organize their Web search results.

This monograph offers fresh ways to think about search-related cognitive processes and describes innovative design approaches to browsers and related tools. For instance, while keyword search presents users with results for specific information (e.g., what is the capital of Peru), other methods may let users see and explore the *contexts* of their requests for information (related or previous work, conflicting information), or the properties that associate groups of information assets (group legal decisions by lead attorney).

We also consider the ways in which these strategies have been evaluated, and suggest methods for improving future evaluations. From our review of cognitive processes, browser design, and evaluations, we reflect on the future opportunities and new paradigms for exploring and interacting with Web search results.

1.1 Scope of the Monograph

With the many faces of the Web, defining the scope of this monograph is challenging. We think of the Web as a distributed, semi-structured, heterogeneous collection and consider interfaces that support search over this collection. We emphasize browser and HTML-based interfaces for Web searching, but we also consider browser or HTML-based interfaces to structured collections to the extent that they embody ideas that are applicable to the Web. Finally, in the history of search there have been two major streams of research: improving the speed and accuracy of keyword searching (information retrieval), and investigating interface developments to enhance search (information seeking or interactive information retrieval). In this report, we survey search environments for the latter of these two streams of research and discuss developments such as: functions to refine a query, alternative representations, and the clustering or grouping of results.

1.2 Reviews & surveys

A number of surveys and reviews address similar areas. Hearst (Hearst, 1999b) surveyed seminal visual search interfaces. She focused on non-Web collections, which had the benefit of a consistent schema, but pointed out that the Web would be a rich area of research for search interfaces. Katifori et al. (Katifori et al., 2007) produced a similar survey that is specifically directed at visualising ontologies. Many of the interfaces covered in the monograph below are also mentioned by Katifori et al. McKiernan (McKiernan, 2003) highlighted working applications of metadata-based and automatically derived visualizations for electronic journals. For the web, Kobayashi and Takeda (Kobayashi and Takeda, 2000) and Yang (Yang, 2005) briefly reviewed visualization techniques as part of larger reviews of online information retrieval systems. JASIST published a Perspectives issue on Visual Information Retrieval Interfaces in memory of Robert R. Korfhage (Rorvig and Lunin Lois, 1999).

1.3 Key Terms

- Document – The target objects that are being sought. On the web, this includes a wide range of media, including: web pages, images, videos, audio, and dynamically generated pages.
- Collection – The collection of documents being searched. This can be a fixed set of an ever-escalating corpus like the web.
- Query – In the simplest case, the query is composed of query terms for a free-text search, which is a special case of the more general case in which queries can include parameters such as those typically found in Advanced Search forms.
- Result Set – The set of documents returned by a Query. Result sets are often ordered by ranking the results. The metric for ranking is varied, and part of the first stream of research mentioned above.
- Result Item – A representation of an individual document in the Result Set. Result items are the richest objects in this analysis. Below are terms regularly associated with a result item:
  - Document – The document that is being represented by the Result Item, or means to retrieve it (often by a unique identifier or URL).
  - Title – Usually the title of the web page, or associated title of the multi-media if known.
  - Summary/Preview – A sample of Document (text taken from a webpage or a still from a video), usually indicating the relationship of the Result Item to the Query.
  - Rank – Based on a chosen metric appropriate for the Collection.
  - Author – The author of the metadata if known.
  - Organization – The associated organization if known.
  - Classification category – If annotated, the associated categories are returned.
- Classification – A set of meta-data, in varying arrangements (such as a hierarchy or graph) and depths (single or multiple levels) that can be used to annotate Result Items, usually when the Collection is fixed and known.
- Category – Individual categories are the parts of a Classification scheme, which are assigned to the Result Item as part of the annotation.

1.4 Web Collection Characteristics

This monograph focuses on exploratory information seeking over the decentralized collection of documents and information access points commonly known as the World Wide Web. The challenging scale of the Web as a collection is one of several characteristics that are relevant here. The “indexable web”, that is, the portion of the web indexed by major search engines, has been estimated at 11.5 billion pages (Gulli and Signorini, 2005).

A more indicative and remarkable characteristic is the heterogeneity of the contents. The Web contains a variety of data and documents. Documents may be in plain text, HTML, XML, PDF, Rich Text Format (RTF), Microsoft Word (.DOC), spreadsheets, and a multitude of specialized or proprietary formats, including microformats (Allsop, 2007). Multiple forms of media, including still images, audio, and video, are widely available and indexed by general purpose as well as specialized search engines. These documents vary from highly structured databases, to semi-structured web pages, to unstructured text.

The metadata available in search results is often limited, in part because of the very heterogeneity of the content just described. There is no universally adopted structure or schema for web search result metadata. Search engines often take the "least common denominator" approach and provide just a title, URL and snippet.

Documents may be static or dynamically generated from an underlying data source such as a database, eCommerce site, RSS feed, inferred user preferences or needs, or advertising algorithm. Indeed, the term document does not fully capture the range of content available. Video and audio may be streamed, allowing the viewer to watch at normal speed, or high or low speeds, or jump to specific parts of the presentation. It may also allow the user to browse to arbitrary URLs on the Web via hyperlinks embedded in the multimedia. This integration of media and Web is also seen in virtual worlds such as Second Life, which enable a user's avatar to walk, run, or fly through 3-D simulated worlds, interacting with other avatars and content in the form of landscapes, buildings, kiosks, furniture, in-world movies, and embedded hyperlinks to out-of-world content such as web pages.

The ease of publishing on the web leads to a wide variety of non-curated content, with a corresponding variety in quality. It allows multiple points of view. Individual documents may be biased or may be outright propaganda. Authorship may be difficult to determine or even disguised. Sources may not be credited. The failure to respect intellectual property rights and plagiarism issues are widely discussed.

Topics vary across a multitude of disciplines and knowledge domains. They range from everyday interests such as news, sports, and weather, to support groups for people struggling with specific medical problems, to highly specialized niches.

2. Framework – Theory of Search

This section presents a three-level model of the information seeking process (Figure 1). The model provides a framework for discussing the work and personal tasks that motivate information seeking, information seeking strategies, stages and tasks, and specific information retrieval tasks, tactics and actions.

2.1 The model

Examining search results is a necessary step within a larger information seeking process, the objective of which is to satisfy a perceived information need or problem (Marchionini, 1995; Marchionini and Shneiderman, 1988; Shneiderman, 1997). In turn, the perceived information need is motivated and initiated by a higher-level work task (Byström and Hansen, 2002, 2005; Järvelin and Ingwersen, 2004) or personally motivated goal (Kari, 2006).

Work tasks are situated in the work organization and reflect organizational culture and social norms, as well as organizational resources and constraints. The work task is similar to Sutcliffe and Ennis' goal or information need (Sutcliffe and Ennis, 1998), or Marchionini's recognition and acceptance of an information problem, but the work task specifically situates these in an organizational context. In the context of the work task, a second level of context is defined, in which information-seeking tasks are identified. These tasks vary as the work task progresses. The third level of context is the information retrieval context, wherein searchers identify sources, issue queries, and examine results. Reflection is inherent in each activity, and each activity except query formulation can return to a previous or higher level activity.

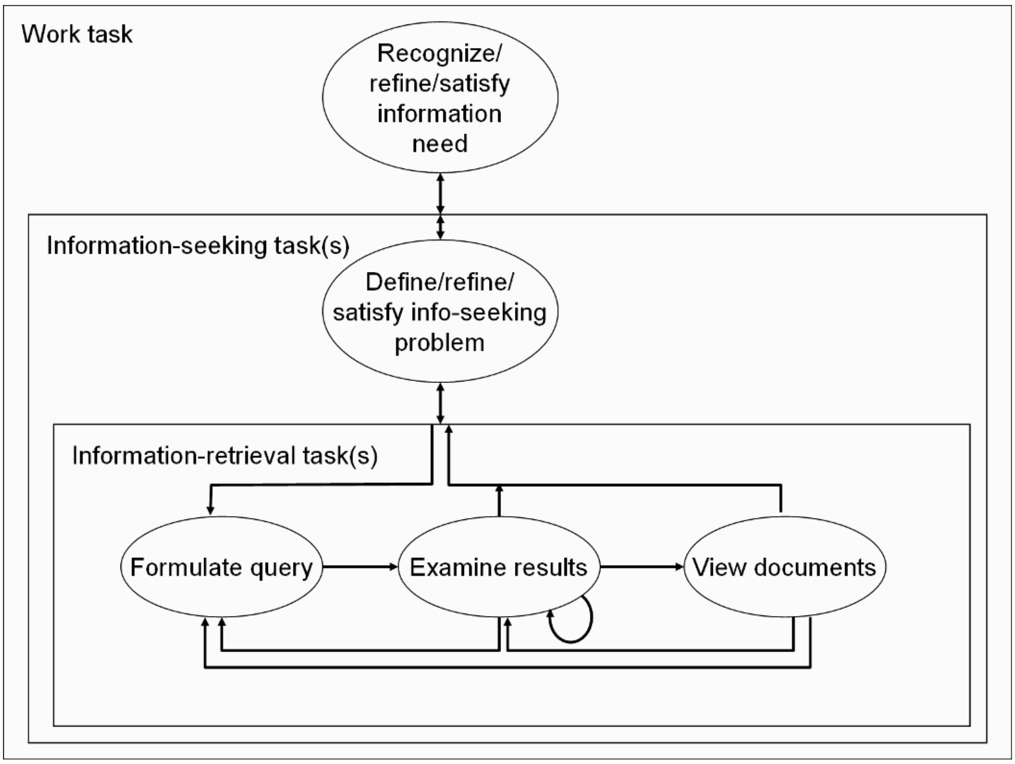

Figure 1. Process model of search in the context of work and information-seeking tasks.

Marchionini's electronic browsing model describes a seven-stage process (Marchionini, 1995). This model combines the Marchionini model with the three-level Byström & Hansen model. The model defines five activities: recognize an information need (to satisfy a work task), define an information-seeking problem (to satisfy the information need), formulate query, examine results, and view documents. It places activities in the context of the three levels of information-seeking and work tasks. It shows how search activities are sequenced within the iterative search process. Each higher-level activity can involve multiple subsidiary activities.

2.2 Work tasks

The process is initiated when a searcher recognizes an information need based on an organizational or personal need and decides to try to satisfy it (Byström and Hansen, 2002, 2005; Marchionini, 1995).

Evolving information needs form a core motivation for information seeking. Dervin and Nilan (Dervin and Nilan, 1986) consider user needs in the context of a sense-making theory of human behaviour. Gaps in knowledge are conceptualized as questions, which can motivate a person to seek information. Belkin (Belkin, 1980) developed the Anomalous States of Knowledge model to explain information seeking behaviour on open-ended questions. The model addresses iteration and refinement of the seeker's knowledge, specification of the problem, and an evolving ability to articulate requests. They often do not have a clear understanding of the need. In fact, they may only be partially conscious of the need, which Taylor (1962) calls a visceral information need.

Usually, users' information needs are initially ill-defined, and undergo a process of refinement. Kuhlthau's model of the stages of the information seeking process tracks cognitive and affective states in a constructive knowledge acquisition process such as writing a paper (Kuhlthau, 1991). At early stages, when the topic is not well-specified, information needs are more fluid. As the topic is refined, information needs become narrower and more fixed. Other researchers have found a similar evolution (Vakkari, 2000).

People do not always choose to seek information. They may prefer to avoid information that is discomforting or troublesome or which is judged not worth the effort (Mooers, 1960). They may also be biased in what information they seek, looking for information that supports preconceived ideas or opinions (McPherson, 1983).

2.3 Information seeking tasks

To satisfy the information need, information seekers will make decisions on what strategies to use and which tools and sources to consult as they undertake one or more information-seeking tasks. Although authors have used varying definitions, a commonly used definition for information seeking strategy is a high level plan for the whole search (Bates, 1979b, a; Marchionini, 1995). Strategies can guide source selection as well as the choice of information seeking tasks and information retrieval tasks and tactics (see below).

Information seeking tasks can be structured as a linear sequence or hierarchical decomposition of tasks. For example, the paper writing process could be modelled as a series of stages (Kuhlthau, 1991), or a medical search could be modelled using a hierarchical decomposition of goals (Bhavnani and Bates, 2002). Each of these tasks requires selecting a source, and then engaging in one or more information retrieval tasks which satisfy a portion of the overall need. From the perspective of an organization, Choo, Detlor and Turnbull (Choo et al., 2000) develop a behavioural model of organizational information seeking on the web by integrating Ellis' (Ellis, 1989) six stages of information seeking (starting, chaining, browsing, differentiating, monitoring, and extracting) with Aguilar's (Aguilar, 1988) four modes of scanning (undirected viewing, conditioned viewing, informal search, and formal search). Each of these tasks helps to satisfy part of an organization's information needs.

Each new piece of information that is gathered can provide new ideas, suggest new directions and change the nature of the information need (Bates, 1989). This leads to behaviour that Bates refers to a *berrypicking,* reflecting the incremental collection of pieces of information that, as a whole, help to satisfy the information need. The choices made at each step are guided by the seeker's assessment of what is most likely to produce useful information as they *forage* for information (Pirolli and Card, 1995). In an environment like the web, this assessment is based on cues such as hyperlink text that provide an *information scent* which help users make cost/benefit trade-offs in their choices. In line with research suggesting that search is made up of sequences that are affected by information discover, Belkin, Cool, Stein and Theil (Belkin et al., 1995), created a set of 'scripts' that describe the typical paths taken by 16 different types of users, including switch points between them for when discoveries change their user type.

2.4 Information retrieval tasks

Typically, within each information retrieval task, the searcher formulates queries, examines results, and selects individual documents to view. As a result of examining search results and viewing documents, the searcher gathers information to help satisfy the immediate information-seeking problem and eventually the higher level information need.

The strategies and tactics that searchers use are affected by the capabilities provided by the search interface, user expertise and the task (Bates, 1990; Golovchinsky, 1997; Navarro-Prieto et al., 1999). Tactics are individual actions or sequences of actions (often called moves) taken to further the search (Bates, 1979b, a; Marchionini, 1995). Experts can use their deeper knowledge of the structure of the information and the task goal to guide their tactics as they search, whereas novices may start with keywords or categories drawn from the surface terms describing the information need and available in the search interface (Navarro-Prieto et al., 1999). Searchers can take numerous actions while examining search results (Bates, 1990; Fidel, 1985; Garcia and Sicilia, 2003; Marchionini, 1995; Shneiderman and Plaisant, 2004; Wildemuth, 2004).

Specific actions supported by a web search interface can be discerned by analyzing the structure of text and hyperlinks on a search result page. For a typical search result page showing a ranked list of results, this yields the set of actions listed in Table 1.

Each action involves cognitive and physical effort and can result in visual changes in the interface or changes in task, domain, or category knowledge (cognitive changes). The visual changes enable cognitive changes by making information visible on the screen. The cognitive changes are necessary to make progress on the information problem and are reflected in transitions between activities. For example, while examining results, searchers may scan a screen of results, causing them to identify additional query terms, causing a transition to the formulate query activity. In this analysis, actions that require visual scanning and/or moving the mouse without clicking are classified as low effort because they involve little physical effort, and they do not result in major changes to the display, thus minimizing cognitive effort. Actions such as selecting a result to view or scrolling the screen require a moderate amount of cognitive or physical effort because they require clicking and reorienting as the visual presentation changes. Issuing a new query requires a high amount of effort because of the cognitive effort required to formulate the query and the need to reorient when the new set of search results is displayed.

Table 1. Actions available to searchers when evaluating a typical search result list.

Action	Effort	Visual changes	Cognitive changes
Scan one screen of results list	Low physical; low-medium cognitive (depends on type of scan)	None	• Identify page to view • Assess results overall • Identify additional query terms • Refine information need • Refine information problem • Extract useful information
Scroll screen	Medium	Shift visible subset of search results	None
Select next or previous page of results	Medium	Shift visible subset of search results	None
Select a result to view specific web page	Medium	Bring web page into view	None
Reformulate query	High	Generate new set of search results	None
View specific web page	Variable	Variable	• Identify additional query terms • Refine information need • Refine information problem • Extract useful information

As an illustration of how the interface affects the actions that are available, consider the addition of a categorized or faceted overview to the list of search results. Adding a categorized overview to search results changes the information that is available and the actions searchers can take with low or moderate physical and cognitive effort. The categorized overviews can provide capabilities such as:

- The overview presentation – a visual or graphical representation of the categories represented by the search results
- Hyperlinks to narrow and broaden the displayed set of search results
- When the pointer is placed over a category, the corresponding search results in that category are highlighted and a pop-up window is displaying with a list of non-empty sub-categories
- When the pointer is placed over a search result, the corresponding categories of which that result is a member are highlighted

Table 2 summarizes the actions afforded by these design elements, the effort required, and their visual and cognitive effects.

Table 2. Additional actions available to searchers when evaluating search results with categorized overviews.

Action	Effort	Visual changes	Cognitive changes
Scan categorized overview	Low physical; low-medium cognitive (depends on type of scan)	None	• Identify category to consider • Assess results overall • Identify additional query terms • Refine information need • Refine information problem • Extract useful information • Assess match between categories and information need
Select category to narrow or broaden results	Medium	Filter visible results, limiting to members of selected category	None
Move pointer over result	Low	Highlight category membership	• Identify categories to consider • Assess results overall • Identify additional query terms • Refine information need • Refine information problem

Move pointer over category	Low	Highlight results in category (currently visible results only) and display subcategories	• Identify page to view • Assess results overall • Identify additional query terms • Refine information need • Refine information problem

The following sections of the report describe different approaches and evaluations of software that provide support for the actions identified in Table 1 and Table 2 in many different ways.

3. Search Environments Review

3.1 Adding Classifications

One of the main streams of research into enhancing search environments has been to use annotations, or classifications, to the documents or collections. One of the main challenges in adding classifications, however, is increasing scale of collections and the associated cost of annotating each document within them. Below we review the different approaches to adding classification to collections and overcoming the challenges briefly stated above.

3.1.1 Hierarchical Classifications

One early research project, the SuperBook interface, showed the benefits of using categories in a document collection, by organizing search results within a book according to the text's table of contents; here the book is the collection and the pages of the book are the documents. An evaluation found that it expedited certain searches without loss of accuracy (Egan et al., 1989). This categorisation approach has been used in many cases and has shown success in both fixed and uncontrolled collections, however the usual approach for the latter is to model the first and allow document owners to assign their own documents into the hierarchy.

An example of a fixed and managed dataset maybe the genre based classification of music in Amazon's online music store (www.amazon.com), where every CD is assigned to one or more categories of music, including: Pop, Rock, and Indie. Here there is incentive for Amazon, as the collection owner, to annotate the collection in this way to make it easier for their clients to find the music they want and, subsequently, encourage sales. Allen (Allen, 1995) investigated two such digital library interfaces, the Dewey Decimal System and the ACM Computer Reviews system, and showed that both used hierarchical classification effectively for organising the collections.

In the same paper, Allen (Allen, 1995) discusses the potential for use of Internet-wide information samples. Examples of hierarchically classified structures for the web, as an ever increasing and unmanaged set of documents, are Google Directory (directory.google.com) and Yahoo Directory (dir.yahoo.com). In both of these directories, the individuals that own or manage documents on the web can submit their websites for inclusion under different parts of the directory. This approach has been popular in the past but has suffered where document owners do not submit their websites for inclusion.

In an attempt to remove this annotation cost and limitation of classification systems, (Kules et al., 2006) took a simple result set of federal agency/department reports, by mapping a set of URL prefixes to a known finite list. This approach took a known pattern from the data and used it to effectively categorize search results by government agency and department. In a more advanced approach (Chen and Dumais, 2000), showed that machine learning techniques can be applied to known categorised data items to automatically assign new web documents into categories. Using the documents that already exist in the LookSmart (www.looksmart.com) directory as a training-set, the machine-learning algorithm successfully and automatically categorised the results of a Yahoo search. The Support Vector Machine (SVM) algorithm (Joachims, 1997) achieved 70% accuracy with the categorisation provided by human participants, where the remaining 30% included partially matching annotations. The user study showed a strong preference for the added categorisation provided by the process. Other approaches to automation exist and are mainly described in the automatic clustering section below.

The benefits of applying such categorisations have been proven numerous times in research. For question answering tasks, Drori and Alon shown that search results augmented with category labels produced the fastest performance and were preferred over results without category labels (Drori and Alon, 2003). Dumais, Cutrell and Chen (Dumais et al., 2001) also studied the effect of grouping search results by a two-level category hierarchy and found that grouping by a well-defined classification speeds user retrieval of documents.

In a similar approach to the book-search mentioned above, the Cha-Cha system organizes intranet search results by an automatically generated web site overview (Figure 2). It reflected the underlying structure of the web site, using the shortest path from the root to each document to dynamically generate a hierarchy for search results. Preliminary evaluations were mixed, but promising, particularly for what users considered "hard-to-find information" (Nation et al., 1997). The WebTOC system (Figure 3) provides a table of contents visualization that supports search within a web site, although no evaluation of its search capability has been reported (Hearst et al., 2002; Yee et al., 2003). WebTOC displays an expandable/collapsible outliner (similar to a tree widget), with embedded coloured histograms showing quantitative variables such as size or number of documents under the branch.

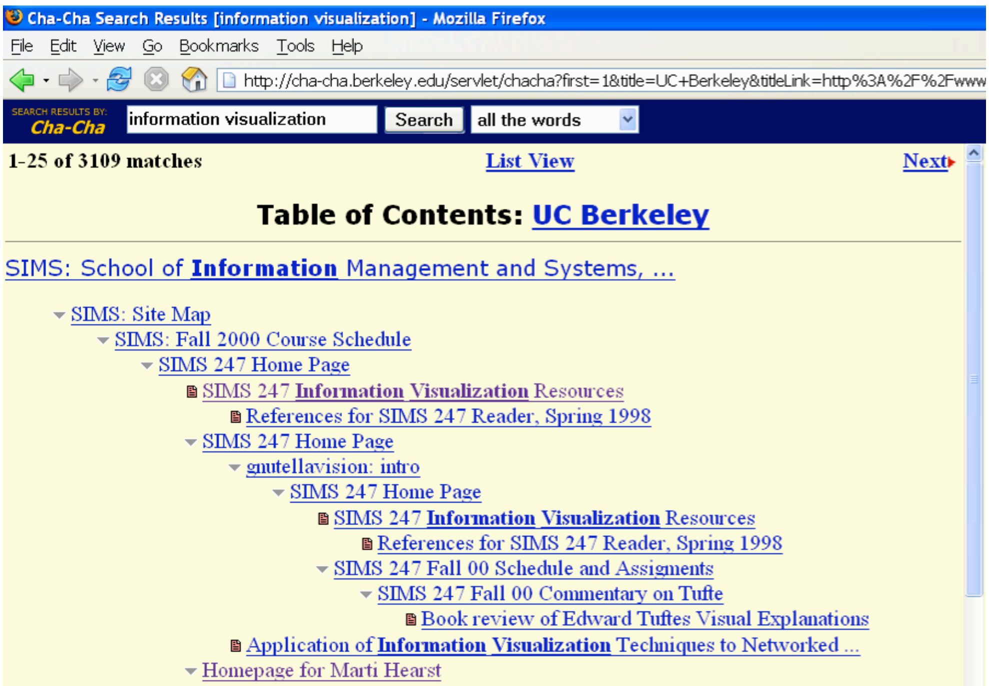

Figure 2. The Cha-Cha system organizes intranet search results by an automatically generated web site overview.

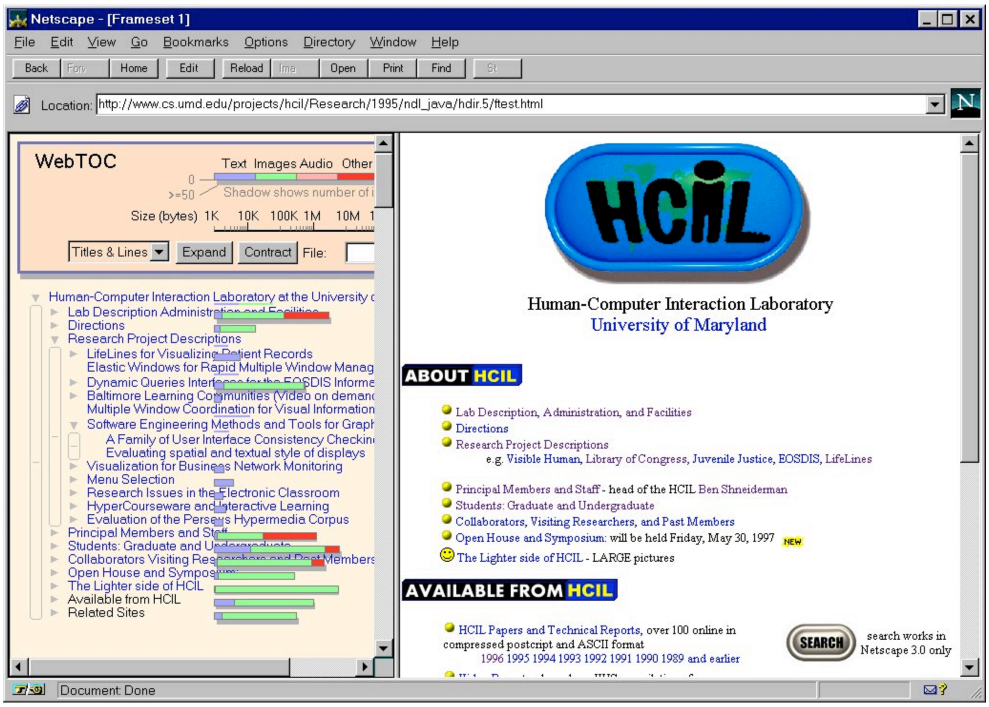

Figure 3. The WebTOC system provides a table of contents visualization that supports search within a web site.

3.1.2 Faceted Classifications

Another approach to helping user find documents based on multiple facets, such as thematic, temporal, and geographic, etc. builds on the traditional library science method called faceted classification. For example, restaurants can be grouped into facets such as by location, price, food type, customer ratings, and size, with 2-20 attribute values per facet. Early versions, called query previews, showed the number of documents having each attribute value. Query previews were updated with new values so as to prevent users from submitting searches that would produce zero results (Donn et al., 1996; Plaisant et al., 1999; Tanin et al., 2000). Facet classification allows users to apply relevant constraints on their search in the context of the current problem and existing knowledge. For example, if users know their own budget and a minimum specification required for a computer, then they can apply constraints in these facets. After applying their constraints they can then see all the relevant computers and possibly choose between them based upon the facets that remain unused.

Many systems have applied this approach in some way, where one hierarchical classification system has not been expressive enough for the documents in a collection. For example, before prototyping a faceted system (express.ebay.com), eBay already allowed users to apply constraints such as price, colour, and size; the available constraints depended, of course, on the type of object being sought. Whereas eBay indexes its own auctions, websites such as shopping.com provide a faceted search over many shopping sites from the whole web.

Flamenco (Figure 4), the source of the code used in eBay Express, is a clear example of the features provided by faceted search. Providing interfaces to fixed collections, including art, architecture, and tobacco documents, Flamenco presents faceted hierarchies to produce menus of choices for navigational searching (Yee et al., 2003).

A selection made in any facet is added to a list of constraints that make it clear to users what is forming the list of results that are being shown.

A usability study compared the Flamenco interface to a keyword search interface for an art and architecture collection for both structured and open-ended, exploratory tasks (Yee et al., 2003). With Flamenco, users were more successful at finding relevant images (for the structured tasks) and reported higher subjective measures (for both the structured and exploratory tasks). The exploratory tasks were evaluated using subjective measures, because there was no (single) correct answer and the goal was not necessarily to optimize a quantitative measure such as task duration.

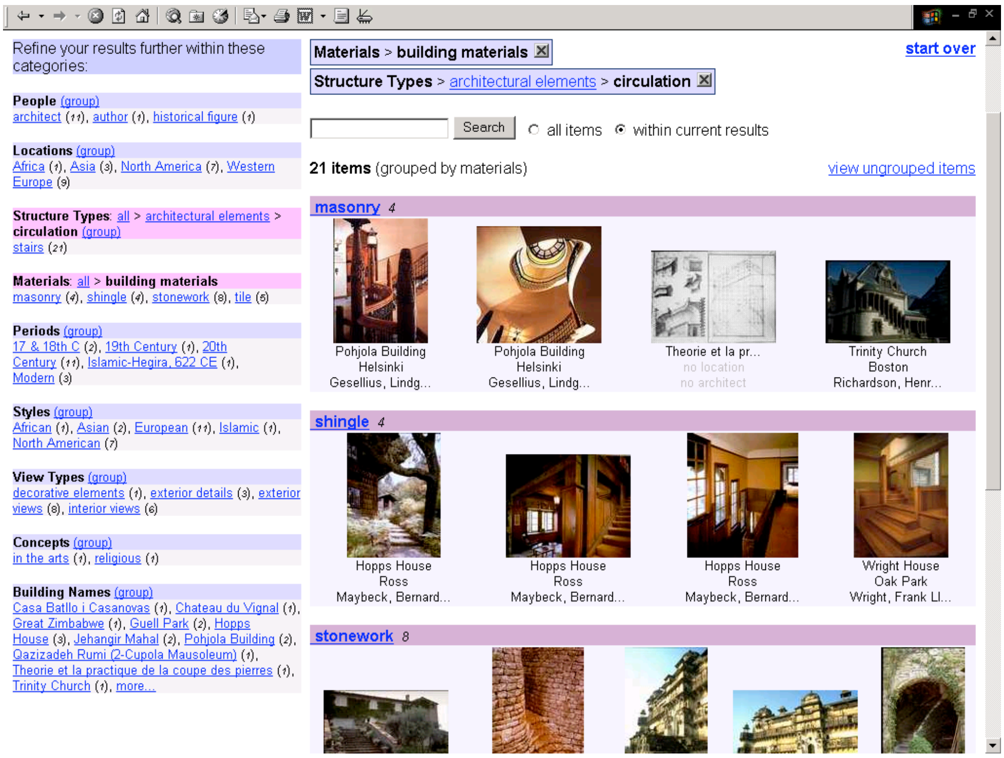

Figure 4. The Flamenco interface permits users to navigate by selecting from multiple facets. In this example, the displayed images have been filtered by specifying values in two facets (Materials and Structure Types). The matching images are grouped by subcategories of the Materials facet's selected Building Materials category.

The success seen by Flamenco, having provided faceted classifications to assist search, have been used in many commercial and academic projects. Endeca Search (www.endeca.com) is a commercial company that provide faceted search, branded as Guided Navigation, to large businesses including Wal-Mart and IBM. epicurious, shown in Figure 5, provides both faceted and keyword search over food recipes in a style that is very similar to Flamenco's experience. The lead researcher of Flamenco, Marti Hearst, refers to epicurious as a good example of faceted browsing in commercial conditions (Hearst et al., 2002).

Figure 5: epicurious is a commercial recipe website that uses Flamenco style faceted search in conjunction with keyword search.

Huynh et al. (Huynh et al., 2007) has developed Exhibit, a faceted search system that is similar to flamenco in many ways, but has some significant developments. One key advance is that, where a selection in Flamenco filters all the facets, a selection in Exhibit filters all the other facets and leaves the facet with the selection unchanged. This provides two benefits: first, users can easily change their selection and second, users can make multiple selections in one facet. This allows users to see, for example, the union of all the red and blue clothes, rather than just red or just blue.

Figure 6: Exhibit faceted search interface

The Rave Browser (Zhang and Marchionini, 2005) takes another approach to faceted search. One notable difference is that multiple selections lead to their intersection of results being displayed. Another feature that the Rave browser provides is a preview of the affect of clicking has on other facets. Graphical representations behind each item in each facet show how many documents can be found by selecting it. When a user hovers over any item in any facet, the size of the bar in the graphical representations are reduced to indicate how many documents will remain under each annotation should the user make the selection. This technique revives the query preview strategy, which is a helpful alternative the simple numeric value indicators (Wilson and schraefel, 2006) that are included in most classification-based systems. Aside from the graphical representation, the preview of the affect by simply hovering (or 'brushing') over the item is a technique that is being included in many new projects. A new version of the Rave Browser is expected in the near future.

Figure 7: Rave Browser interface

mSpace (schraefel et al., 2006) represents yet another type of faceted search: column interfaces. Like iTunes, mSpace presents facets in a left-to-right set of columns. Each column is fully populated so that users can still make a selection in any facet, but then only the columns to the right filter. This allows the facets to represent additional information that would otherwise be lost in faceted search. In their classical music example, where the facets are Era, Composer, Arrangement, and Piece (from left to right), if users select a Composer, they see a filtered list of the arrangements (s)he used and then a list of the pieces (s)he composed. When users make a second selection in the Arrangement, other forms of faceted search would remove all the composers that did not use that arrangement. In mSpace, users are still able to see all the arrangements that the selected composer used, and all the pieces of the selected arrangement.

The functionality of mSpace as described above, would mean that there is a type of information that is not conveyed by mSpace but seen in other forms of faceted search: the Era of the selected composer. This missing information is not a problem in mSpace, unlike other column-faceted browsers like iTunes, because the related items in facets to the left of a selection are highlighted. This allows users to see which era the selected composer is in, whilst still allowing them to get the added facts provided. The affect of this leftward highlighting, named Backward Highlighting (Wilson et al., 2007) is that users incidentally learn more about the structure of a collection, which can make search easier for future searches in the same domain.

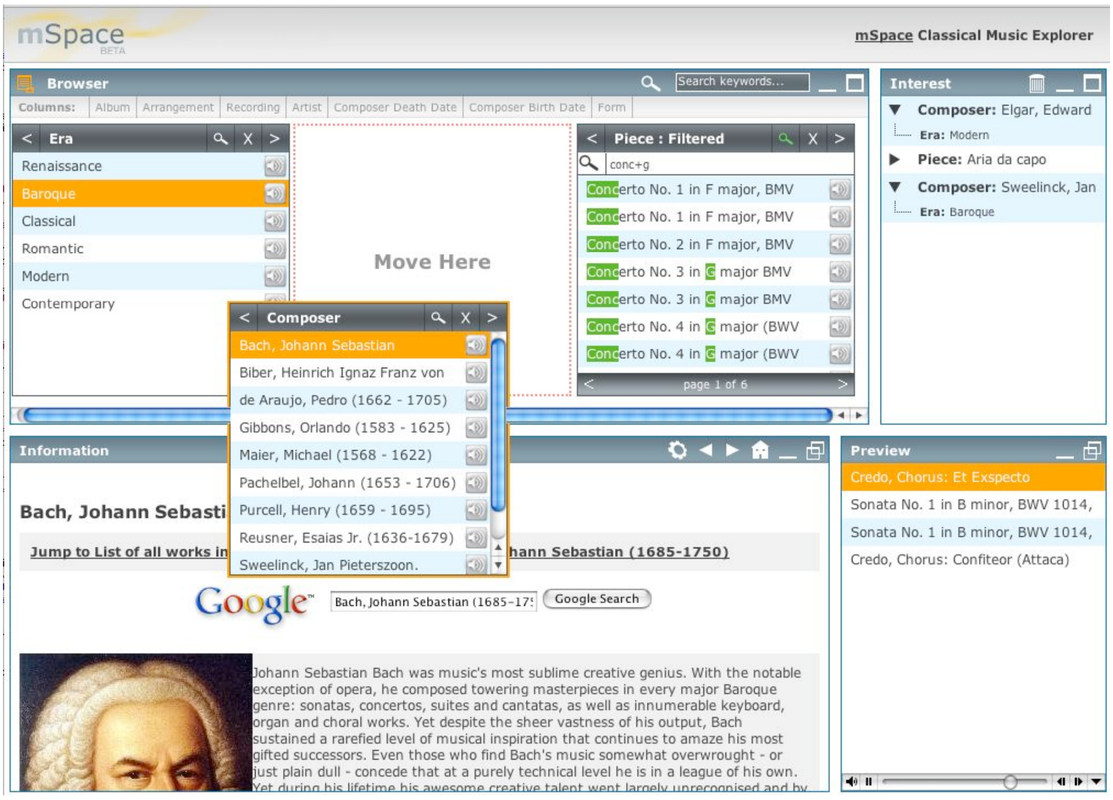

Figure 8: mSpace faceted column browser

Finally, given the importance placed on the direction and order of facets in columns, mSpace allows users to reorder, remove and supplement the facets shown, using any of the facets that are available from the collection. This allows users to say that they would rather know all the composers that used a given arrangement than all of the arrangements used by a given composer. Another unique aspect of mSpace is that it assumes that a selection that is different to a previous selection in a facet is a change of selection and not a multiple selection. Consequently, users can quickly compare the difference between two items in a facet. The default assumption that users are changing their selection supports the concept of answering subjunctive questions about the collection (Lunzer and Hornbæk, 2003), which means simply to compare the outcomes of multiple actions.

A recent longitudinal study of mSpace (Wilson and schraefel, 2008a) has indicated that this more complex form of faceted search is easy to learn and is thereafter perceived as a very powerful system, receiving positive subjective views. A log analysis showed that 50% of participants used facets in their first visit to the site and 90% in their second visit. Over the whole month-long period, there were more interactions with facets than individual keyword searches. Further, given that facets allow users to produce complicated queries than basic keyword searches, faceted searches represented two times the number of Boolean searches and three times the number of advanced searches.

Each of the faceted classification examples so far have been on fixed collections, but some research into faceted browsing has been looking at unknown document sets like the web. The SERVICE (SEarch Result Visualization and Interactive Categorized Exploration) search system (see http://www.cs.umd.edu/hcil/categorizedovervew) couples the typical ranked list of web search results list with automatically generated facets (Kules and Shneiderman, to appear). Clicking on a category filters (or narrows) the displayed results to just the pages within that category. Moving the pointer over a category highlights the visible search results in that category in yellow. Moving the pointer over a result highlights all the categories in the overview that contain the result.

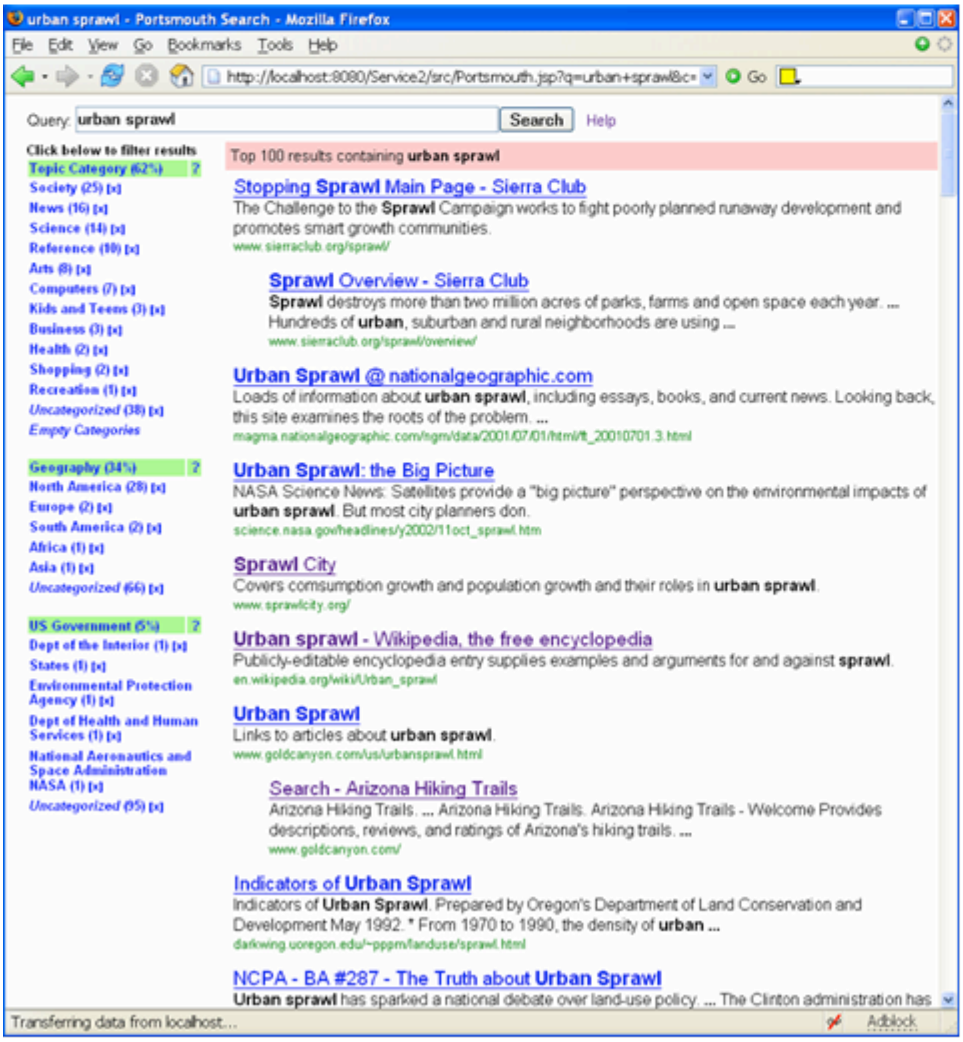

Figure 9. The SERVICE web search interface (Kules and Shneiderman, to appear).

The facets are automatically generated by applying fast-feature classifiers (Kules, 2006b) over the top 100 results of a Google query, and organised them into known possible categories drawn from the Open Directory Project (ODP) and a database of US Government web sites (http://www.lib.lsu.edu/gov/tree): Topic, Geography and US Government. A similar project: Dyna-Cat (Pratt, 1997) also automatically produced facets for sets of search results, showed that not only was there improvements in objective and subjective measures, that users were 50% faster in fact finding tasks using Dyna-cat over typical ranked list keyword search interfaces.

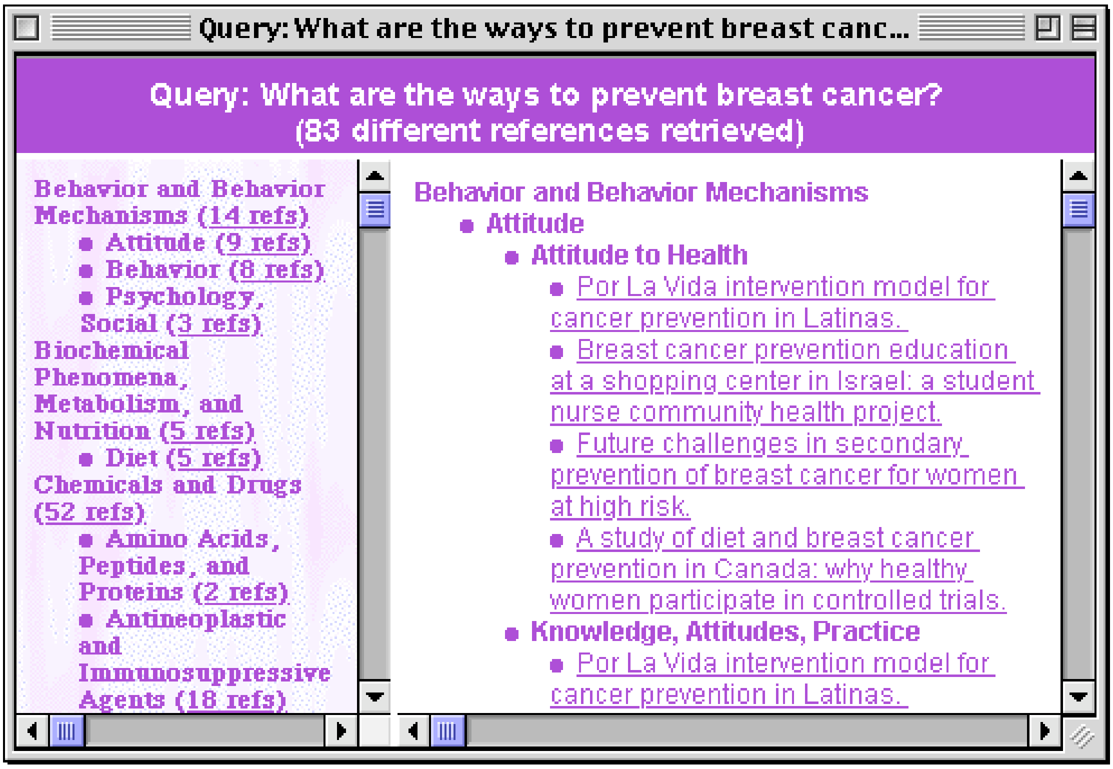

Figure 10: Dyna-Cat search interface

Northern Light (www.northernlight.com), a commercial search service, provides a similar capability by grouping results in their Custom Search Folders. Exalead (exalead.com) is another project that successfully organizes search results according to categories drawn the Open Directory Project, and presents them along side search results in a publicly available web search engine. PunchStock search, shown in Figure 11, presents a faceted view of personal image collections. The NCSU library, Figure 12 also uses facets to enhance the search if their collection.

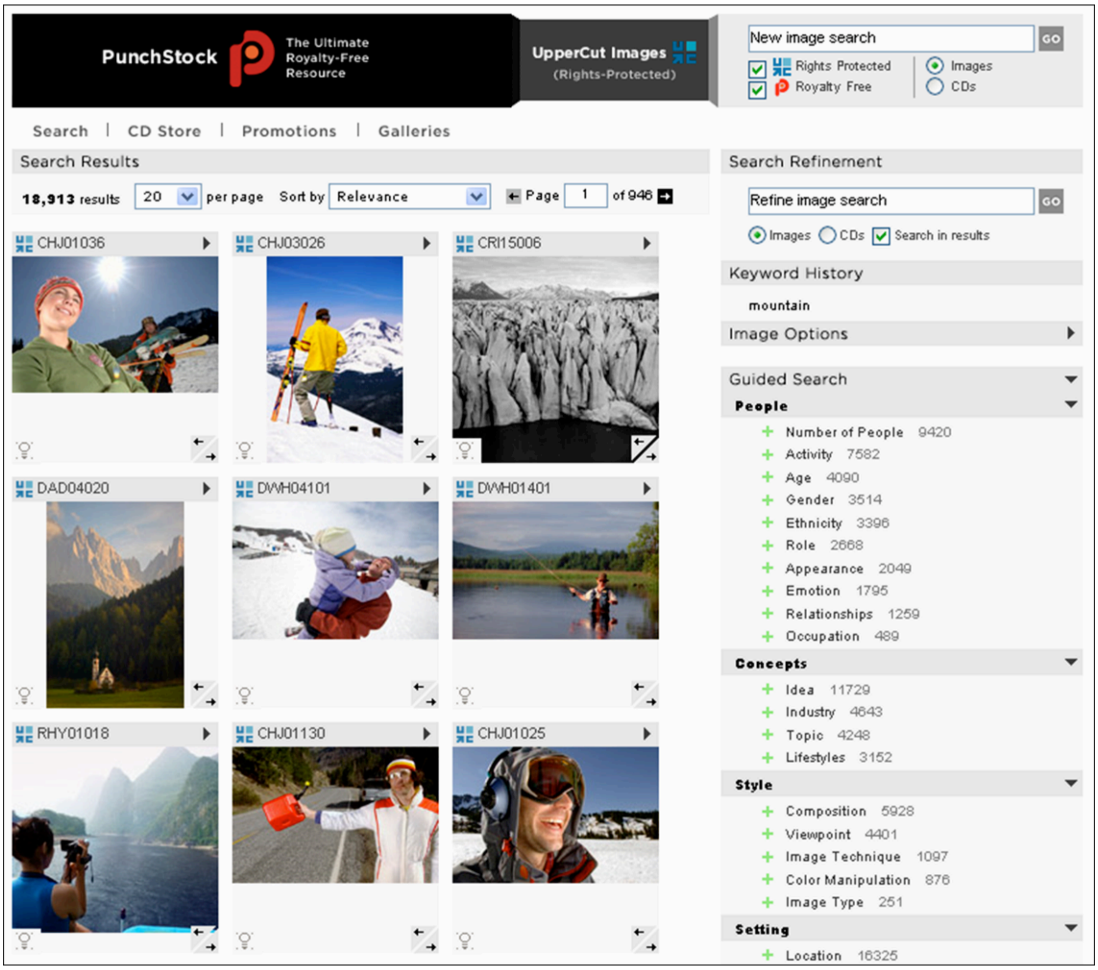

Figure 11. The PunchStock photo search interface provides categorized overviews of photo search results.

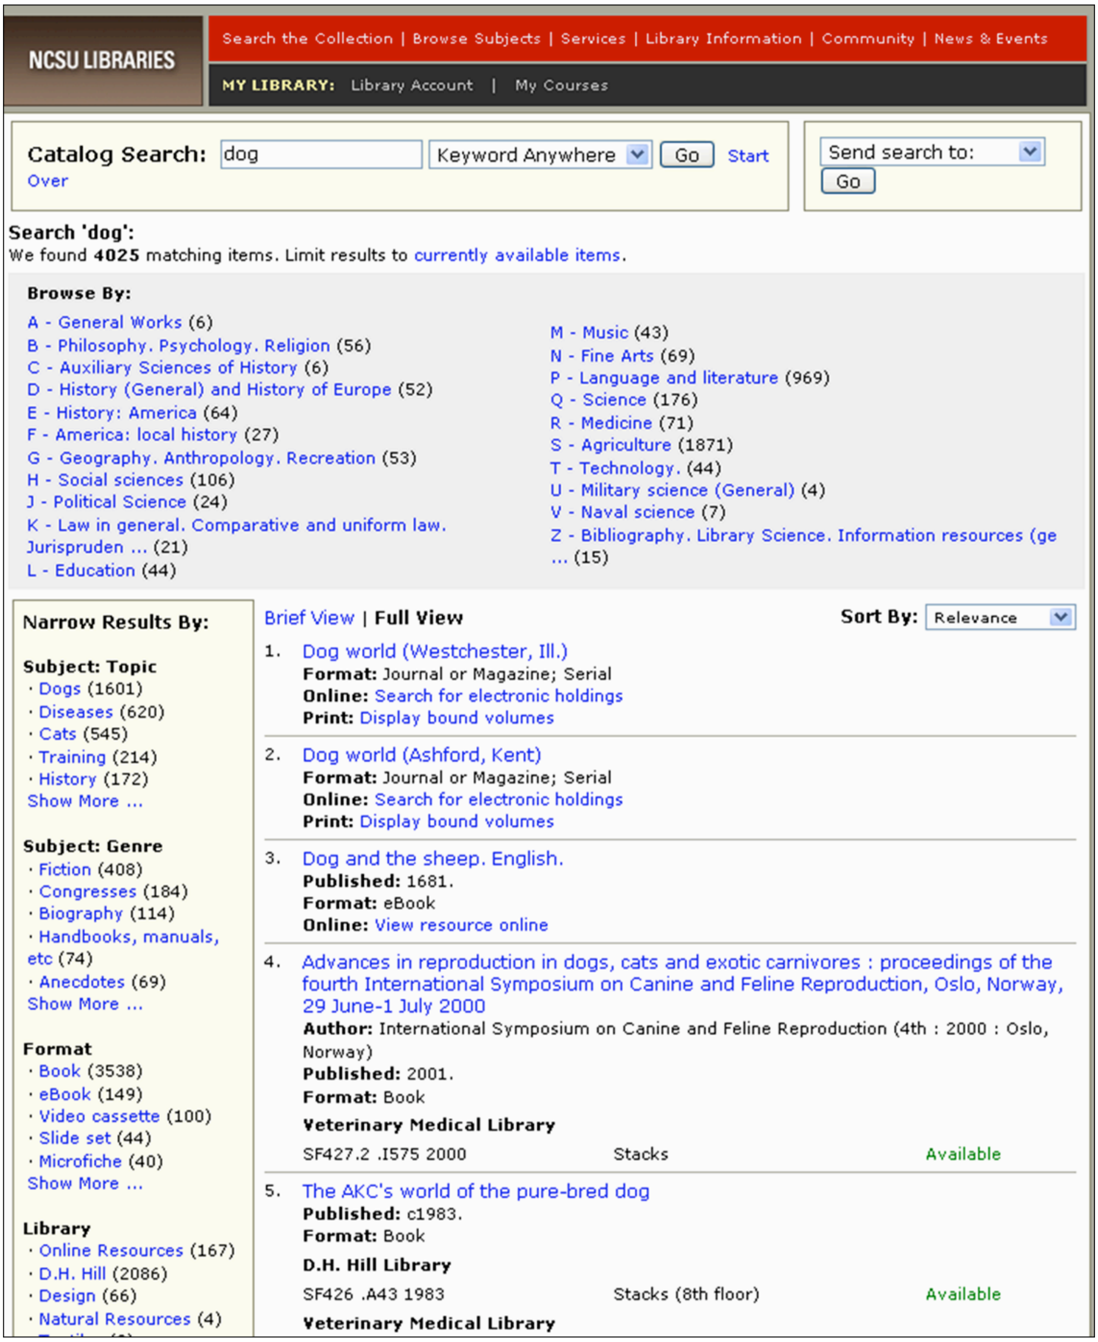

Figure 12. The NCSU library catalog provides categorized overviews of search results using subject headings, format, and library location.

3.1.3 Automatic Clustering

Where collections have been large or unmanaged, faceted classifications have been less successfully applied. The examples where it has been used above provide very generic facets, and often the facets that would be important depend on the query or subject of the users' varying goals. Another approach is to automatically identify attributes of a collection or result set that are important, rather than explicitly producing labelled annotations. This approach is called clustering, and has shown success under popular information retrieval metrics such as precision and recall (Hearst, 1999a) or task completion time (Rivadeneira and Bederson, 2003).

A one-level clustered overview was found helpful when the search engine failed to place desirable web pages high in the ranked results, possibly due to imprecise queries (Stoica and Hearst, 2004). Clusty (www.clusty.com) uses the clustering technique to produce an expandable overview of labelled clusters (Figure 13). The benefits of clustering include domain independence, scalability, and the potential to capture meaningful themes within a set of documents, although results can be highly variable (Efron et al., 2004).

Generating meaningful groups and effective labels, however, is a recognized problem (Shneiderman et al., 1998) and where possible (usually, but not exclusively in fixed or managed collections) having an information architect design optimal annotations or labels, will provide a better search interface (Hearst, 2006). Greene et al. (Greene et al., 2000) extracted a category hierarchy from WordNet (Miller, 1995) using keywords from the document collection. Similarly, Drori and Alon (Drori and Alon, 2003) investigated the use of semi-automated methods, combining k-means clustering and statistical classification techniques to generate a set of categories that span the concepts of the Bureau of Labor Statistics web pages and assign all to these categories. They found that concept learning based on human-supplied keywords performed better than methods using the title or full-text.

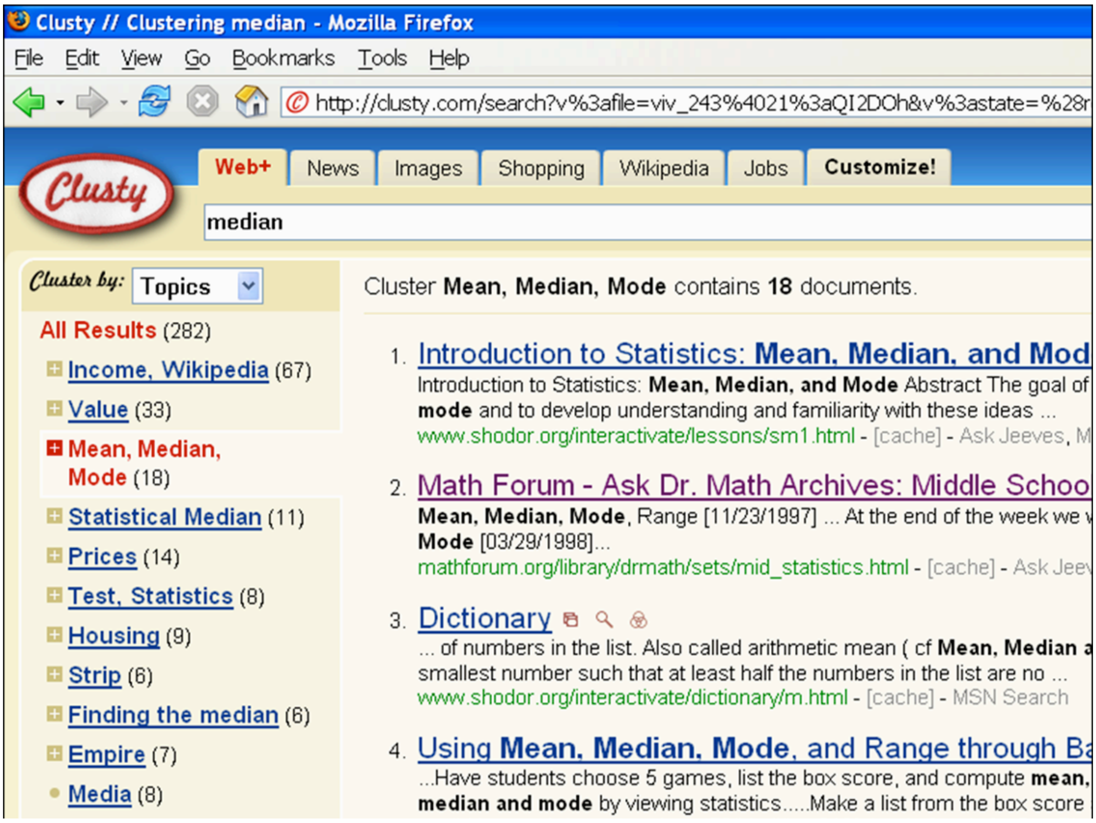

Figure 13. The Clusty metasearch engine uses automated clustering to produce an expandable overview of labeled clusters.

3.1.4 Social Classifications

Recent developments in web technologies, such as web2.0, have led to a rise in social classifications such as bookmarks, tagging, and collaborative feedback ratings. For example, the Yoople search engine allows users to explicitly move the search results up or down to provide relevance feedback into the weightings used in future searches. Very little research, however, has explicitly proven the benefits, even though the notion has become increasingly popular on the web. Some research is emerging, Millen et al. (Millen et al., 2006) have investigated the use of social bookmarking in enterprise search software and an eight-week field trial has shown positive results.

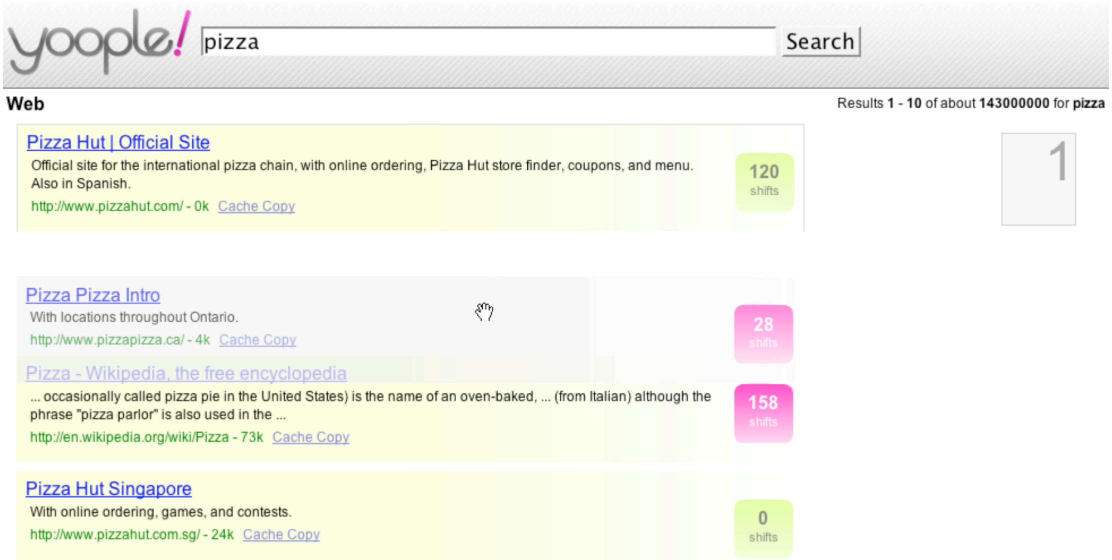

Figure 14: Yoople Search, users can drag items in the results list to indicate where they think a result should appear.

In Yoople (Figure 14) mentioned above, the classification produced is only numeric and affects the background search algorithms. Social tagging, however, has allowed communities of people to develop flat classification schemes for collections. Flickr (www.flickr.com), a photo archiving website, depends almost entirely, although photo names and descriptions can be included, on user tagging to return images that relate to a keyword search.

A negative side to such tagging classification schemes is that they are hard to present to users. Therefore, they are usually only used to aid keyword search rather than to help users interactively browse through documents, like in some of the faceted and category based systems listed above. A popular example of presenting a flat set of tags to a user has been to foreground popular tags in a tag cloud Figure 15 (Hassan-Montero and Herrero-Solana, 2006).

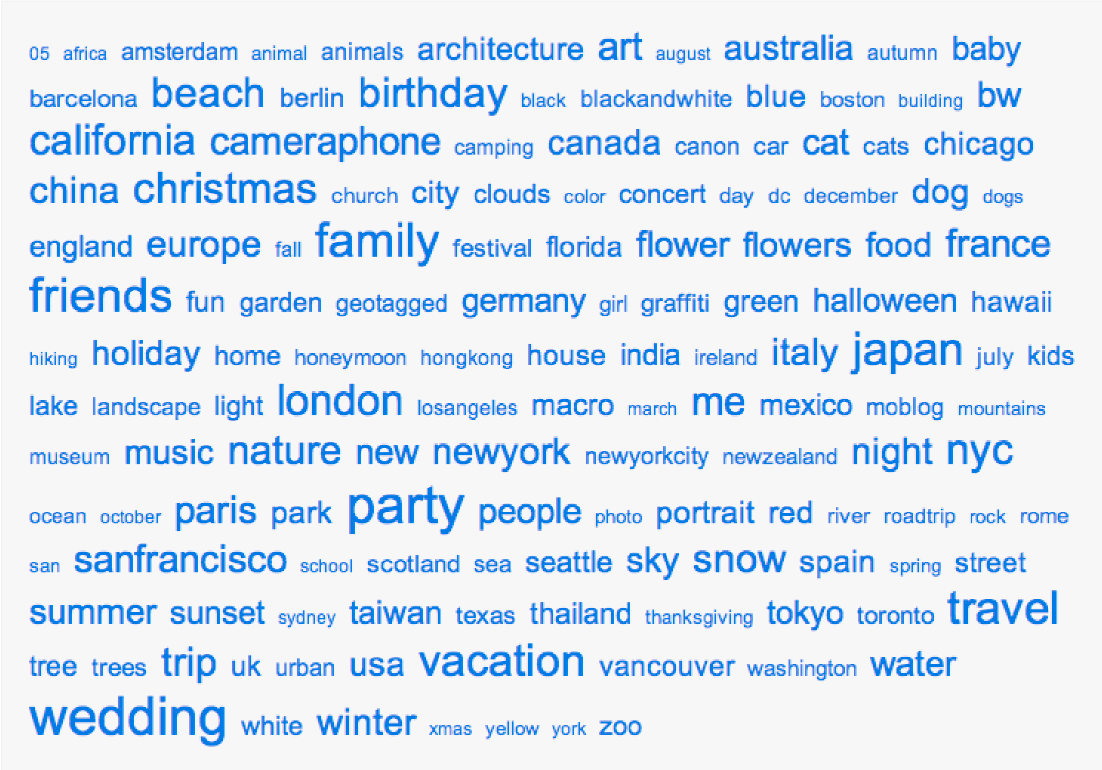

Figure 15: A tag cloud taken from Flickr website, included in (Godwin-Jones, 2006).

Some systems are trying to take the benefits of tagging to produce what are known as folksonomies (Hotho et al., 2006) or community driven classification systems. Wu et al. (Wu et al., 2006) present some design prototypes that overcome some of the challenges in converting flat tags into a structure classification system. While this research is in its very early stages, systems such as mSpace and the Phlat browser (Figure 16), allow users to use tags in conjunction with faceted search to improve search (Cutrell and Dumais, 2006). While both tagging and faceted search has proven successful, no empirical evidence, known at the time of writing, has been produced to show that their combination provides specific advantages.

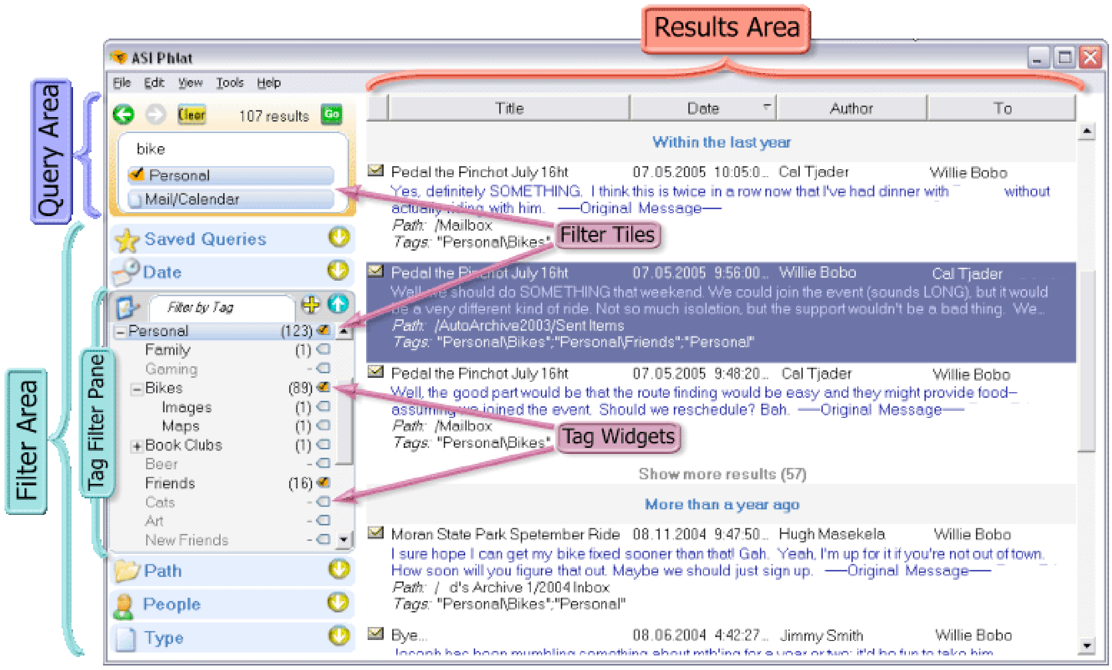

Figure 16: Phlat personal information management browser

3.2 Result Organisation

At the opposite end to allowing users to define their own search query, which has been discussed above, is the presentation of the result sets and items that will help users to identify the specific result item(s) that help them achieve their goals. There is ample research, discussed in section 2.2.1, which specifically looks at the best way to present a single result in a linear set. In the following subsections, however, we discuss the research that looks beyond a linear results list and visualises a result set to help users find the specific results they are looking for.

3.2.1 Result Lists

The most common method to show a result set, seen in most web search engines, is to provide a simple list of results. The order that the result items are listed is determined by some metric, usually based on how relevant the result is to the search terms used. The importance of constructing this algorithm well has been motivated many times in research where users regularly enter single term ambiguous queries (Wang et al., 2003), and view only a few results (Jansen and Spink, 2003) and rarely stray past the first page of results (Beitzel et al., 2004). The accuracy and efficiency of such algorithms, however, has been research for many years in the information retrieval space, and is not covered in the scope of this report; (Becks et al., 2002) present a good discussion of result ordering, including occasions where ranked relevance may not be the most appropriate ordering.

The representation of each result has also received much research and is especially important when research has shown that the acceptance of search systems has been significantly reduced by unclear or confusing representations (www.antarctica.net**)**. In Google, each search result has: name (also a link to the finished document), a sample of text from the document, an URL for the document, and the size of the document. Additional information, such as previous visit dates and counts, can be added if a user account is available. Work by (Chen et al., 1998) investigated the structure of result representations and provides a framework for considering the presentation of search results. In one of the more significant publications in this area, White et al. (White et al., 2002) showed that best objective results occurred when the text sample in a representation included the query-terms used in the original search. This allows users to see the context of their query in each result item, so that they can best judge its value in viewing the whole document. This was backed up by Drori and Alon (Drori and Alon, 2003), who showed objective and subjective benefits for both query-relevant snippets and associated categories.

Other search engines, including ASK (www.ask.com) and Exalead (www.exalead.com - Figure 17), also include thumbnails of the document represented by each result. Research into the benefits of page thumbnails (Woodruff et al., 2001) has shown most advantage when users are returning to a previous search to find a result from the previous session. With these conditional objective benefits, however, subjective results have been positive.

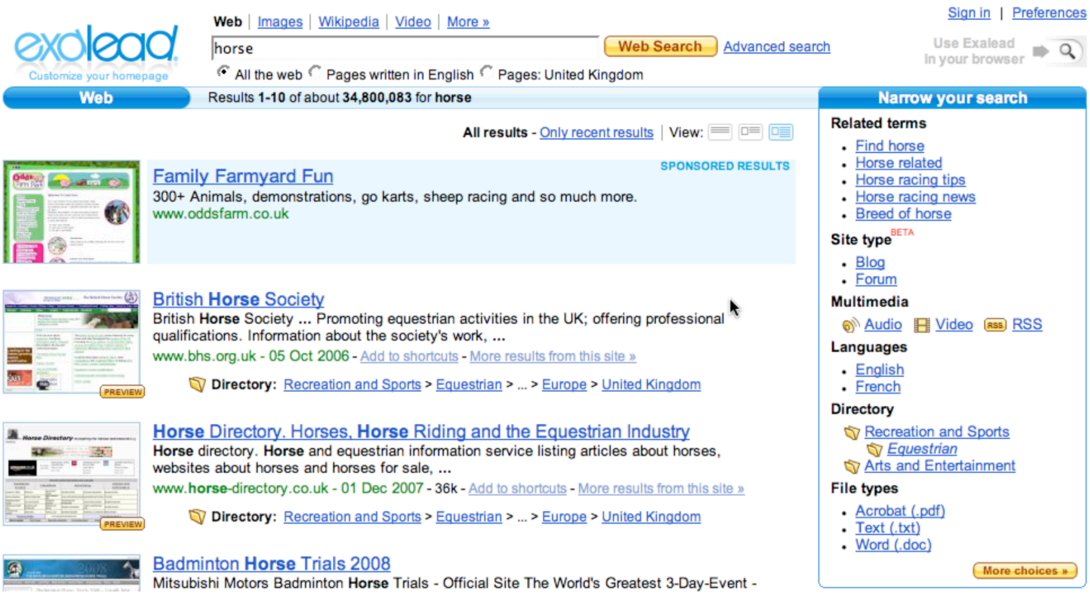

Figure 17: Exalead search provides images and categories with each search result, and facets to help users narrow the results

3.2.2 2D Result Representations

With the aim of providing useful overview visualisations of result sets to users, much research has aimed at taking information visualisation techniques and applying them to search results. One example is Self-organising maps (SOMs), originally produced by Kohonen (Kohonen, 2001). SOMs automatically produce a two-dimensional (2D) visualisation of data items, according to the attributes that they share, using an unsupervised machine-learning algorithm. Consequently, the SOM approach applies well when visualising automatically clustered documents.

Au et al. (Au et al., 2000) used SOMs have been used to support exploration of a document space to search for patterns and gain overviews of available documents and relationships between documents. Chen, Houston, Sewell, & Schatz (Chen et al., 1998) compared a SOM with the Yahoo! Entertainment category, for both browsing and searching tasks. They found that recall improved when searchers were allowed to augment their queries with terms from a thesaurus generated via a clustering-based algorithm. Other similar work by Kules (Kules, 2006b) has also shown positive results in using SOMs to overview search.

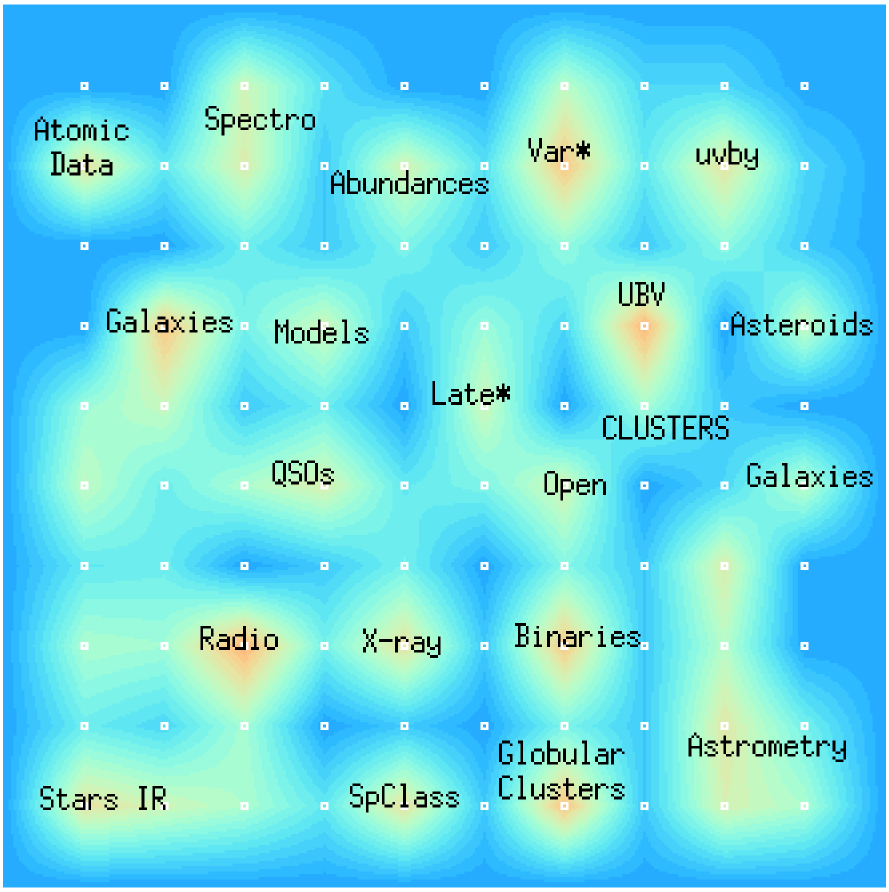

Figure 18: A SOM of search results, image taken from http://vizier.u-strasbg.fr

An alternative 2D organization on result sets is known as a treemap (Ginsburg, 2004; Lamping and Rao, 1996; Perugini et al., 2004). In a treemap, the clusters of documents are grouped using a space-filling algorithm. As shown in Figure 19, the 2D space is divided up into the top-level items in the hierarchy produced by the clustering algorithm, where the number of search results found within the category dictates the size of each part. Each of these sections is divided up into their children within the hierarchy, where the number of search results again dictates the size of the subparts. This process is repeated as necessary according to the hierarchy and space available. With this, users can access categories at any level of the hierarchy and see the results associated with that cluster.

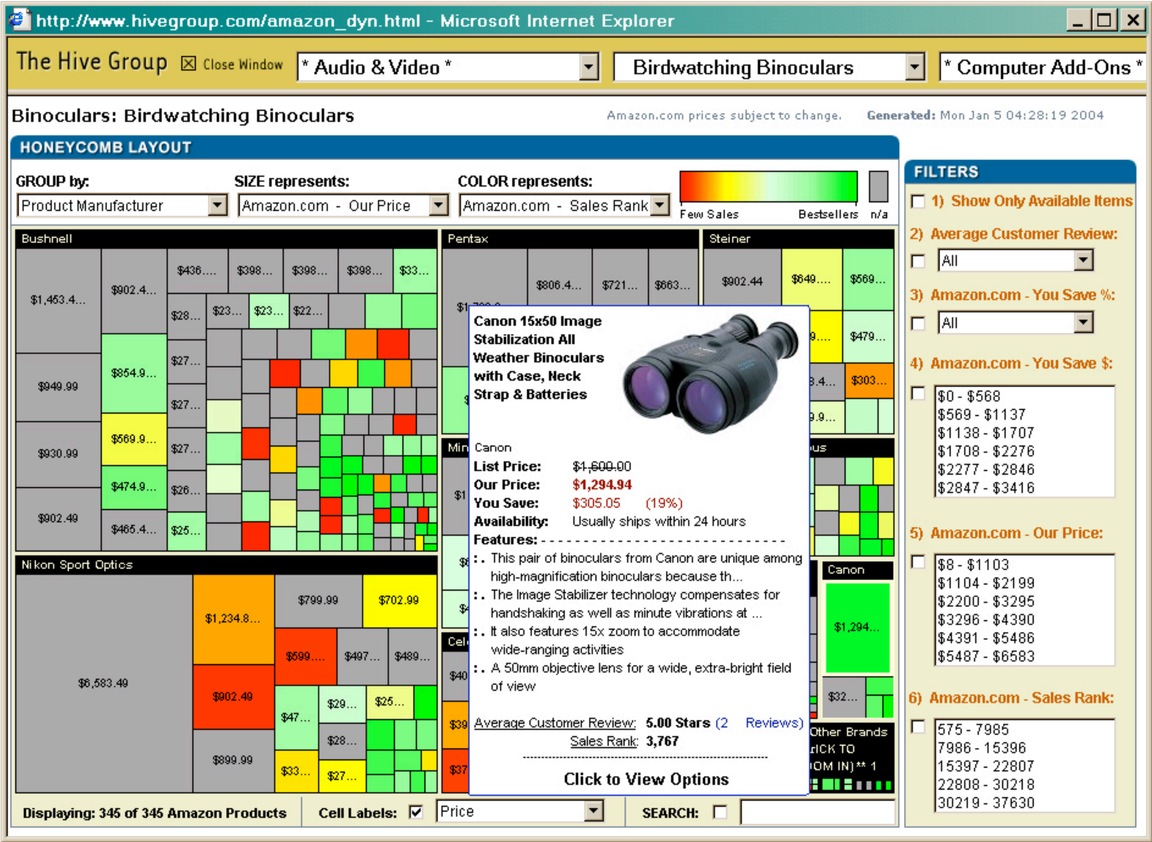

Figure 19. The HiveGroup built a demonstration shopping application that allows users to explore Amazon.com products. (image from http://www.cs.umd.edu/hcil/treemap-history/hive-birdwatching.jpg).

Figure 20. This overview of web search results uses a treemap. Nesting is used to show both top and second-level categories simultaneously. The top 200 results for the query "urban sprawl" have been categorized into a two-level government hierarchy, which is used to present a categorized overview on the left. The National Park Service has been selected to filter the results. The effect on the right side is to show just the three results from the Park Service (Kules, 2006b).

Using another information visualisation alternative, Citiviz (Figure 21) displays the clusters in search results using a hyperbolic tree (Rivadeneira and Bederson, 2003) and a scatterplot (Shneiderman et al., 2000). The Technical Report Visualizer prototype (Kunz, 2003) allows users to browse a digital library by one of two user-selectable hierarchical classifications, also displayed as hyperbolic trees and coordinated with a detailed document list. Chen (Chen, 2006) implements clustered and time sliced views of a research literature to display the evolving research fronts over time. Chen's work concludes that there a number of benefits in visualising results like this but that the metrics used to build and display the visualisations are very important to the acceptance by users.

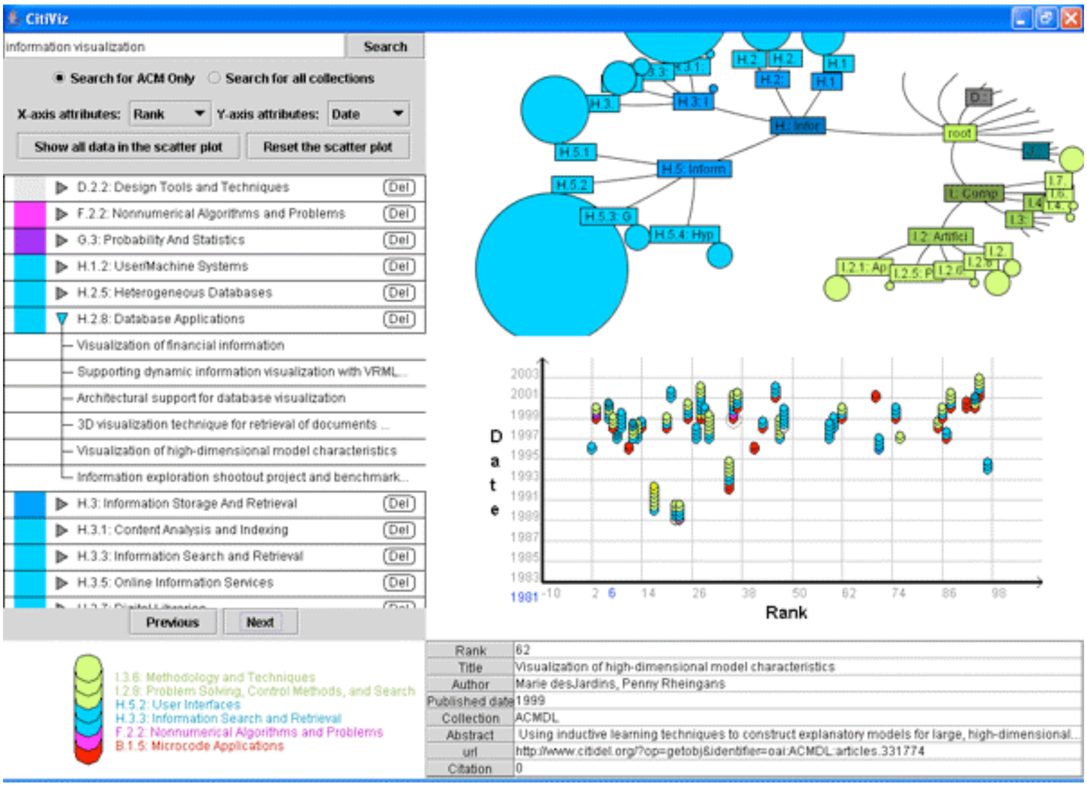

Figure 21. The CitiViz search interface visualizes search results using scatterplots, hyperbolic trees, and stacked discs. The hyperbolic tree, stacked disks, and textual list on the left are all based on the ACM Computing Classification System.

Several web search (or metasearch) engines, including Grokker (www.grokker.com), Kartoo (www.kartoo.com), and FirstStop WebSearch (www.firststopwebsearch.com) incorporate visualizations similar to both treemaps and hyperbolic trees. Kartoo Figure 23 produces a map of clusters that can be interactively explored using Flash animations. Grokker clusters documents into a hierarchy and produces an Euler diagram, a coloured circle for each top-level cluster with sub-clusters nested recursively (Figure 22). Users explore the results by "drilling down" into clusters using a 2-D zooming metaphor. It also provides several dynamic query controls for filtering results. Unlike treemaps, however, the circular form often leads to wasted space in the visualisation. Further, this early version of the Grokker interface has been found to compare poorly with textual alternatives (Eaton and Zhao, 2001). The authors found that the textual interfaces were significantly preferred. The conclusions were that web search results lack "1)… a natural spatial layout of the data; and 2)… good small representations," which makes designing effective visual representations of search results challenging. Refined visual structures making better use of space and built around meaningful classifications and colouring may ameliorate this problem, as illustrated by promising interfaces like WebTOC, which uses a familiar hierarchy of classification and a colour-coded visualisation of each site's content. Many users appreciate the visual presentation and animated transitions, so designing them to be more effective could lead to increased user acceptance.

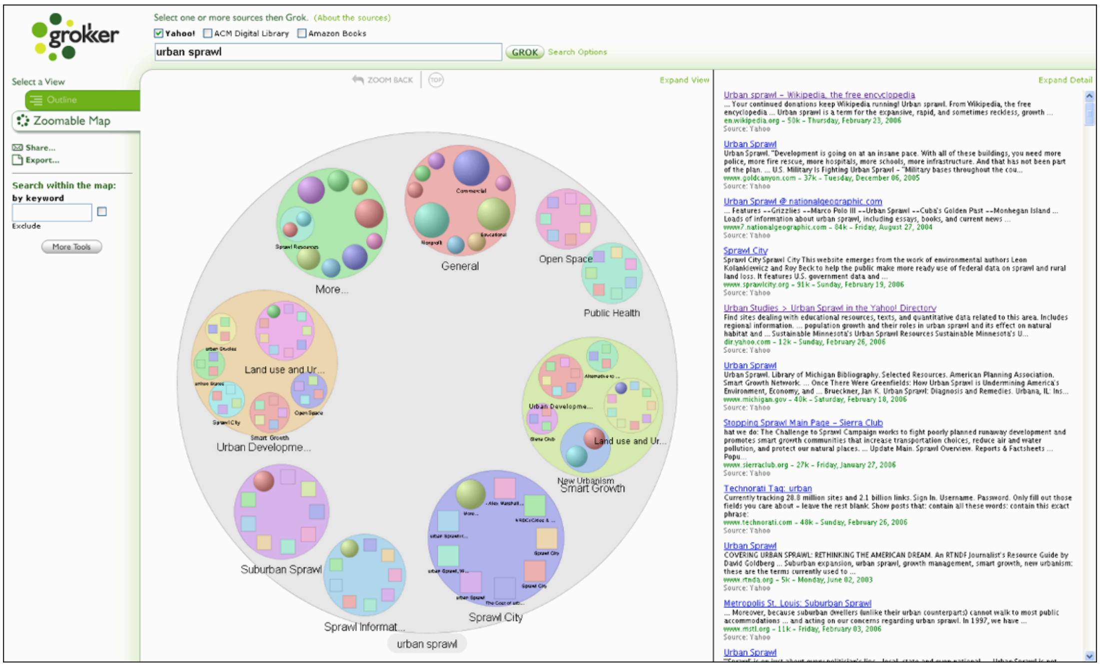

Figure 22. Grokker clusters documents into a hierarchy and produces an Euler diagram, a colored circle for each top-level cluster with sub-clusters nested recursively.

Figure 23. Kartoo generates a thematic map from the top dozen search results for a query, laying out small icons representing results onto the map.

In support of the conclusions about cluster map visualisations from Cai (Cai, 2002), early information visualisation research by Tufte (Tufte, 1990) states that graphical visualisations need to have clear and meaningful axis to be effective for users. Other approaches to 2D graphical representations have focussed on representing results in meaningful and configurable grids, where each axis is a specific metric, such as time, location, theme, size, or format, etc. The GRiDL prototype (Figure 24) displays search result overviews in a matrix using two hierarchical categories (Terveen et al., 1999). The users can easily identify interesting results by cross-referencing the two dimensions. The List and Matrix Browsers provide similar functionality (Grewal et al., 1999; McCrickard and Kehoe, 1997). The Scatterplot browser (Amento et al., 1999), also allows users to see results in a visualisation that is closer to a graph, where results are plotted on two configurable axis. Without specific regions plotted in the scatterplot visualisation, it is harder for users to view subsets of results, but easier to identify single result items. Informal evaluations of these interfaces have been promising, although no extensive studies of the techniques have been published.

Figure 24. The GRiDL prototype displays search results within the ACM Digital Library along two axes. In this screenshot, documents are organized by ACM classification and publication year. Individual dots are colored by type of document.

Another stream of research has focussed on attributes of documents that can form dimensions that leverage common knowledge. For example, GeoVIBE tightly coupled a geographic layout with an abstract layout of documents relative to defined points of interest (Robertson et al., 1998). A similar, and familiar dimension, is time, and research has also looked at visualising documents and classifications on a timeline (André et al., 2007). Finally, another clear metaphor is simply distance. Benford et al. (Benford et al., 1999) displayed results in a series of concentric circles, where the distance from the centre of the circles represents the relevance of the documents to the query.

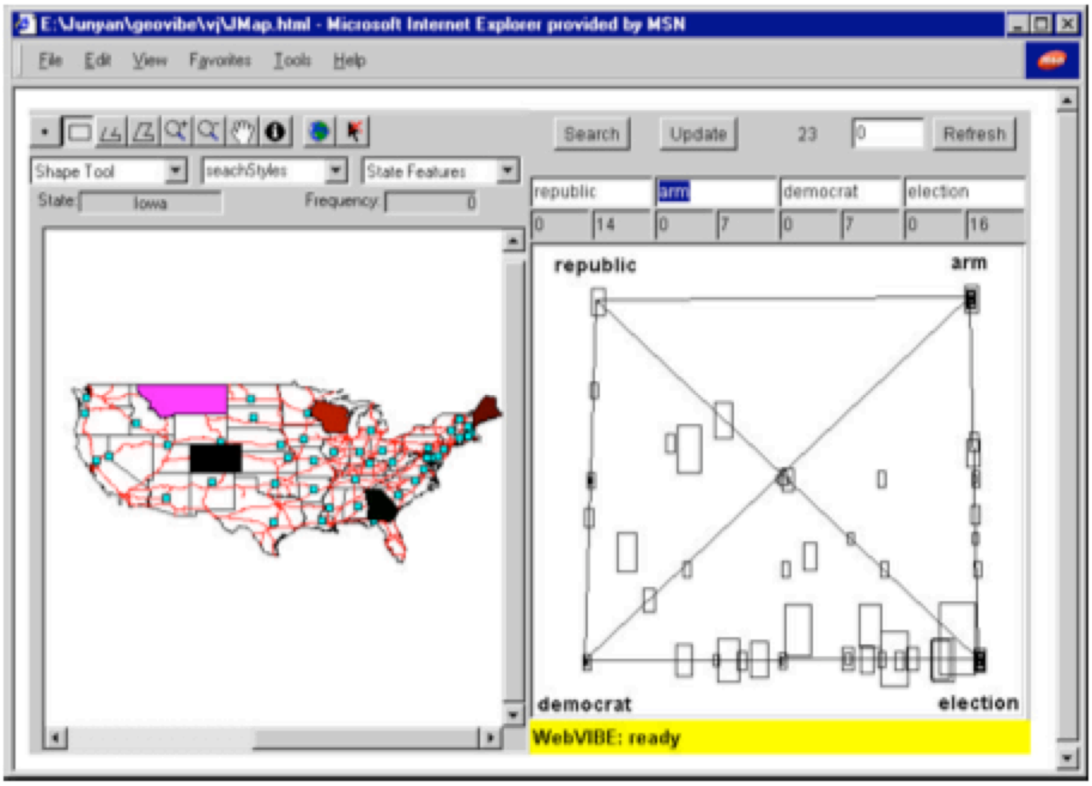

Figure 25: GeoVIBE, a search interface using a geography visualization for results.

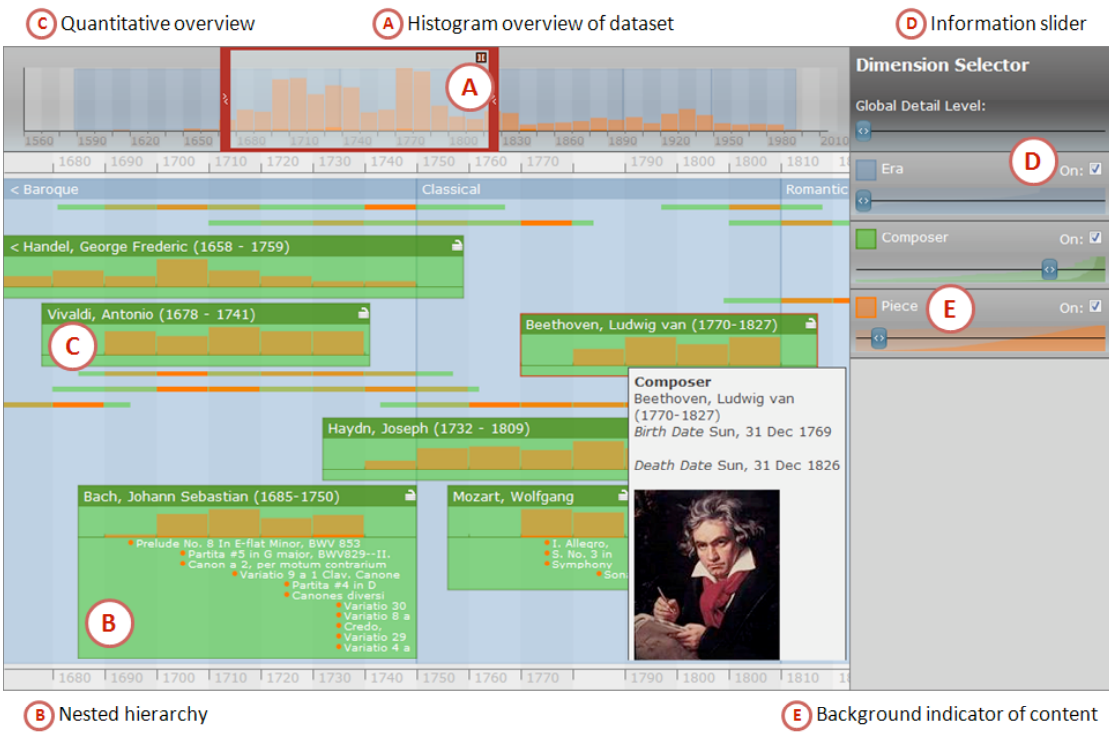

Figure 26: Continuum represents documents in nested categorizations on a timeline.

Each of the visualisations above have focussed on visualising a result set, so that users can identify areas of results that may be relevant to them. A different approach is to provide an overview visualisation to help direct people to relevant items in linear result sets. Query term similarity allows searchers to explore the contribution that each query term makes to the relevance of each result by displaying the terms and results in a two- or three-dimensional space (Börner, 2000). Similar work, by Hoeber and Yang (Hoeber and Yang, 2006), provides an overview of 100 search results, displaying a grid, with query terms as the horizontal axis and the 100 results as the vertical axis. Users are able to click on parts of the grid with high colour intensity to assist users in finding relevant documents that may not be in the top 100 results. The approach of query and result set visualisations has shown significant search benefits for users, and augments, rather than visualises, result sets.

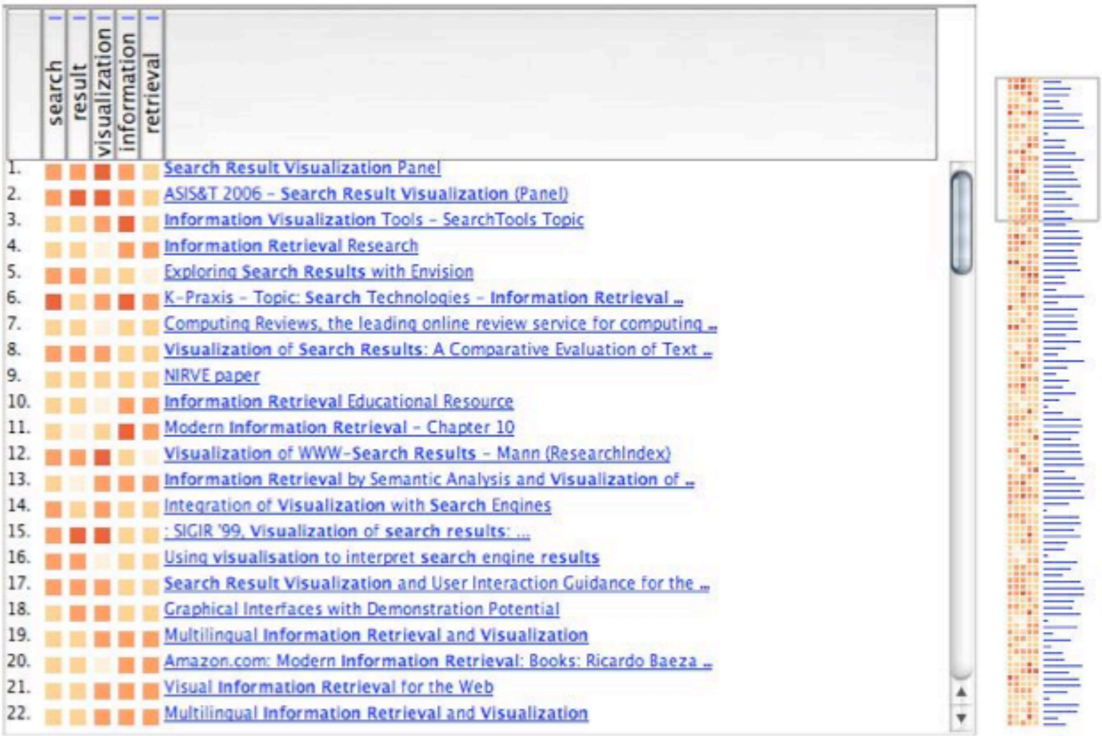

Figure 27: Hotmaps, a 2D visualization of how query terms relate to search results.

The strength of the more successful 2D visualisations has been when the use of space is clear and meaningful to users. An alternative approach, to computer generated visualisations or clusters of documents, is to allow users to arrange documents under their own parameters. The TopicShop interface permits users to meaningfully arrange sites on a canvas, where the clusters are meaningful to the individual and, when returning to a document, spatial consistency makes it easier for users to remember where they placed the result. More recent research in this area has explored 3D spaces, and is described in the section below.

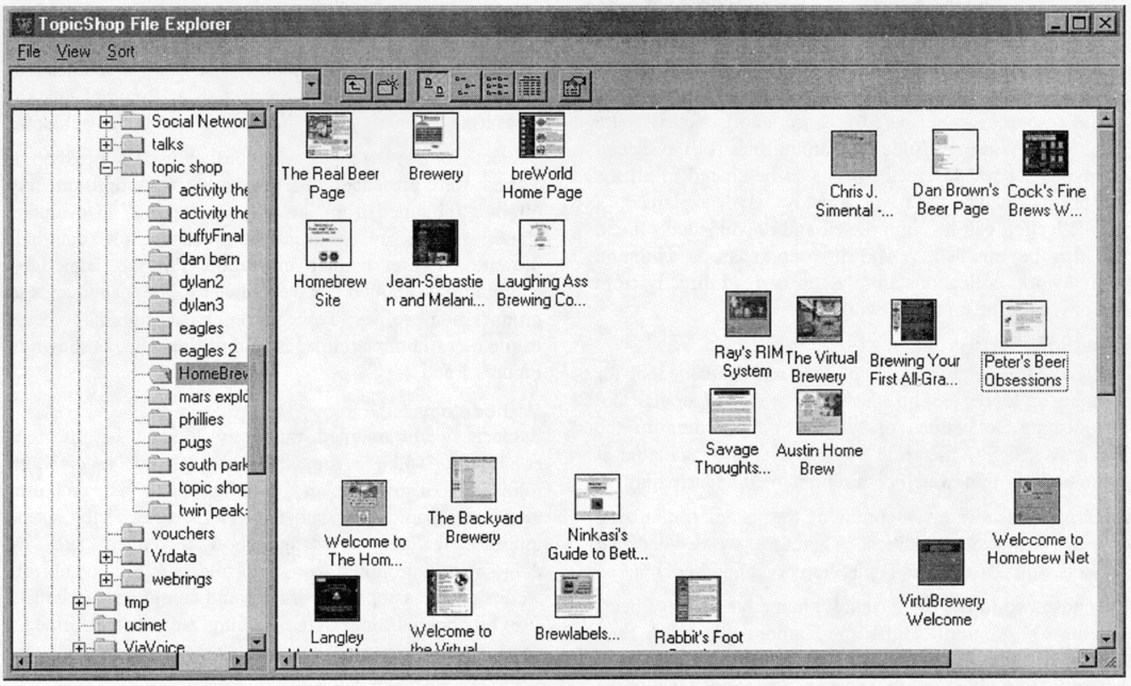

Figure 2: TopicShop Explorer, icons view. Each site is represented by a large thumbnail image and the site title. Users can organize sites by arranging them spatially, a technique especially useful in the early stages of exploration.

Figure 28. The TopicShop Explorer interfaces combines a hierarchical set of topics with a user-controlled spatial layout of sites within each topic (shown here) or a detailed list of titles and attributes . (image from PDF)

Similar to the ideas behind TopicShop, (Furnas and Rauch, 1998) presents a set of design principles for information spaces and instantiates them in the NaviQue workspace. In NaviQue, a single information surface is used to display queries, results, information structures, and ad hoc collections of documents. Users can initiate a query simply by typing onto a blank section of the workspace. Results are clustered around the query text. Kerne et al. (Kerne et al., 2006) furthered this work by considering the spatial layout of different multimedia types for constructing personal collections.

Figure 29: combinFormation automatically searches the web for multimedia relating to a query term and displays them on a 2D plane. Users can drag and drop the items to make arrange their own sub-collections.

3.2.3 3D Result Representations

When presenting research into 3D visualisations, it is first important to consider research into the benefits and weakness of 3D over 2D. Although research has shown success for some situations, such as simulations and 3D CAD/CAM design (Shneiderman, 2003), research into the visualisation of information has shown that a third dimension can inhibit users and make interfaces more confusing (Risden et al., 2000; Sutcliffe and Patel, 1996). Aspects of 3D representations, such as occlusion and cost of visualising depth, must be considered (Cugini et al., 1996). Further, research by Modjeska (Modjeska, 2000) has shown that 25% of the population struggle with 3D visualisations displayed on a 2D device, such as a computer screen. Investigation by Sebrechts et al. (Sebrechts et al., 1999) also showed that participants were significantly slower at using a 3D interface, unless they had significant computer skills. Considering these challenges, however, the research described below highlights some of the ideas that have been proposed for 3D visualisations.

The Data Mountain browser, like the TopicShop explorer described above, allows users to arrange documents on a plane, creating a 2 ½ dimension view with perspective, reduced size for distant items, and occlusion. In the Data Mountain, however, users have a 3D plane to arrange the documents (Robertson et al., 1998). Subsequent studies of the Data Mountain have shown to be unproductive (Cockburn and McKenzie, 2000, 2002). Another approach that has been expanded to 3D environments is the hyperbolic tree (Munzner, 1997).

Figure 30. Data Mountain provides an inclined plane upon which web pages are organized. (image from PDF)

As part of a discussion on applying 3D visualizations to web search results, (Benford et al., 1999) describe the VR-VIBE system. Each query is manually or automatically positioned in a 3-D space. Documents are positioned near the queries for which they are relevant. If a document is relevant to multiple queries, it is positioned between them. Each document's overall relevance is shown by the size and shade of its representative icon.

Cone Trees (Robertson et al., 1991) were designed to display a hierarchy in a 3D view, using depth to increase the amount of the tree that is visible. Users are then able to rotate the tree using a smooth animation, although this was found confusing by a number of participants in an evaluation and user studies have not shown advantages; supporting the concerns noted by Modjeska.

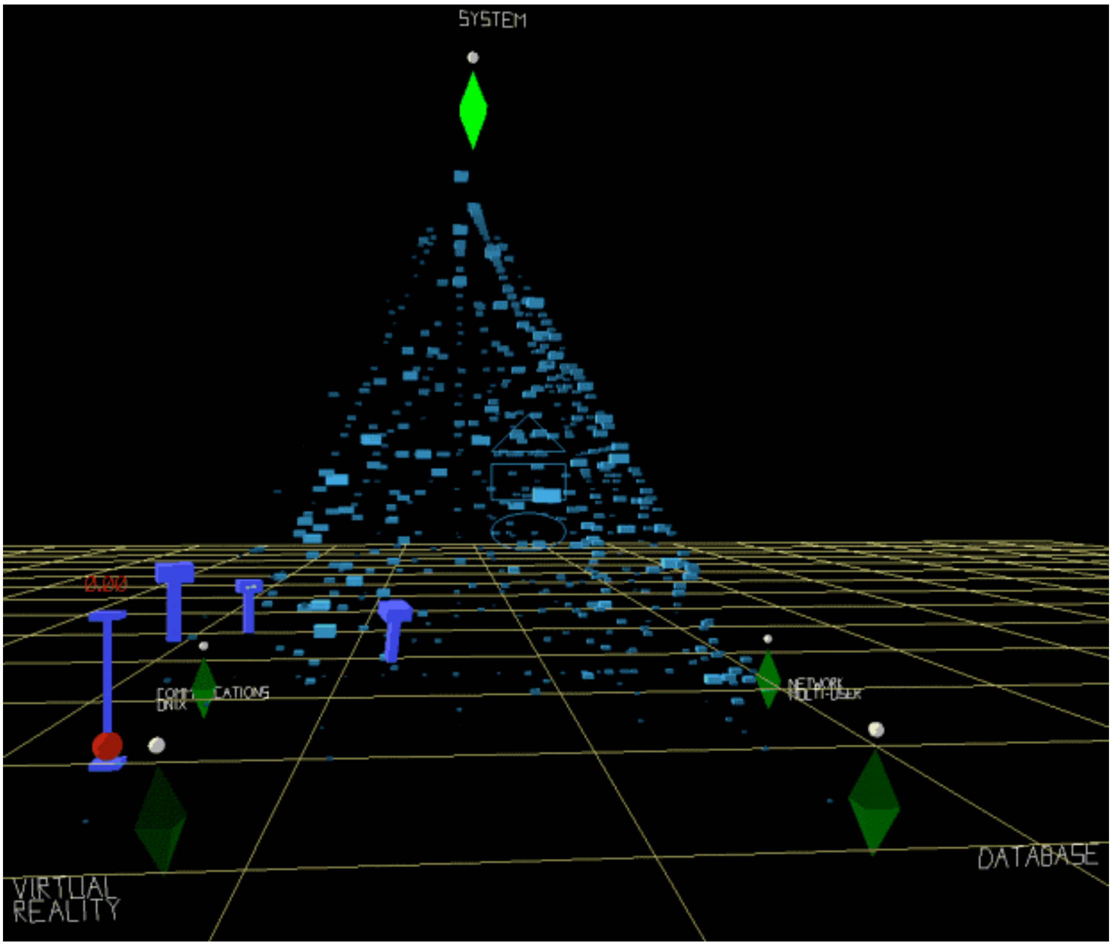

Figure 31: VR-VIBE represents search results in a 3D environment, where results are displayed in relative proximity to the keywords they match.

Finally, to fully embrace 3D environments (Börner, 2000) used Latent Semantic Analysis and clustering to organize and display a set of documents extracted from a digital library in a 3-D space. A multi-modal, virtual reality interface, called the CAVE enables users to explore the document collection; users control their movement through the CAVE with a special input device called a Wand. This requirement for a special input device, however, makes it an unrealistic option for web environments where most users will have typical mouse and keyboard input.

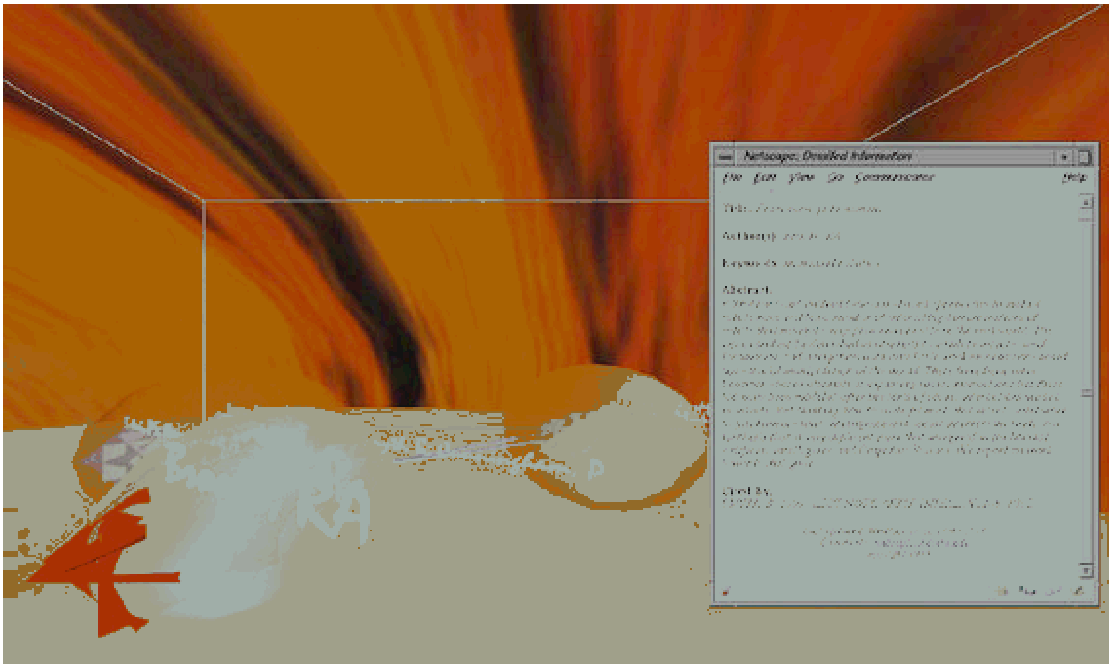

Figure 32. Search results being visualized using the CAVE explorer (image from PDF).

3.3 Additional Functions

3.3.1 Previews

One of the challenges users can face when presented with classification systems is deciding which categories to select in order to find the documents they require. This heavily depends on the label used to represent a category, but if users do not recognise a one or more terms, it can become a block for them during search (schraefel et al., 2004).

Above, we noted that the Rave Browser previews to users the affect that a selection will have on the remaining items and facets, by temporarily reducing individual bar graph representations; selecting the item will make this change permanent. This notion has been used in research into the mSpace browser (schraefel et al., 2004) to help users understand the terms presented in a facet. When a user hovers over an item in an mSpace column facet, they are presented with a multimedia preview cue, which provides an example of the documents included within the category. A user evaluation showed that this preview cue significantly supported users in finding documents. Further research by Endeca (Kotelly, 2007) is investigating the ability to automatically summarise the result items and present users with a text description that is typical of the cluster.

3.3.2 More like this

Search engines typically provide users with a link by each result to see more results that are similar according to the server's metrics. Although no specific research has been published, Endeca's search interface allows users to express a similar desire to see related documents, but allows them to express the dimension by which they want to see similarity. For example, they can choose to see more documents of a similar category, or size, or creation date.

3.3.3 Connections

Networks have been used to for knowledge discovery tasks, displaying connections between related documents and between related literatures. Stepping Stones visualizes search results for a pair of queries, using a graph to show relationships between the two sets of results (Das-Neves et al., 2005).

Similarly, Beale et al. (Beale et al., 1997) visualizes sequences of queries using a force-directed layout of node-link diagrams. Queries and their resulting documents are represented as nodes, with links between a query and its results. When a document is in the result set for multiple queries, it is linked to each query.

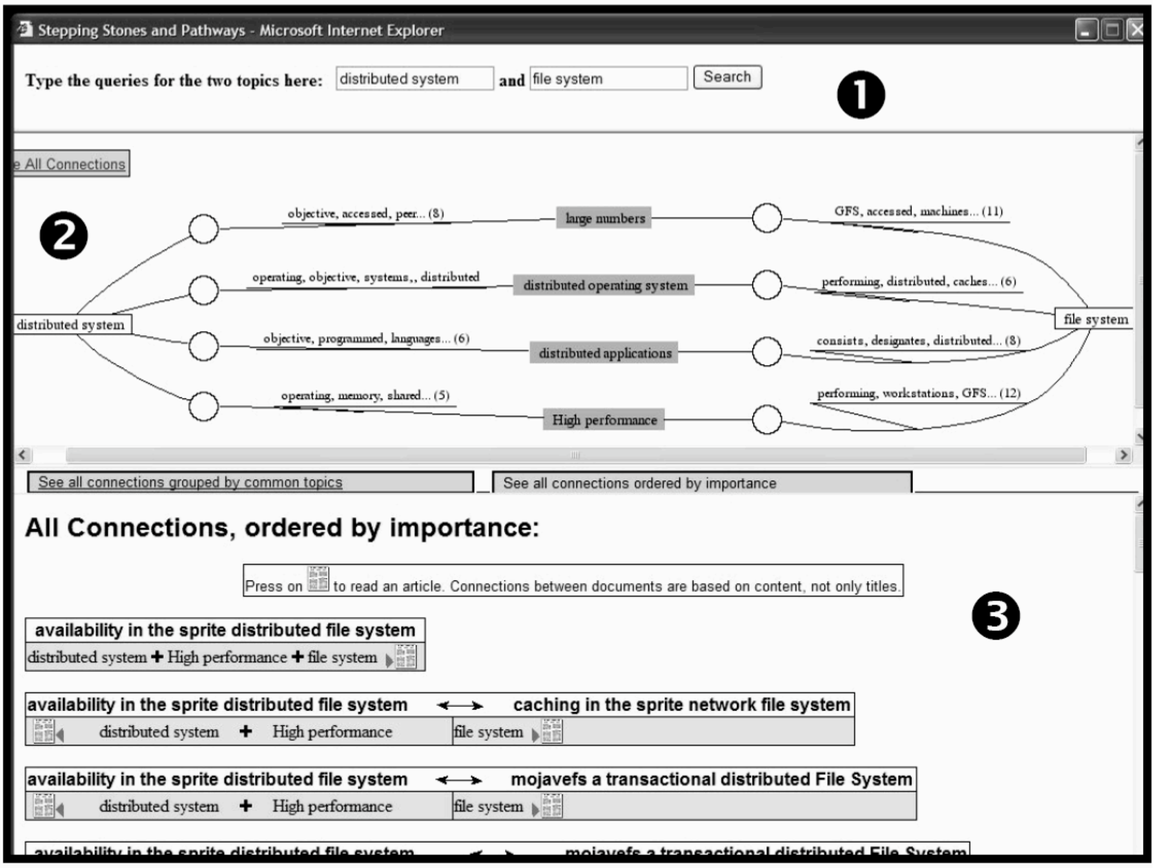

Figure 33. The Stepping Stones interface visualizes the results of a pair of queries (1) as a graph (2) of topic nodes connected to the queries. Individual documents, shown in (3), cover portions of the path between the queries. (image from PDF)

In some representations, single documents can appear in multiple collections or sub-collections. When these different groups are arranged into regions that correspond to document attributes, often lines or colours are used to show related items. Such regions can include such as a geographic map or a hierarchy of courts (Kang and Shneiderman, 2006; Shneiderman and Aris, 2006). Figure 34 shows a visualisation, where arcs that jump between the regions represent single items in different groups.

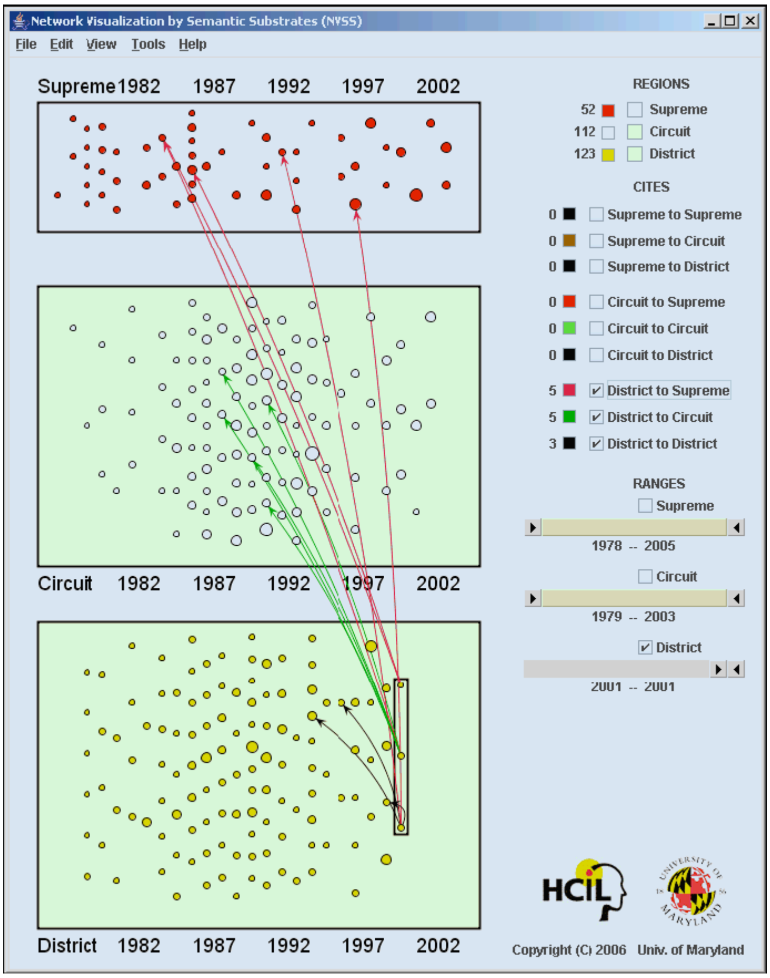

Figure 34. This visualization arranges 287 court cases into three regions (Supreme Court, Circuit Court, and District Court). Within each region, cases are arranged chronologically. Of the 2032 citations, the selection box in the District region limits the display to just the citations from the three District court cases in 2001. This shows the greater importance of older Supreme Court opinions (Shneiderman and Aris, 2006). (image from PDF)

One challenge of representing links across multiple sets or dimensions is the limited number of methods to make connections. In Figure 34, the arcs are in various colours. Such use of colours and arrows can be limited. For example, it may be hard to display how a single item in 3 different groups may be connected, or how an item relates to two parent levels of a hierarchy; such challenges become more important when hierarchical and faceted classification schemes are involved. In Continuum, relationships across multiple levels of hierarchy are shown by clustering and nesting results rather than keeping them in separate panes, see Figure 26 above.

3.3.4 Animation

Animation on the web is a debated area of research. Although lots of research has aimed at developing animation tools for the web (Faraday and Sutcliffe, 1999; Haajanen et al., 1997), other research has investigated the limitations of technologies such as Flash (Holzinger and Ebner, 2003). Research by Robertson, Cameron, Czerwinski and Robbins, however, has shown that there are advantages of using animation to support users through transitions in an interface (Robertson et al., 2002). Further, recent research by (Baudisch et al., 2006) has presented a tool called Phosphor that of highlights, with a fading period, any changes that occur in an interface. We can see examples of this sort of animated support on the web through technologies such as AJAX (Paulson, 2005). Facebook, for example, allows users to post comments with AJAX, where the page does not refresh and both colour and smooth spatial changes indicate to users that the interface is changing to accommodate their actions. Research into mSpace, which uses a lot of AJAX technology, also identified a similar emphasis on smooth transitions during discussions with participants (Wilson and schraefel, 2008a).

3.3.5 Semantic Zooming

Semantic zooming is characterised by the progressive inclusion of additional information, as opposed to simply enlarging information. In the design of NaviQue (Furnas and Rauch, 1998), users can zoom into clusters and into information structures. The semantic zooming replaces items with representative icons as users zoom out, and replaces icons with the specific items when zooming in to show more detail. A history of all ad hoc collections is automatically maintained in one corner of the display. Users can drag and drop collections into a "pocket" in another corner so that they are immediately accessible, without need to pan or zoom the workspace.

As a form of animation, effective zooming on the web is challenging and must be done to enhance a website and not affect its usability. In maintaining a typical search result list style, another option presented by research is a technique called WaveLens, which allows users to zoom in more on a single result, such that the information snippet expands to reveal more lines of text from the source document (Paek et al., 2004).

3.3.6 Alternative Inputs

Multimedia can often be hard to describe with words, and querying over such collections requires that the documents be annotated. While some research is aimed at automatically increasing the amount of annotation of multimedia (Bloehdorn et al., 2005; Jeon et al., 2003), other approaches have examined the querying medium. Query-by-example (Kato et al., 1992; Santini and Jain, 1998) is a strand of research, where users can submit audio to find related audio, images to find similar images, video to find similar video etc. For example, Shazam (www.shazam.com) is a mobile phone service that users can ring while music is playing and responds by sending an SMS message with the name and artist of a song. Similarly, Retrievr (labs.systemone.at/retrievr/) is a service that allows users to construct a simple image with paint tools, and finds images from Flickr (www.flickr.com) that match the approximate shapes and colours. Although some research has discussed practical applications for systems such as query-by-humming, to search for music (Kosugi et al., 2000), the challenge for the web is providing a means of input for users. Where Google provides a keyword box and Retrievr provides a sketch box with paint tools, it can become difficult to allow audio input, without requiring users to make recordings and upload them. Similar problems arise for video-querying for video.

4. Evaluation Methods

This section examines evaluation methods that have been used to assess systems different parts of information systems. We break this evaluation discussion down into the three levels included in the framework from Section 2.

4.1 Information retrieval evaluations

The main concerns for information retrieval research have been to assess the quality of indexing methods and the algorithms that match documents to the queries provided by users. The TREC conferences (Harman, 1997) have been arranged to evaluate document retrieval systems for domains like the Web, SPAM, video, legal documents and genomics; for collections with interesting or important characteristics such as very large collections; and for specialized tasks such as cross-language IR, filtering, finding novel documents, and question answering.

To support the TREC evaluations, predefined collections of documents, relevant queries and human relevance assessments, were produced and used as benchmark standards across any studies. This provided the opportunity to not only evaluate, but also compete for most improved retrieval times or retrieval accuracy within the community. The most commonly used measurements in TREC were precision and recall (see http://trec.nist.gov/tracks.html). Yang (Yang, 2005) and Kobayashi (Kobayashi and Takeda, 2000) provide extensive reviews of evaluations focusing on the techniques and evaluations used for information retrieval.

4.2 Information seeking task evaluations

The evaluations of information retrieval have very different aims to the evaluations of information seeking. Information retrieval evaluations have focussed on metrics of the server side, but have consequently been criticized for a narrow conceptualization of the information need, relevance and interaction (Borlund, 2003b). Information seeking research, therefore, focuses on evaluating systems for how they meet the needs of users.

Recent work by Wilson and schraefel (Wilson and schraefel, 2007; Wilson et al., to come) has proposed an evaluation framework for information systems to systematically assess their support for the needs of a range of known types of users. The framework first measures the functionality of the search system by the way it supports known tactics and moves employed with information (Bates, 1979b, a). The framework then uses a novel mapping to summarise the measured support to the of different types of users (Belkin et al., 1993), whose conditions vary on dimensions such as previous knowledge and intended use. This evaluation framework was later refined and validated to show that its results could accurately predict the results of user studies (Wilson and schraefel, 2008b). With the confidence provided by the validation, the framework can be used to identify weaknesses in the design of a system, or new function alone, so that it can be resolved before user studies are carried out. Further, it can be used to inform the design of user studies, so that they accurately test the desire features.

One of the exploratory studies used to validate the framework produced by Wilson et al. was an information seeking evaluation of faceted browsers (Capra et al., 2007); faceted browsing is described in Section 3.1.2. The study used three types of tasks to evaluate the system, and performed both a within participants study and a between participants study to get qualitative and quantitative results, respectively. The first type of task was a simple lookup task, which could be answered by using only one facet of the annotation. The second task type was a complex lookup, which involved multiple facets. The final task was exploratory, where users were asked to learn and produced a summarised report about a certain topic. This third type of task is a good example of something that has not been included in Information Retrieval research, which, as mentioned above, has focussed on matching documents to queries. Instead, by asking users to carry out learning tasks, we can assess the system for other types of information seeking activities, such as comparison, synthesis, and summarisation.

As part of their discussion of higher-level problems encountered by users of information visualization systems, Amar and Stasko (Amar and Stasko, 2004) discusses tasks that may be part of or require exploratory search.

Another interesting contribution to information seeking evaluations is on the discussion of time as a measurement by Capra *et al*. Although their tasks were timed, they suggest that time may not be a useful metric for exploratory tasks, as extended use of a system could mean that users have discovered increasing amounts of relevant information. In contrast to the information retrieval view that finding the answer quicker is more important, finishing an exploratory task early may indicate that a search system does not provide effective support for browsing. A separate measure would have to be in place to know that users had learnt all there is to know about a topic, and then speed can be used as a measurement.

The suitability of relevance in exploratory search conditions may also be in question for some information seeking evaluations. In systems that use faceted classifications, for example, each document with a particular annotation has an equal weighting and thus every document suggested as a result of selecting a particular part of the classification will be equally relevant. Instead, Spink *et al*. (Spink, 2002; Spink and Wilson, 1999), have been designing a metric that tries to measure the progress of users in achieving their goal. Although designed for feedback to users, the rate of progress for similar tasks on different systems could be used to assess their support.

Koshman (Koshman, 2005) evaluated the VIBE (Visual Information Browsing Environment) prototype system, which graphically represents search results as geometric icons within one screen display. As part of understanding different approaches to information seeking, the researchers sought to understand how expert and novice use and performance differed through the use of a quasi-experimental within-participants design. Borland discusses more about information seeking evaluation (Borlund, 2003b).

Having discussed search tasks and measurements that may be unique to or important for information seeking, carefully controlled user studies can still be performed to evaluate systems in terms of information seeking. Käki and Aula (Käki, 2005) provide a good example of a study that employs a within-subjects design, balanced task sets, time limitations, pre-formulated queries, cached result pages and limiting access to result documents.

Finally, the study approach that investigates user interaction with software over a long period of time, such as longitudinal studies or studies that are repeated periodically with the same participants, provides a unique type of insight into realistic human behaviour. One of the main arguments against the short task-oriented studies is the lack of realism. Further, the results of such user studies are often based on the participants' first or early responses to new designs, compared to their familiarity with existing software. By studying interaction over time, we can begin to evaluate changes in search tactics and subjective views as users adopt and adapt to new interfaces (Kules, 2006a; Vakkari, 2000; Wilson and schraefel, 2008a).

4.3 Work task evaluations

Evaluating search systems and the work task level requires measuring the success of work task style problems. For example, Allen (Allen, 1998) investigated the interaction of spatial abilities with 2-D data representations. The work-level task given to participants was to read an article and then use the system to find a few good articles. At this level, many different information-seeking activities can be used, but ultimately the system was assessed on its support for achieving the overall goal. Other research has used similar work-task oriented measurements, including producing creating instructional material (Kabel et al., 2004) and writing newspaper articles (Kules and Shneiderman, to appear), in order to show that users explored more or learnt more with different information retrieval systems.

One of the challenges of evaluation, especially at the information seeking and work task level, is that human participants interpret the tasks based on their own experiences and knowledge. There is a tension between the need to make results reliable and replicable and the need to make the task realistic for the participants. (Borlund, 2003a) advocated addressing this by incorporating participant-provided information needs into the experimental session along with a researcher-provided need. If the results for both tasks are consistent, the researcher can conclude that the researcher-provided task is realistic and thus reap the benefit of a realist but tightly controlled task. Other research has used the previous actions of users to inform the realism of search tasks in system evaluations (Elbassuoni et al., 2007; Elsweiler and Ruthven, 2007)

Aside from the individuality of work task understanding, work tasks vary dramatically depending on the domain. Consequently, some research has addressed domain specific work task evaluations to show that an evaluation approach can be applied accurately in multiple scenarios (Gong and Kieras, 1994; Huvila and Widén-Wulff, 2006).

Jävelin and Ingwersen (Järvelin and Ingwersen, 2004) argue that the main challenge ahead for evaluating systems at the work task level, is that research needs to explicitly integrate context in terms of the high level work task, the specific information seeking task and the systems context. In particular, they note that the Web "is not a single coherent unit but appears quite different for different actors, tasks, and domains."

5. Conclusions

In this report, we have covered research into visualising search results on the web in three main sections. First, the theory of search, information retrieval, and information seeking has been covered to provide sources that may help to inform new designs of information systems. By better understanding what drives the use of a system, we can identify and respond to the requirements of users. Second, a survey of existing information systems and prototypes is provided to give examples of systems and functions that have been shown to successfully enhance the search experience for users and examples of research that has found the limitations of design ideas. Finally, we have presented a survey of evaluation techniques that have been applied to understand how systems and designs support users in search.

6. References

AGUILAR, F. J. (1988) *General Managers in Action,* New York, NY, Oxford University Press.

ALLEN, B., Year. Information space representation in interactive systems: relationship to spatial abilities. *Proceedings of the third ACM conference on Digital libraries*. Pittsburgh, Pennsylvania, United States, June 23-26, p. 1-10. ACM Press.

ALLEN, R., Year. Two digital library Interfaces that exploit hierarchical structure. *Proceedings of Electronic Publishing and the Information Superhighway*. Boston, MA, USA, May 30-June 2, p. 134–141.

ALLSOP, J. (2007) *Microformats: Empowering Your Markup for Web 2.0*, friends of ED.

AMAR, R. & STASKO, J., Year. A Knowledge Task-Based Framework for Design and Evaluation of Information Visualizations. *Proceedings of the IEEE Symposium on Information Visualization*. Austin, Texas, USA, October 10-12, p. 143-150. IEEE Computer Society.

AMENTO, B., HILL, W., TERVEEN, L., HIX, D. & JU, P. (1999) An empirical evaluation of user interfaces for topic management of Web sites. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems.* Pittsburgh, Pennsylvania, United States, ACM Press.

ANDRÉ, P., WILSON, M. L., RUSSELL, A., SMITH, D. A. & OWENS, A., Year. Continuum: designing timelines for hierarchies, relationships and scale. *Proceedings of the 20th annual ACM symposium on User interface software and technology*. Newport, Rhode Island, USA, October 6-10, p. 101-110. ACM Press.

AU, P., CAREY, M., SEWRAZ, S., GUO, Y. & RÜGER, S. (2000) New paradigms in information visualization. *Proceedings of the 23rd annual international ACM SIGIR conference on Research and development in information retrieval.* Athens, Greece, ACM Press.

BATES, M. J., 1979a Idea tactics. *Journal of the American Society for Information Science,* 30 (5)**,** p. 280-289.

BATES, M. J., 1979b Information search tactics. *Journal of the American Society for Information Science,* 30 (4)**,** p. 205-214.

BATES, M. J., 1989 The Design of Browsing and Berrypicking Techniques for the Online Search Interface. *Online Review,* 13 (5)**,** p. 407-424.

BATES, M. J., 1990 Where should the person stop and the information search interface start. *Information Processing and Management,* 26 (5)**,** p. 575-591.
BAUDISCH, P., TAN, D., COLLOMB, M., ROBBINS, D., HINCKLEY, K., AGRAWALA, M., ZHAO, S. & RAMOS, G., Year. Phosphor: explaining transitions in the user interface using afterglow effects. *Proceedings of the 19th annual ACM symposium on User Interface Software and Technology*. Montreux, Switzerland, October 15-18, p. 169-178. ACM Press.
BEALE, R., MCNAB, R. J. & WITTEN, I. H., Year. Visualising sequences of queries: a new tool for information retrieval. *Proceedings of the IEEE Conference on Information Visualisation*. Phoenix, AZ, USA, October 18-25, p. 57-63. IEEE Computer Society.
BECKS, A., SEELING, C. & MINKENBERG, R., Year. Benefits of document maps for text access in knowledge management: A comparative study. *Proceedings of the ACM Symposium on Applied Computing*. Madrid, Spain, March 10-14, p. 621-626. ACM Press.
BEITZEL, S. M., JENSEN, E. C., CHOWDHURY, A., GROSSMAN, D. & FRIEDER, O., Year. Hourly analysis of a very large topically categorized web query log. *Proceedings of the 27th annual international ACM SIGIR conference on Research and Development in Information Retrieval*. Sheffield, UK, July 25-29, p. 321-328. ACM Press.
BELKIN, N. J., 1980 Anomalous states of knowledge as a basis for information retrieval. *Canadian Journal of Information Science,* 5 133-143.
BELKIN, N. J., COOL, C., STEIN, A. & THIEL, U., 1995 Cases, scripts, and information-seeking strategies: on the design of interactive information retrieval systems. *Expert Systems with Applications,* 9 (3)**,** p. 379-395.
BELKIN, N. J., MARCHETTI, P. G. & COOL, C., 1993 Braque: design of an interface to support user interaction in information retrieval. *Information Processing & Management,* 29 (3)**,** p. 325-344.
BENFORD, S., TAYLOR, I., BRAILSFORD, D., KOLEVA, B., CRAVEN, M., FRASER, M., REYNARD, G. & GREENHALGH, C., 1999 Three dimensional visualization of the World Wide Web. *ACM Computing Surveys,* 31 (4)**,** p. 25.
BHAVNANI, S. K. & BATES, M. J., Year. Separating the knowledge layers: Cognitive analysis of search knowledge through hierarchical goal decompositions. *Proceedings of the American Society for Information Science and Technology Annual Meeting*. Medford, NJ, November 18, p. 204-213. Information Today.
BLOEHDORN, S., PETRIDIS, K., SAATHOFF, C., SIMOU, N., TZOUVARAS, V., AVRITHIS, Y., HANDSCHUH, S., KOMPATSIARIS, I., STAAB, S. & STRINTZIS, M. G., Year. Semantic Annotation of Images and Videos for Multimedia Analysis. *Proceedings of the 2nd European Semantic Web Conference*. Heraklion, Crete, Greece, May 29-June 2, p. 592-607. Springer.
BORLUND, P., 2003a The IIR evaluation model: A framework for evaluation of interactive information retrieval systems. *Information Research,* 8 (3)**,** p. paper no. 152.
BORLUND, P., 2003b The IIR evaluation model: a framework for evaluation of interactive information retrieval systems. *Information Research,* 8 (3)**,** p. Paper 152.
BÖRNER, K., Year. Visible Threads: A Smart VR Interface to Digital Libraries. *Proceedings of IST/SPIE's 12th Annual International Symposium on Visual*

Data Exploration and Analysis. San Jose, CA, USA, January 22-28, p. 228-237. SPIE.
BYSTRÖM, K. & HANSEN, P. (2002) Work tasks as units for analysis in information seeking and retrieval studies. IN BRUCE, H., FIDEL, R., INGWERSEN, P. & VAKKARI, P. (Eds.) *Emerging Frameworks and Methods.* Greenwood Village, CO, Libraries Unlimited.
BYSTRÖM, K. & HANSEN, P., 2005 Conceptual framework for tasks in information studies: Book Reviews. *Journal American Society Information Science and Technology,* 56 (10)**,** p. 1050-1061.
CAI, G., Year. GeoVIBE: A Visual Interface for Geographic Digital Libraries. *Proceedings of the First Visual Interfaces to Digital Libraries Workshop*. Roanoke, VA, USA, June 28, p. 171-187. Springer-Verlag.
CAPRA, R., MARCHIONINI, G., OH, J. S., STUTZMAN, F. & ZHANG, Y., Year. Effects of structure and interaction style on distinct search tasks. *Proceedings of the ACM/IEEE-CS Joint Conference on Digital libraries*. Vancouver, British Columbia, Canada, June 18-23, p. 442-451. ACM Press.
CHEN, C., 2006 CiteSpace II: Detecting and visualizing emerging trends and transient patterns in scientific literature. *Journal of the American Society for Information Science and Technology,* 57 (3)**,** p. 359-377.
CHEN, H. & DUMAIS, S., Year. Bringing order to the Web: automatically categorizing search results. *Proceedings of the SIGCHI conference on Human factors in computing systems*. The Hague, The Netherlands, April 1-6, p. 145 - 152. ACM Press.
CHEN, H., HOUSTON, A. L., SEWELL, R. R. & SCHATZ, B. R., 1998 Internet browsing and searching: User evaluations of category map and concept space techniques. *Journal of the American Society for Information Science,* 49 (7)**,** p. 582-608.
CHOO, C. W., DETLOR, B. & TURNBULL, D. (2000) *Web Work: Information Seeking and Knowledge Work on the World Wide Web,* Dordrecht, The Netherlands, Kluwer Academic Publishers.
COCKBURN, A. & MCKENZIE, B., Year. An Evaluation of Cone Trees. *Proceedings of the British Computer Society Conference on Human-Computer Interaction*. Sheffield, UK, September 5-8, p. Springer-Verlag.
COCKBURN, A. & MCKENZIE, B., Year. Evaluating the effectiveness of spatial memory in 2D and 3D physical and virtual environments. *Proceedings of the SIGCHI conference on Human factors in computing systems: Changing our world, changing ourselves*. Minneapolis, Minnisota, USA, April 20-25, p. 203-210.
CUGINI, J., PIATKO, C. & LASKOWSKI, S., Year. Interactive 3D Visualization for Document Retrieval. *Proceedings of the Workshop on New Paradigms in Information Visualization and Manipulation, ACM Conference on Information and Knowledge Management*. Rockville, Maryland, USA, November 12-16, p.
CUTRELL, E. & DUMAIS, S. T., 2006 Exploring personal information. *Communications of the ACM,* 49 (4)**,** p. 50-51.
DAS-NEVES, F., FOX, E. A. & YU, X., Year. Connecting topics in document collections with stepping stones and pathways. *Proceedings of the 14th ACM International Conference on Information and Knowledge Management*. Bremen, Germany, October 31 - November 5, p. 91-98. ACM Press.

DERVIN, B. & NILAN, M. (1986) Information needs and uses. IN WILLIAMS, M. (Ed.) *Annual Review of Information Science and Technology.* White Plains, New York, Knowledge Industries.
DONN, K., PLAISANT, C. & SHNEIDERMAN, B., Year. Query previews in networked information systems. *Proceedings of the Third Forum on Research and Technology Advances in Digital Libraries*. Washington, DC, USA, May 13-15, p. 120-129. IEEE Computer Society.
DRORI, O. & ALON, N., 2003 Using documents classification for displaying search results list. *Journal of Information Science,* 29 (2)**,** p. 97-106.
DUMAIS, S., CUTRELL, E. & CHEN, H. (2001) Optimizing search by showing results in context. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, Seattle, WA.* New York, ACM Press.
EATON, C. & ZHAO, H. (2001) Visualizing Web Search Results.
EFRON, M., ELSAS, J., MARCHIONINI, G. & ZHANG, J., Year. Machine learning for information architecture in a large governmental website. *Proceedings of the 4th ACM/IEEE-CS Joint Conference on Digital libraries*. Tuscon, AZ, USA, June 7-11, p. 151-159. ACM Press.
EGAN, D. E., REMDE, J. R., GOMEZ, L. M., LANDAUER, T. K., EBERHARDT, J. & LOCHBAUM, C. C., 1989 Formative design evaluation of SuperBook. *ACM Transactions on Information Systems,* 7 (1)**,** p. 30-57.
ELBASSUONI, S., LUXENBURGER, J. & WEIKUM, G., Year. Adaptive Personalization of Web Search. *Workshop on Web Information Seeking and Interaction at SIGIR*. Amsterdam, Netherlands, July 27, p.
ELLIS, D., 1989 A behavioral model for information retrieval system design. *Journal of Information Science,* 15 (4-5)**,** p. 237-247.
ELSWEILER, D. & RUTHVEN, I., Year. Towards task-based personal information management evaluations. *Proceedings of the 30th annual international ACM SIGIR conference on Research and Development in Information Retrieval*. Amsterdam, The Netherlands, July 23-27, p. 23-30.
FARADAY, P. & SUTCLIFFE, A., Year. Authoring animated Web pages using “contact points”. *Proceedings of the SIGCHI conference on Human factors in computing systems: the CHI is the limit.* Pittsburgh, PA, USA, May 15-20, p. 458-465. ACM Press.
FIDEL, R., 1985 Moves in online searching. *Online Review,* 9 (1)**,** p. 61-74.
FURNAS, G. W. & RAUCH, S. J., Year. Considerations for information environments and the NaviQue workspace. *Proceedings of the Third ACM Conference on Digital libraries*. Pittsburgh, PA, USA, June 23-26, p. 79-88. ACM Press.
GARCIA, E. & SICILIA, M.-Á., 2003 User interface tactics in ontology-based information seeking. *Psychology Journal,* 1 (3)**,** p. 242-255.
GINSBURG, M., 2004 Visualizing digital libraries with open standards. *Communications of the Association for Information Systems,* 13 336-356.
GODWIN-JONES, R., 2006 EMERGING TECHNOLOGIES Tag Clouds in the Blogosphere: Electronic Literacy and social Networking. *Language Learning & Technology,* 10 (2)**,** p. 8-15.
GOLOVCHINSKY, G. (1997) Queries? Links? Is there a difference? *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, Atlanta, GA.* New York, ACM Press.
GONG, R. & KIERAS, D., Year. A validation of the GOMS model methodology in the development of a specialized, commercial software application.

Proceedings of the SIGCHI Conference on Human Factors in Computing Systems. Boston, Massachusetts, USA, April 24-28, p. 351-357. ACM Press.
GREENE, S., MARCHIONINI, G., PLAISANT, C. & SHNEIDERMAN, B., 2000 Previews and overviews in digital libraries: Designing surrogates to support visual information-seeking. *Journal of the American Society for Information Science,* 51 (3)**,** p. 380-393.
GREWAL, R., JACKSON, M., WALLIS, J. & BURDEN, P., Year. Using visualisation to interpret search engine results. *The Active Web 1999*. Staffordshire, UK, January 20, 1999, p.
GULLI, A. & SIGNORINI, A., Year. The indexable web is more than 11.5 billion pages. *Proceedings of the 14th International Conference on the World Wide Web*. Chiba, Japan, May 10-14, p. 902-903. ACM Press.
HAAJANEN, J., PESONIUS, M., SUTINEN, E., TARHIO, J., TERASVIRTA, T. & VANNINEN, P., Year. Animation of user algorithms on the Web. *Proceedings of IEEE Symposium on Visual Languages*. Isle of Capri, Italy, September 23-26, p. 356-363. IEEE Computer Society.
HARMAN, D. K. (1997) The TREC conferences. *Readings on information retrieval.* San Francisco, CA, USA, Morgan Kaufmann Publishers Inc.
HASSAN-MONTERO, Y. & HERRERO-SOLANA, V., Year. Improving Tag-Clouds as Visual Information Retrieval Interfaces. *Proceedings of International Conference on Multidisciplinary Information Sciences and Technologies*. Mérida, Spain, October 25-28, p. Wiley Periodicals, Inc.
HEARST, M. (1999a) The use of categories and clusters for organizing retrieval results. IN STRZALKOWSKI, T. (Ed.) *Natural Language Information Retrieval.* Boston, Kluwer Academic Publishers.
HEARST, M. (1999b) User interfaces and visualization. IN BAEZA-YATES, R. & RIBEIRO-NETO, B. (Eds.) *Modern Information Retrieval.* Reading, MA, Addison-Wesley.
HEARST, M., ELLIOT, A., ENGLISH, J., SINHA, R., SWEARINGEN, K. & YEE, P., 2002 Finding the flow in web site search. *Communications of the ACM,* 45 (9)**,** p. 42-49.
HEARST, M. A., 2006 Clustering versus faceted categories for information exploration. *Communications of the ACM,* 49 (4)**,** p. 59-61.
HOEBER, O. & YANG, X. D., Year. The visual exploration of web search results using HotMap. *Proceedings of the International Conference on Information Visualization*. London, UK, July 5-7, p. 157-165. IEEE Computer Society.
HOLZINGER, A. & EBNER, M., Year. Interaction and Usability of Simulations & Animations: A case study of the Flash Technology. *Proceedings of IFIP TC13 International Conference on Human Computer Interaction*. Zürich, Switzerland, September 1-5, p. 777-780. IOS Press.
HOTHO, A., JASCHKE, R., SCHMITZ, C. & STUMME, G., 2006 Information retrieval in folksonomies: Search and ranking. *The Semantic Web: Research and Applications,* 4011 411–426.
HUVILA, I. & WIDÉN-WULFF, G., Year. Perspectives to the classification of information interactions: the Cool and Belkin faceted classification scheme under scrutiny. *Proceedings of the 1st international conference on Information interaction in context*. Copenhagen, Denmark, October 18-20, p. 144-152. ACM Press.

HUYNH, D. F., MILLER, R. & KARGER, D., Year. Exhibit: Lightweight Structured Data Publishing. *Proceedings of the World Wide Web Conference*. Banff, Alberta, Canada, May 8-12, p. 737-746. ACM Press.

JANSEN, B. & SPINK, A., Year. An analysis of web information seeking and use: documents retrieved versus documents viewed. *Proceedings of the 4th international conference on Internet computing*. Las Vagas, NV, USA, June 23-26, p. 65-69.

JÄRVELIN, K. & INGWERSEN, P., 2004 Information seeking research needs extension towards tasks and technology. *Information Research,* 10 (1)**,** p. paper 212.

JEON, J., LAVRENKO, V. & MANMATHA, R., Year. Automatic image annotation and retrieval using cross-media relevance models. *Proceedings of the 26th annual international ACM SIGIR conference on Research and development in informaion retrieval*. Toronto, Canada, July 28 - August 1, p. 119-126. ACM Press.

JOACHIMS, T. (1997) *Text Categorization with Support Vector Machines: Learning with Many Relevant Features*, Springer.

KABEL, S., DE HOOG, R., WIELINGA, B. J. & ANJEWIERDEN, A., 2004 The added value of task and ontology-based markup for information retrieval. *Journal of the American Society for Information Science and Technology,* 55 (4)**,** p. 348-362.

KÄKI, M., Year. Findex: search result categories help users when document ranking fails. *Proceeding of the SIGCHI Conference on Human Factors in Computing Systems*. Portland, OR, USA, April 2-7, p. 131-140. ACM Press.

KANG, H. & SHNEIDERMAN, B., 2006 Exploring personal media: A spatial interface supporting user-defined semantic regions *Journal of Visual Languages & Computing,* 17 (3)**,** p. 254-283.

KARI, J., 2006 Evolutionary information seeking: a case study of personal development and Internet searching. *First Monday,* 11 (1)**,** p.

KATIFORI, A., HALATSIS, C., LEPOURAS, G., VASSILAKIS, C. & GIANNOPOULOU, E., 2007 Ontology visualization methods - a survey. *ACM Computing Surveys,* 39 (4)**,** p. Article No. 10.

KATO, T., KURITA, T., OTSU, N. & HIRATA, K., Year. A sketch retrieval method for full color image database-query byvisual example. *Proceedings of 11th IAPR International Conference on Pattern Recognition*. The Hague, The Netherlands, August 30 - September 2, p. 530-533. IEEE Computer Society.

KERNE, A., KOH, E., DWORACZYK, B., MISTROT, J. M., CHOI, H., SMITH, M. S., GRAEBER, R., CARUSO, D., WEBB, A., HILL, R. & ALBEA, J., Year. combinFormation: a mixed-initiative system for representing collections as compositions of image and text surrogates. *Proceedings of the 6th ACM/IEEE-CS joint conference on Digital libraries*. Chapel Hill, NC, USA, June 11-15, p. 11-20. ACM Press.

KOBAYASHI, M. & TAKEDA, K., 2000 Information Retrieval on the Web. *ACM Computing Surveys,* 32 (2)**,** p. 144-173.

KOHONEN, T. (2001) *Self-Organizing Maps*, Springer.

KOSHMAN, S., 2005 Testing user interaction with a prototype visualization-based information retrieval system. *Journal of the American Society for Information Science and Technology,* 56 (8)**,** p. 824-833.

KOSUGI, N., NISHIHARA, Y., SAKATA, T., YAMAMURO, M. & KUSHIMA, K., Year. A practical query-by-humming system for a large music database.

Proceedings of the eighth ACM international conference on Multimedia. New York, NY, USA, July 30 - August 2, p. 333-342. ACM Press.
KOTELLY, B., Year. Resonance: Introducing the concept of penalty-free deep look ahead with dynamic summarization of arbitrary results sets. *Workshop on Human-Computer Interaction and Information Retrieval*. MIT CSAIL, Cambridge, Mssachusetts, USA, October 23, p.
KUHLTHAU, C. C., 1991 Inside the search process: Information seeking from the user's perspective. *Journal of the American Society for Information Science,* 42 (5)**,** p. 361-371.
KULES, B., Year. Methods for evaluating changes in search tactics induced by exploratory search systems. *ACM SIGIR 2006 Workshop on Evaluating Exploratory Search Systems*. Seattle, WA, August 11, p.
KULES, B. (2006b) Supporting Exploratory Web Search with Meaningful and Stable Categorized Overviews. University of Maryland, College Park.
KULES, B., KUSTANOWITZ, J. & SHNEIDERMAN, B., Year. Categorizing web search results into meaningful and stable categories using Fast-Feature techniques. *Proceedings of the Sixth ACM/IEEE-CS Joint Conference on Digital Libraries*. Chapel Hill, NC, June 11-15, p. 210-219. ACM Press.
KULES, B. & SHNEIDERMAN, B., to appear Users can change their web search tactics: Design guidelines for categorized overviews. *Information Processing & Management*.
KUNZ, C., Year. SERGIO - An Interface for context driven knowledge retrieval. *Proceedings of eChallenges*. Bologna, Italy, October 22-24, p.
LAMPING, J. & RAO, R., 1996 The Hyperbolic Browser: A focus + context technique for visualizing large hierarchies. *Journal of Visual Languages and Computing,* 7 (1)**,** p. 33-55.
LUNZER, A. & HORNBÆK, K., Year. Side-By-Side Display and Control of Multiple Scenarios: Subjunctive Interfaces for Exploring Multi-Attribute Data. *Proceedings of Australian Special Interest Group Conference on Computer-Human Interaction*. Brisbane, Australia, November 26-28, p.
MARCHIONINI, G. (1995) *Information Seeking in Electronic Environments*, Cambridge University Press.
MARCHIONINI, G. & SHNEIDERMAN, B., 1988 Finding facts vs. browsing knowledge in hypertext systems. *Computer,* 21 (1)**,** p. 70-80.
MCCRICKARD, D. S. & KEHOE, C. M., Year. Visualizing Search Results using SQWID. *Proceedings of the Sixth International World Wide Web Conference*. Santa Clara, California, USA April 7-11, p.
MCKIERNAN, G., 2003 New Age navigation: Innovative information interfaces for electronic journals. *The Serials Librarian,* 45 (2)**,** p. 87-123.
MCPHERSON, K., 1983 Opinion-Related Information Seeking: Personal and Situational Variables. *Personality and Social Psychology Bulletin,* 9 (1)**,** p. 116.
MILLEN, D. R., FEINBERG, J. & KERR, B., Year. Dogear: Social bookmarking in the enterprise. *Proceedings of the SIGCHI conference on Human Factors in computing systems*. Montréal, Québec Canada, April 22-27, p. 111-120. ACM Press.
MILLER, G. A., 1995 WordNet: A Lexical Database for English. *Communications of the ACM,* 38 (11)**,** p. 39.
MODJESKA, D. K. (2000) Hierarchical Data Visualization In Desktop Virtual Reality. University of Toronto.

MOOERS, C. N., 1960 'Mooers' Law or Why some Retdeval systems are used and Others are not. *American Documentation,* 11 (3)**,** p.
MUNZNER, T., Year. H3: laying out large directed graphs in 3D hyperbolic space. *Proceedings of IEEE Symposium on Information Visualization*. Phoenix, AZ, USA, October 18-25, p. 2-10. IEEE Computer Society.
NATION, D. A., PLAISANT, C., MARCHIONINI, G. & KOMLODI, A., Year. Visualizing websites using a hierarchical table of contents browser: WebTOC. *Proceedings of the Third Conference on Human Factors and the Web*. Denver, CO, USA, June 12, p.
NAVARRO-PRIETO, R., SCAIFE, M. & ROGERS, Y., Year. Cognitive Strategies in Web Searching. *Proceedings of the Fifth Conference on Human Factors and the Web*. Gaithersburg, MD, USA, June 3, p.
PAEK, T., DUMAIS, S. & LOGAN, R., Year. WaveLens: a new view onto Internet search results. *Proceedings of the 2004 conference on Human factors in computing systems*. Vienna, Austria, April 24-29, p. 727-734. ACM Press.
PAULSON, L. D., 2005 Building rich web applications with Ajax. *Computer,* 38 (10)**,** p. 14-17.
PERUGINI, S., MCDEVITT, K., RICHARDSON, R., MANUEL PEREZ-QUIONES, SHEN, R., RAMAKRISHNAN, N., WILLIAMS, C. & FOX, E. A., Year. Enhancing usability in CITIDEL: multimodal, multilingual, and interactive visualization interfaces. *Proceedings of the 4th ACM/IEEE-CS joint conference on Digital libraries*. Tuscon, AZ, USA, June 7-11, p. 315 - 324. ACM Press.
PIROLLI, P. & CARD, S., Year. Information foraging in information access environments. *Proceedings of the SIGCHI conference on Human factors in computing systems*. Denver, CO, USA, May 7-11, p. 51-58. ACM Press.
PLAISANT, C., SHNEIDERMAN, B., DOAN, K. & BRUNS, T., 1999 Interface and data architecture for query preview in networked information systems. *ACM Transactions on Information Systems,* 17 (3)**,** p. 320-341.
PRATT, W., 1997 Dynamic organization of search results using the UMLS. *American Medical Informatics Association Fall Symposium,* 480 4.
RISDEN, K., CZERWINSKI, M. P., MUNZNER, T. & COOK, D. B., 2000 Initial examination of ease of use for 2 D and 3 D information visualizations of Web content. *International Journal of Human-Computers Studies,* 53 (5)**,** p. 695-714.
RIVADENEIRA, W. & BEDERSON, B. B. (2003) A Study of Search Result Clustering Interfaces: Comparing Textual and Zoomable User Interfaces. University of Maryland, HCIL.
ROBERTSON, G., CAMERON, K., CZERWINSKI, M. & ROBBINS, D., Year. Polyarchy visualization: visualizing multiple intersecting hierarchies. *Proceedings of the SIGCHI conference on Human factors in computing systems*. Minneapolis, Minnesota, USA, April 20-25, p. 423-430. ACM Press.
ROBERTSON, G., CZERWINSKI, M., LARSON, K., ROBBINS, D. C., THIEL, D. & DANTZICH, M. V., Year. Data mountain: using spatial memory for document management. *Proceedings of the 11th annual ACM symposium on User interface software and technology*. San Francisco, CA, USA, November 1-4, p. 53-162. ACM Press.
ROBERTSON, G. G., MACKINLAY, J. D. & CARD, S. K., Year. Cone Trees: animated 3D visualizations of hierarchical information. *Proceedings of the SIGCHI conference on Human factors in computing systems: Reaching*

through technology. New Orleans, Louisiana, USA, April 27 - May 2, p. 189-194. ACM Press.
RORVIG, M. & LUNIN LOIS, F., 1999 Introduction and Overview: Visualization, Retrieval, and Knowledge. *Journal of the American Society for Information Science,* 50 (9)**,** p. 790-793.
SANTINI, S. & JAIN, R., Year. Beyond query by example. *Proceedings of the sixth ACM international conference on Multimedia*. Bristol, UK, September 12-16, p. 345-350.
SCHRAEFEL, M. C., WILSON, M. L. & KARAM, M. (2004) Preview Cues: Enhancing Access to Multimedia Content. School of Electronics and Computer Science, University of Southampton.
SCHRAEFEL, M. C., WILSON, M. L., RUSSELL, A. & SMITH, D. A., 2006 mSpace: improving information access to multimedia domains with multimodal exploratory search. *Communications of the ACM,* 49 (4)**,** p. 47-49.
SEBRECHTS, M. M., CUGINI, J. V., LASKOWSKI, S. J., VASILAKIS, J. & MILLER, M. S., Year. Visualization of search results: a comparative evaluation of text, 2D, and 3D interfaces. *Proceedings of the 22nd annual international ACM SIGIR conference on Research and development in information retrieval*. Philadelphia, PA, USA, July 27-31, p. 3-10. ACM Press.
SHNEIDERMAN, B., 1997 Designing information-abundant web sites: issues and recommendations. *International Journal of Human-Computer Studies,* 47 (1)**,** p. 5-29.
SHNEIDERMAN, B., 2003 Why not make interfaces better than 3D reality? *Computer Graphics and Applications, IEEE,* 23 (6)**,** p. 12-15.
SHNEIDERMAN, B. & ARIS, A., 2006 Network Visualization by Semantic Substrates. *IEEE Transactions on Visualization and Computer Graphics,* 12 (5)**,** p. 733-740.
SHNEIDERMAN, B., BYRD, D. & CROFT, W. B., 1998 Sorting out searching: A user-interface framework for text searches. *Communications of the ACM,* 41 (4)**,** p. 95-98.
SHNEIDERMAN, B., FELDMAN, D., ROSE, A. & GRAU, X. F., Year. Visualizing digital library search results with categorical and hierarchial axes. *Proceedings of the Fifth ACM International Conference on Digital Libraries*. San Antonio, TX, June 2-7, p. 57-66. ACM Press.
SHNEIDERMAN, B. & PLAISANT, C. (2004) *Designing the User Interface: Strategies for Effective Human-Computer Interaction,* Boston, Pearson/Addison-Wesley.
SPINK, A., 2002 A user-centered approach to evaluating human interaction with web search engines: an exploratory study. *Information Processing & Management,* 38 (3)**,** p. 401-426.
SPINK, A. & WILSON, T. D., Year. Toward a theoretical framework for information retrieval (IR) evaluation in an information seeking context. *Proceedings of Multimedia Information Retrieval Applications*. Glasgow, UK, April 14-16, p. Paper 9. BCS.
STOICA, E. & HEARST, M. A., Year. Nearly-Automated Metadata Hierarchy Creation. *Proceedings of the Human Language Technology Conference of the North American Chapter of the Association for Computational Linguistics*. Boston, MA USA, May 2-7, p.

SUTCLIFFE, A. & ENNIS, M., 1998 Towards a cognitive theory of information retrieval. *Interacting with Computers,* 10 (3)**,** p. 321-351.
SUTCLIFFE, A. G. & PATEL, U., Year. 3D or not 3D: Is it Nobler in the Mind? *Proceedings of the BCS Human Computer Interaction Conference on People and Computers XI*. London, UK, August 20-23, p. 79-94. Springer-Verlag.
TANIN, E., PLAISANT, C. & SHNEIDERMAN, B., Year. Browsing Large Online Data with Query Previews. *Proceedings of the Symposium on New Paradigms in Information Visualization and Manipulation*. Washington D.C., USA, November 10, p. ACM Press.
TAYLOR, R. S., 1962 The process of asking questions. *American Documentation,* 13 (4)**,** p. 391-396.
TERVEEN, L., HILL, W. & AMENTO, B., 1999 Constructing, organizing, and visualizing collections of topically related Web resources. *ACM Transactions on Computer-Human Interaction,* 6 (1)**,** p. 67-94.
TUFTE, E. R. (1990) *Envisioning information*, Graphics Press Cheshire.
VAKKARI, P., Year. Cognition and changes of search terms and tactics during task performance: A longditudinal case study. *Proceedings of the International Conference on Computer Assisted Information Retrieval (RAIO)*. Paris, France, April 12-14, p.
WANG, P., BERRY, M. & YANG, Y., 2003 Mining londitudinal web queries: trends and patterns. *Journal of the American Society for Information Science and Technology,* 54 (8)**,** p. 742-758.
WHITE, R. W., RUTHVEN, I. & JOSE, J. M., Year. Finding relevant documents using top ranking sentences: an evaluation of two alternative schemes. *Proceedings of the 25th annual international ACM SIGIR conference on Research and development in information retrieval*. Tempere, Finland, August 11-15, p. 57-64. ACM Press.
WILDEMUTH, B. M., 2004 The effects of domain knowledge on search tactic formulation. *Journal of the American Society for Information Science and Technology,* 55 (3)**,** p. 246-258.
WILSON, M. & SCHRAEFEL, M. C., Year. Bridging the gap: Using IR models for evaluating exploratory search interfaces. *SIGCHI 2007 Workshop on Exploratory Search and HCI*. San Jose, CA, USA, 29 April, p.
WILSON, M. L., ANDRÉ, P., SMITH, D. A. & SCHRAEFEL, M. C. (2007) Spatial Consistency and Contextual Cues for Incidental Learning in Browser Design. School of Electronics and Computer Science, University of Southampton.
WILSON, M. L. & SCHRAEFEL, M. C., Year. mSpace: what do numbers and totals mean in a flexible semantic browser. *Proceedings of the 3rd International Semantic Web User Interaction Workshop*. Athens, GA, USA, November 5-9, p.
WILSON, M. L. & SCHRAEFEL, M. C., Year. A longitudinal study of both faceted and keyword search. *Proceedsing of the ACM/IEEE-CS Joint Conference on Digital Libraries (in submission)*. Pittsburgh, PA, USA, June 16-20, p. IEEE Computer Society.
WILSON, M. L. & SCHRAEFEL, M. C., Year. A Validated Framework for Measuring Interface Support for Interactive Information Seeking. *ACM Conference on Research and Development in Information Retrieval (in submission)*. Singapore, July 20-24, p. ACM Press.

WILSON, M. L., SCHRAEFEL, M. C. & WHITE, R. W., to come Evaluating Advanced Search Interfaces using Established Information-Seeking Models. *Journal of the American Society for Information Science and Technology*.

WOODRUFF, A., FAULRING, A., ROSENHOLTZ, R., MORRSION, J. & PIROLLI, P., Year. Using thumbnails to search the Web. *Proceedings of the SIGCHI conference on Human factors in computing systems*. Seattle, Washington, USA, March 31 - April 5, p. 198-205. ACM Press.

WU, H., ZUBAIR, M. & MALY, K., Year. Harvesting social knowledge from folksonomies. *Proceedings of the seventeenth conference on Hypertext and hypermedia*. Odense, Denmark, August 23-25, p. 111-114. ACM Press.

YANG, K., 2005 Information Retrieval on the Web. *Annual Review of Information Science and Technology,* 39 33-80.

YEE, K.-P., SWEARINGEN, K., LI, K. & HEARST, M., Year. Faceted metadata for image search and browsing. *Proceedings of the ACM Conference on Human factors in Computing Systems*. Fort Lauderdale, FL, USA, April 5-10, p. 401-408. ACM Press.

ZHANG, J. & MARCHIONINI, G., Year. Evaluation and evolution of a browse and search interface: relation browser. *Proceedings of the National Conference on Digital Government Research*. Atlanta, GA, USA, May 15-18, p. 179-188. Digital Government Research Center.