

# An Empirical Evaluation of User Interfaces for Topic Management of Web Sites

Brian Amento<sup>1,2</sup>, Will Hill<sup>1</sup>, Loren Terveen<sup>1</sup>, Deborah Hix<sup>2</sup>, and Peter Ju<sup>1</sup>

<sup>1</sup>AT&T Labs - Research  
180 Park Avenue, P.O. Box 971  
Florham Park, NJ 07932 USA  
{pju, willhill, terveen}@research.att.com

<sup>2</sup>Department of Computer Science  
Virginia Tech  
Blacksburg, VA 24061 USA  
{brian, hix}@cs.vt.edu

## ABSTRACT

*Topic management* is the task of gathering, evaluating, organizing, and sharing a set of web sites for a specific topic. Current web tools do not provide adequate support for this task. We created the *TopicShop* system to address this need. TopicShop includes (1) a webcrawler that discovers relevant web sites and builds site profiles, and (2) user interfaces for exploring and organizing sites. We conducted an empirical study comparing user performance with TopicShop vs. Yahoo™. TopicShop subjects found over 80% more high-quality sites (where quality was determined by independent expert judgements) while browsing only 81% as many sites and completing their task in 89% of the time. The site profile data that TopicShop provides – in particular, the number of pages on a site and the number of other sites that link to it – was the key to these results, as users exploited it to identify the most promising sites quickly and easily.

## KEYWORDS

information access, information retrieval, information visualization, human-computer interaction, computer supported cooperative work, social filtering

## INTRODUCTION: THE TOPIC MANAGEMENT TASK

An important task that many web users perform is gathering, evaluating, and organizing relevant information resources for a given topic; we call this *topic management*. Sometimes users investigate topics of professional interest, at other times topics of personal interest. Users may create collections of web information resources for their own use or for sharing with coworkers or friends. For example, someone might gather a collection of web sites on wireless telephony as part of a report they're preparing for their boss and a collection on the X-Files as a service to their fellow fans. Librarians might prepare topical collections for their clients, and teachers for their students [1].

Topic management is a difficult task that is not supported

well by current web tools. A common way to find an initial set of (potentially) relevant resources is to use a search engine like AltaVista or a directory like Yahoo. At this point, however, a user's work has just begun: the initial set usually is quite large, consisting of dozens to hundreds of sites of varying quality and relevance, covering assorted aspects of the topic. Users typically want to select a manageable number – say 10 to 20 – of high-quality sites that cover the topic. With existing tools, users simply have to browse and view resources one after another until they are satisfied they have a good set, or, more likely, they get tired and give up. Browsing a web site is an expensive operation, both in time and cognitive effort. And bookmarks, probably the most common form of keeping track of web sites, are a fairly primitive organizational technique.

We have designed a system called *TopicShop* to support the topic management task directly. TopicShop includes three main components:

- a webcrawler, which discovers web sites relevant to a user-specified set of seed sites and creates site profiles which can be used to inform user evaluation of sites;
- a Java applet, which serves as a front end to the webcrawler; it lets users specify seed sites for a topic, track the progress of the webcrawler, and explore sites as they are discovered and profiled; and
- the TopicShop Explorer, an interface that lets users view thumbnail images of sites and site profiles, sort sites by their properties, organize sites into categories, both spatially and by using folders, and easily share results by exchanging files via email, disks, etc.

This paper reports on an empirical evaluation of the TopicShop Explorer. Since Yahoo is a state of the art, widely used tool for exploring collections of web resources on specific topics, we chose to investigate how subjects performed on topic management tasks using either TopicShop or Yahoo.

In the remainder of the paper, we discuss related work, describe the features of TopicShop, and discuss the design and results of the study in detail.

## RELATED WORK

Our work aims to create algorithms for extracting useful information from the web and interfaces to aid users in comprehending and organizing web sites. We have written previously [13] about information structures and analysis algorithms we developed. We also performed studies showing that links between web sites carry useful information, in particular, that connectivity correlates with expert quality judgements [14]. The work reported here builds on this previous work, but focuses on user interfaces. In particular, it investigates two major issues:

- what useful information can we provide to users to inform their evaluation of web resources, enabling them to browse fewer and higher quality sites?
- how can we design interfaces that help users easily and effectively explore collections of resources and organize resources for later use by themselves and others?

Others have looked at these questions. For example, Abrams, Baecker, and Chignell [1] carried out a study of how several hundred web users used bookmarks. Bookmarks were a very popular way to create personal information spaces of web resources. They observed a number of strategies for organizing bookmarks, including a flat ordered list, a single level of folders, and hierarchical folders. They also made four design recommendations to help users manage their bookmarks more effectively. First, bookmarks must be easy to organize, e.g., via automatic sorting techniques. Second, visualization techniques are necessary to provide comprehensive overviews of large sets of bookmarks. Third, rich representations of sites are required; many users noted that site titles are not accurate descriptors of site content. Finally, tools for managing bookmarks must be well integrated with web browsers.

Our focus is different than Abrams et al. Bookmarks typically are gathered opportunistically, even casually, as users happen to encounter interesting sites, and bookmark files usually span many different topics. In contrast, we are interested in situations where users are explicitly engaged in gathering and organizing a collection of related resources for a specific topic. Nonetheless, many of Abrams et al.'s specific observations and recommendations are relevant to our work. In particular, as we illustrate later, TopicShop supports all the user organizational strategies they observed and follows their design recommendations.

A number of researchers have created interfaces to support users in managing collections of information resources. SenseMaker [2] focuses on supporting users in the contextual evolution of their interest in a topic. They attempt to make it easy to evolve a collection, e.g., expanding it by query-by-example operations or limiting it by applying a filter. Scatter/Gather [11] supports the browsing of large collections of text, allowing users to iteratively reveal topic structure and locate desirable

documents. Card, Robertson, and York [3] describe the WebBook, which uses a book metaphor to group a collection of related web pages for viewing and interaction, and the WebForager, an interface that lets users view and manage multiple WebBooks. Mackinlay, Rao, and Card [8] developed a novel user interface for accessing articles from a citation database. The central UI object is a "Butterfly", which represents one article, its references, and its citers. The interface makes it easy for users to browse from one article to a related one, group articles, and generate queries to retrieve articles that stand in a particular relationship to the current article.

Other researchers have investigated techniques for determining information about web sites to aid user comprehension of the sites. For example, Kleinberg [7] defines algorithms that identify *authoritative* and *hub* pages within a hypertext document. Authorities and hubs are mutually dependent: a good authority is a page that is linked to by many hubs, and a good hub is one that links to many authorities. Several systems, including WebQuery [4] and twURL [15], group or sort pages by their degree (the number of other pages in the collection they are connected with). Pirolli, Pitkow, and Rao [10] developed a categorization algorithm that used hyperlink structure, text similarity, and user access data to categorize web pages into various functional roles, such as "head", "index", and "content". Later Pitkow and Pirolli [12] experimented with clustering algorithms based on co-citation analysis [5], in which pairs of documents were clustered based on the number of times they were both cited by a third document.

While our work shares many goals of previous work, it differs in several respects. First, as discussed elsewhere [13], we designed novel information structures and algorithms for building collections of resources. Second, we seek to provide exploration and organization techniques that are both powerful and very easy to use; we attempt to achieve this goal with a user interface that is an enhanced version of the Microsoft Windows™ file explorer that users already are familiar with. Finally, and most important, we performed an empirical study that has yielded both quantitative results that show the utility of TopicShop and qualitative results that help us better understand the topic management task and how to support it more effectively.

## THE TOPICSHOP SYSTEM

### The Webcrawler and Java Applet

Topic management obviously begins with the identification of a collection of relevant resources. We have developed an algorithm [13] that takes as input a user-specified set of web sites (the "seeds") and follows links from the seeds to construct a graph of the seed sites and closely connected sites. Sites reached by following links from the seeds are likely to be on the same topic.

For purposes of this paper, however, the discovery of new sites is not relevant; what is relevant is that our webcrawler

heuristically groups web pages into web sites and builds profiles of the sites it fetches.

Pages are grouped into sites using heuristics that look at the directory structure of URLs. For example, if the crawler encounters a link to the URL `http://a/b/page1.html`, and `http://a/b/index.html` is a site known to the crawler, it records this URL as part of the site. Further, if the link was encountered while the crawler was analyzing the site `http://x/y/`, a link is recorded from the site `http://x/y/` to the site `http://a/b/index.html`.

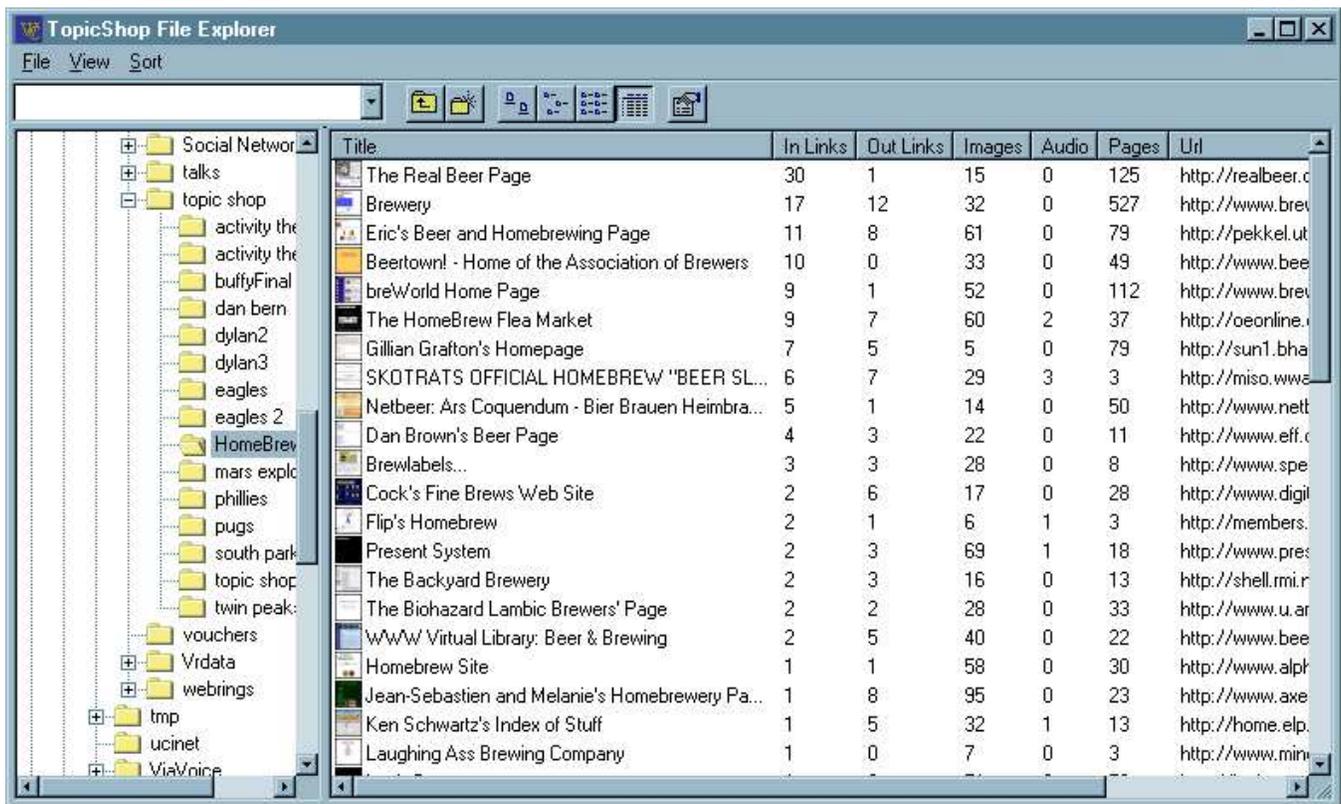
Site profiles are built by fetching a fairly large number of pages from each site. Profiles contain the following data:

- title (of the site's root page),
- a thumbnail image (of the site's root page),
- links to and from other sites,
- internal html pages, images, audio files, and movie files.

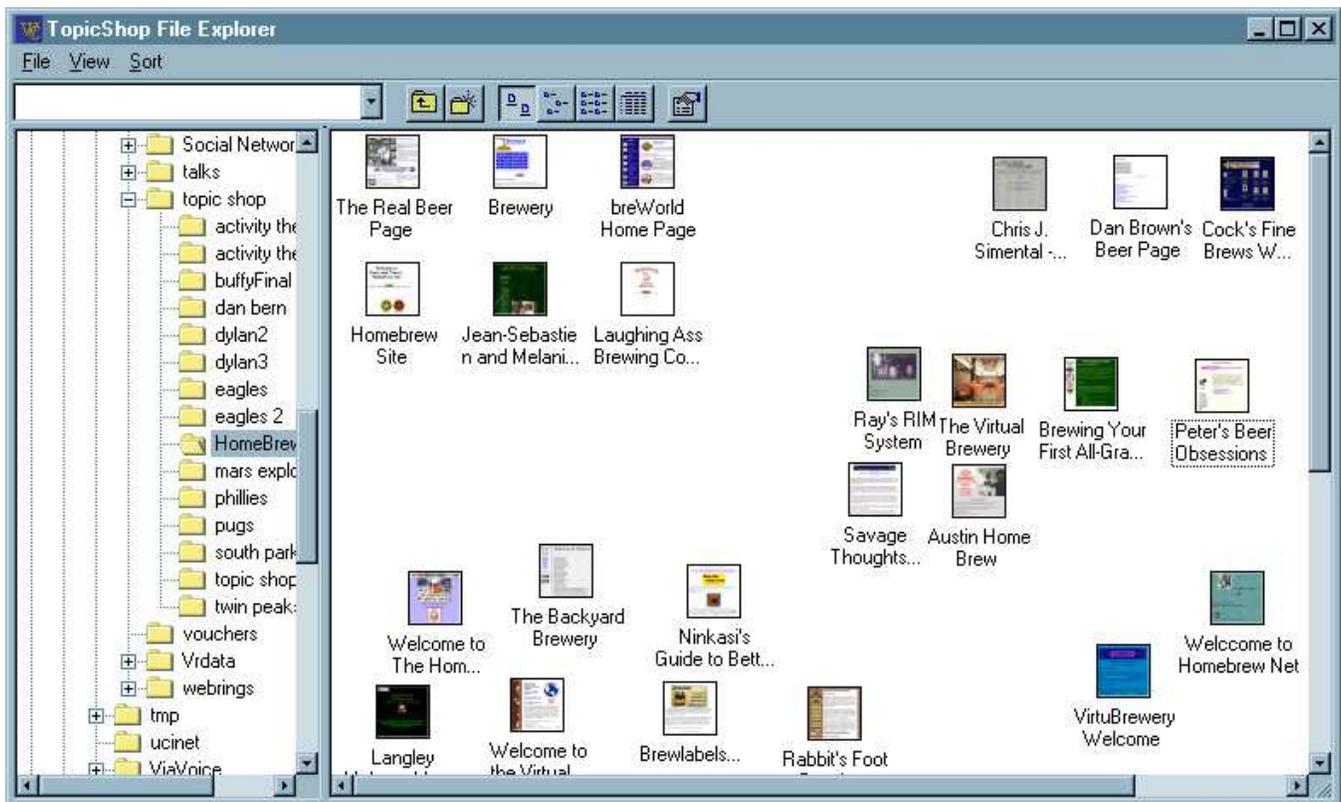
Users build a TopicShop collection by identifying a set of seed sites on a topic of interest (e.g., by finding a topic-specific index page or using a search engine) and then running a Java applet that sends the seed pages to the webcrawler running on our server. As the crawler runs, it continually sends its results to the applet; thus, the user can begin viewing and exploring sites immediately. When a crawl is completed, the server creates a special type of Microsoft Windows™ file for each site that contains the site profile data. The server then compresses all the files into one archive file which is downloaded to the user's computer.

### The TopicShop Explorer: Evaluating, Managing, and Sharing Collections of Resources

Once users save and uncompress the downloaded file, they view and manage their collections using the TopicShop Explorer, a customized version of the normal Windows file explorer. The TopicShop Explorer is a small Windows executable that interprets and processes site profile files.



**Figure 1: TopicShop Explorer, details view.** Each web resource (site) is represented by a small thumbnail image, the site title, properties of the site itself – number of internal pages, images, and audio files – and the site's relations to other sites – number of in links and out links. By clicking on a column, users can sort by the appropriate property.



**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.**

Users can view their collections in two different ways: details (Figure 1) or icons (Figure 2). The main feature of the details view is that it shows site profile information, and the main feature of the icons view is that users can arrange icons spatially. We will explain the user interface properties further as we consider the three main design goals of the TopicShop Explorer.

### Design Goals

1. *Make relevant but invisible information visible.* We hypothesize that making site profile information visible will significantly inform users in evaluating a collection of sites. No longer must they decide to visit sites — a time-consuming process — based solely on titles and (sometimes) brief textual annotations. (A chief complaint of subjects in the Abrams et al [1] study was that titles were inadequate descriptors of site content — and that was for sites that users already had browsed and decided to bookmark.) Instead, users can choose to visit only sites that have been endorsed (linked to) by many other sites or sites that are rich in a particular type of content (e.g., images or audio files). In addition to site profile data, the thumbnail images also are quite useful; most notably, they are effective visual identifiers for sites users already have visited.
2. *Make it simple for users to explore and organize resources.* In the details view, users can sort resources

by any property (e.g. number of in-links, out-links, images, etc.) simply by clicking on the label at the top of the column that displays that property. In either view, right-clicking on a site brings up a window that shows the profile data from which the numbers in the columns are derived, e.g., lists of all sites that link to the selected site and all internal pages of the site. Double-clicking on a site will send the user's default web browser to that site.

Users can organize resources both spatially (in the icons view) and by creating subfolders and moving resources into the subfolders. Nardi & Barreau [9] found that users of graphical file systems preferred spatial location as a technique for organizing their files. We believe spatial organization is particularly useful early in the exploration process while users are still discovering important distinctions among resources and user-defined categories have not yet explicitly emerged. As categories do become explicit, users can create folders to contain sites in each of the categories.

3. *Integrate topic management into a user's normal computing and communications environment.* The TopicShop Explorer may not look like a novel interface at all; interestingly enough, this was an explicit goal. We wanted it to be as similar to the normal Windows Explorer as possible so Windows

users could apply all their existing knowledge, thus minimizing learning time and maximizing ease of use. Further, this decision makes it very easy for collections of resources to be shared. Since a collection is just a normal Windows folder containing files (of the special type that we designed), they can be shared in all the normal ways. As we already have explained, a collection can be compressed and downloaded. It can also be emailed. And if users share a common network, collections simply can be read directly from any machine on the network.

## **EMPIRICAL STUDY**

We wanted a suitable baseline topic management tool for comparison to TopicShop. Yahoo is the most popular tool for locating collections of web sites (according to the Media Metrix PC metering company, results available at <http://searchwnginewatch.com/reports/mediamatrix.html>).

Bookmark lists are a common means of organizing collections of web resources. Therefore, we decided that subjects would use either TopicShop or Yahoo/bookmarks.

We chose two topics for the study: home brewing (of beer) and the TV program “Buffy the Vampire Slayer” – each contained about 60 sites in their corresponding Yahoo category. Our choice of these topics was influenced by the fact that pursuing special interests, including hobbies and media fandom, is one of the main ways people use the web. To quantify this, we studied a set of approximately 770K queries issued to the Magellan search engine between March 1997 and August 1998. We determined that 42% of the queries had to do with entertainment topics, including media fandom (categorization was done on the 515 most popular query strings, which cumulatively accounted for 96K query instances; categorization was done by two independent raters, inter-rater reliability of 87%.)

### **Design**

The experiment was a 2x2, between subjects design, with topic (home brewing or Buffy) and user interface (TopicShop or Yahoo) as factors. Sixteen members of our lab volunteered to participate, giving four subjects per each of four conditions. None of the subjects had seen TopicShop before, although some were familiar with the general concepts.

The two main metrics we wanted to measure were the quality of resources users gathered and the amount of effort (time and total number of sites browsed) required. To give a quality baseline, four experts for each topic were presented a list of the sites (in random order) on that topic; only titles were presented, no Yahoo annotations or TopicShop profile data. This meant that the experts had to browse each site and evaluate it based on its content and layout. Each expert collected the 20 “best” sites. For this study, we defined “best” as a set of sites that collectively provided a useful and comprehensive overview for someone wanting to learn about the topic. During analysis, we used

the “expert intersection”, the set of resources that all experts for a given topic selected, as the yardstick for measuring the quality of resources selected by the subjects.

Subjects for a given topic, whether they used TopicShop or Yahoo, were presented with the same set of approximately 60 sites (obtained from the Yahoo category) to evaluate. Yahoo subjects saw (as usual) site titles and, for about half the sites, a brief textual annotation. For the TopicShop condition, we applied our webcrawler to the Yahoo sites to produce site profiles; TopicShop subjects thus had access to site tiles, thumbnail images, and profile data, as shown in Figures 1 and 2.

### **Methodology**

Subjects were assigned randomly to one of the four conditions. To begin the experiment, subjects received 15 minutes of instruction and training in the task and user interface. TopicShop subjects were shown the basic interface features and taught how to collect sites by dragging and dropping icons into folders. Yahoo subjects were shown a sample list of sites and taught how to collect sites by bookmarking. After training, subjects performed a short task to ensure that they were comfortable with collecting and organizing sites.

For the main task, subjects investigated the sites for their assigned topic by using the interface (TopicShop or Yahoo) and browsing to sites. Subjects within a single topic were presented with the same collection of sites in both interface conditions. They were asked to choose the 15 “best” (as defined previously) sites and rank them by quality. Subjects were asked to complete the task in 45 minutes and were kept informed of the elapsed time. Clearly, there is a relationship between time on task and quality of results: the more time spent, the better results one can expect. By limiting the amount of time, we hoped to focus on any differences in the quality of results (i.e., the sites users selected) between the two interfaces. And people don’t spend unlimited amounts of time browsing, so we wanted to see whether users could find high-quality sites in a limited amount of time.

The task ended when subjects were satisfied with their collections of sites. Subjects then completed a short questionnaire. Finally, an informal interview was conducted to reveal strategies subjects used to perform the task, their reactions to the interface, and what could help them to complete the task more effectively.

### **Results**

We first compared the set of resources chosen by each subject to the expert intersection. For each topic, the expert intersection contained 12 resources. For the Buffy topic, Yahoo subjects selected an average of 5.0 sites that were in the expert intersection, while TopicShop subjects selected 7.5 expert-endorsed sites. For home brewing, Yahoo subjects matched 4.3 sites and TopicShop subjects matched 9.3. Overall, Yahoo subjects selected 4.6 sites from the

expert intersection, while TopicShop subjects selected over 80% more, or 8.4 ( $p < 0.05$ ). These results are summarized in Table 1. Notice that choosing sites at random would result in obtaining 3 sites in the expert intersection. (Users selected 15 out of 60 sites, or 25%; 25% of the 12 sites in the expert intersection is 3 sites.) The Yahoo score of 4.6 is not that much better than random selection. This probably is due to task time limit of 45 minutes. If Yahoo subjects had had unlimited time, undoubtedly they would have been able to find more high quality sites. To sum up, we see that TopicShop users found significantly better resources in the time given to complete the task.

#### Mean Number of High-Quality sites Identified

<i>Topic</i>	<i>Interface Type</i>	
	<b>Yahoo</b>	<b>TopicShop</b>
<b> Buffy </b>	5.0	7.5
<b> Home brewing </b>	4.3	9.3
<b> Average over Topic </b>	<b> 4.6 </b>	<b> 8.4 </b>

**Table 1: Expert intersection analysis**

It also is revealing to examine the amount of work subjects performed to complete their tasks. A study of data from the search engine Excite (51,473 queries, 18,113 users) showed that 86% of all users looked at no more than 30 pages returned in response to their query [6]. In our study, Yahoo users browsed an average of 44 sites, while TopicShop subjects visited about 36, or about 19% less. Further, the task of constructing a high-quality collection of resources is more difficult than doing a simple search; the task is global, since one is trying to develop a comprehensive overview of a topic, so more sites must be considered. By providing additional dynamic data up front, TopicShop enables users to make better decisions about which sites to immediately rule out and which to investigate further. Yahoo users can rely only on textual annotations, which are provided by site maintainers. While these annotations are sometimes helpful, they can be out-of-date or self-promotional, so are not necessarily good indications of the perceived quality of a site.

We also analyzed time on task. We did not expect a large difference since we gave users a (soft) limit of 45 minutes to complete the task and kept them aware of elapsed time during the experiment. Still, TopicShop subjects took about 11% less time (41.5 minutes vs. 46.6 minutes for Yahoo).

While the differences in time and effort were not statistically significant (although we hope they will be in a larger study we are preparing to conduct), they do show that TopicShop subjects did not obtain better quality results at the cost of more work.

The questionnaire gave us data on what information subjects found most useful in evaluating a site. TopicShop

subjects were asked to rank the utility of the site profile attributes, including the title and the number of in-links, out-links, images, audio files, and pages on the site. Subjects ranked these properties on a scale of 1 (most useful) to 7 (least useful). Three of these properties — in-links (2.00), title (2.75), and number of pages (3.00) — were ranked most highly. The other four properties had an average score greater than 5. Even though many subjects noted that title is not a very good indication of quality, it still was perceived as one of the most useful site properties. In interviews, subjects explained that titles were useful mainly as memory aids for sites. Thus, subjects considered the number of endorsements (in links) and the size of a site (in pages) to be the most useful indicators of quality.

The questionnaire also asked subjects what additional information would have helped them in evaluating sites. Six of the eight Yahoo subjects said that the number of links between sites would be very useful. One subject even made it a point to go to the links page of every site visited to see not only what sites were linked to, but also to read any annotations or recommendations made by the site author. Thus, link information was rated as highly useful by those subjects who saw it and as desirable by those subjects who did not.

#### User Exploration Strategies

Most Yahoo subjects, lacking any better options, simply looked through the initial 60 sites in alphabetical order, reverse alphabetical order, or sometimes a combination of the two. A few users tried reading all the titles and annotations to make some judgements about the sites before browsing them; however, many times the initial judgement of a site proved inaccurate once it was browsed, so even these users often reverted to exhaustive alphabetical search. Of course, users still read annotations as they proceeded methodically through the list of sites, but did not rely on the annotations to decide which sites to browse. Users also often browsed a few sites at random to try to cover a good sample of the available sites.

TopicShop subjects used different strategies, ones that were informed by the data in the TopicShop Explorer. They spent more time prior to browsing sites on exploration within the TopicShop interface, sorting the columns and watching how the arrangement of sites changed. They were mainly looking for sites that appeared near the top in multiple sorts. Many also attempted to get a rough idea of how sites were distributed in each column. Eventually, subjects tended to proceed by selecting a property they thought was useful and evaluating the first few sites in that column. After they exhausted the quality sites in the column, they would move on to another column and continue. Some subjects would also visit some sites at the low end of the data columns to convince themselves that the profile data could be trusted.

As evidence of the influence of the TopicShop explorer on user strategies, let us consider the overlap in sites selected

by subjects. TopicShop subjects arrived at a much larger common set of sites. The intersection for the eight TopicShop subjects across both topics was 9.5 sites, while the eight Yahoo subjects averaged an intersection of only 2.5 sites. It makes sense that TopicShop users would agree with each other quite a bit, even more than they agreed with the experts, since they relied on the same data and tended to pursue the same strategies for selecting resources.

To better evaluate the utility of the TopicShop data, we created purely automated versions using the “gather from the top of the column” strategy. We defined six sets of sites mechanically: five of the sets consisted simply of the top 15 sites for each numeric site profile attribute, and the sixth consisted of the top three sites on each attribute.

Recall that the Yahoo subjects had an overall average expert intersection of 4.6 (out of 12). All the automated TopicShop strategies performed better, with an average expert intersection of 5.6. We found it surprising and noteworthy that a purely mechanical strategy using only the TopicShop site profile data could outperform human subjects who had to rely only on Yahoo’s site titles and annotations. (Again, we assume that the task time limit was a factor; with enough time to browse and evaluate site content, we expect that people would outperform these mechanical strategies. Of course, who has enough time?)

We also observed a common, but unproductive strategy: nearly all subjects initially assumed that personal home pages (as determined by the title and site location) would be of low quality. They supposed that they could immediately eliminate these sites and select only from the resulting, smaller subset. However, subjects quickly realized that this was not true – after visiting a few personal pages, they found that some were of quite high quality, so subjects abandoned this strategy.

### **Design Implications**

Observations, interviews, and questionnaires suggested four significant design improvements to the TopicShop user interface.

The first design improvement is to add additional site profile data about the front page of a site. We observed that most subjects made their judgement of a site by viewing only the front page of the site. It makes sense that the “front door” page of a site should be both attractive and representative of the site as a whole – after all, the site author designs it to be the initial impression a visitor to the site experiences. The front page usually gives a good idea of the amount and type of content available on the site as well as the production quality.

Users navigated to a total of 639 web sites, and looked at only the front page of over half the sites. And of the 240 sites that users selected for their collection of the best resources, users browsed only the front page of 91. Among the 399 sites that users rejected, 285 sites were rejected after browsing the front page. Overall, users viewed an

average of 2.39 pages per site. Thus, we see that a user’s initial impression of a site is extremely important. Therefore, we will analyze site front door pages and present the results as additional site profile data in a future design.

The second design change we will incorporate into TopicShop is better methods for creating subcategories of a topic. A key need that subjects in both interface conditions discussed was support for lightweight, flexible categorization. As users explore sites, they create rough mental groupings, using site similarity, site type (general information sites, specific subtopic sites, personal sites, etc.), or even site layout.

While TopicShop lets users create folders and group subcategories of sites within folders, our observations of subjects showed that this seems to be too much overhead for users when they are starting out. Their mental groupings remain indistinct until they have encountered a sufficient number and variety of sites to enable them to articulate the organizing principle of their categories. Further, categories may be split or combined several times in early stages of exploration. And while the icons view of TopicShop does support this flexible, lightweight categorization (and several subjects used and liked it), this view hides the important site profile data from immediate view. We have two potential design solutions that could be added to TopicShop to better support categorization.

Linked views are one solution to this problem. One window would show the icons view, another would show the details view, with user selections mirrored in both windows. Users then could spatially arrange sites as they form opinions about the types of sites within a topic, while simultaneously sorting sites based on profile data. As users develop firm categories, they could create folders to hold sites within each category.

Another potential design solution is a coloring scheme. Users could assign a color to a group of sites at any time, then add other sites as they continue to browse. Then, whenever sites are sorted, they would be sorted first by color (i.e., group), then whatever other property the user specified. This would let users quickly create groups and still keep all sites in a single window. Again, when users are satisfied that a group really is a category, a folder can be created to contain it.

A third improvement to the design of TopicShop is to add two levels of annotations. One of the TopicShop design goals was to make it easy to reuse and share topical collections. Subjects affirmed that this was important. In support of this desire, all 16 subjects mentioned that they wanted to record comments about sites as they visited and collected them. Comments can be recorded for individual sites as well as user defined categories. These comments would be useful both to the original users when they returned to their collections in the future and to people with whom they shared the collections. The comments would

explain why sites were selected, why they were considered to be of high quality, and what they were good for.

The final design change involves sorting techniques within TopicShop. Currently, sorting in TopicShop is limited to a single column, but subjects expressed a desire for several more powerful sorting techniques. First, they wanted to combine several columns, e.g., sorting by the sum of in-links and out-links. Second, they wanted to be able to do a multi-level sort. For example, one might want to sort sites primarily by number of pages, then break ties by using another property, e.g., the number of in-links.

## CONCLUSIONS

As the amount of information on the web continues to grow, tools that support users in finding and managing collections of topical resources will become increasingly significant. The focus must move from compiling collections to helping users comprehend and manage them. Our goal is to reduce the time users must spend sifting through “relevant” – but poor quality – sites and increase the amount of time they can devote to exploring high-quality information.

We showed that TopicShop users can find nearly twice as many high-quality web sites while considering fewer sites and taking less time than Yahoo users. By mining the rich data that already exist in the structure of web sites and content of their pages, TopicShop helps users quickly identify small and manageable subsets of web resources. Finally, the user exploration strategies and additional user needs that our study revealed will lead to significant improvements in the design of TopicShop and similar tools.

## ACKNOWLEDGMENTS

We thank all the subjects who participated in our studies. We also appreciate design suggestions and many other useful comments from Julia Hirschberg, Lynn Cherny, Erik Ostrom, Steve Whittaker, and Bonnie Nardi.

## REFERENCES

1. Abrams, D., Baecker, R., and Chignell, M. Information Archiving with Bookmarks: Personal Web Space Construction and Organization, in *Proceedings of CHI'98* (Los Angeles CA, April 1998), ACM Press, 41-48.
2. Baldonado, M.Q.W., and Winograd, T. An Information-Exploration Interface Supporting the Contextual Evolution of a User's Interests, in *Proceedings of CHI'97* (Atlanta GA, March 1997), ACM Press, 11-18.
3. Card, S.K., Robertson, G.C., and York, W. The WebBook and the Web Forager: An Information Workspace for the World-Wide Web, in *Proceedings of CHI'96* (Vancouver BC, April 1996), ACM Press, 111-117.
4. Carrière, J., and Kazman R. WebQuery: Searching and Visualizing the Web through Connectivity, in *Proceedings of WWW6* (Santa Clara CA, April 1997).
5. Garfield, E. *Citation Indexing*. ISI Press, Philadelphia, PA, 1979.
6. Jansen, B. J., Spink, A., Bateman, J., and Saracevic, T. Searchers, the Subjects They Search, and Sufficiency: A Study of a Large Sample of EXCITE Searches, submitted to *WebNet'98*.
7. Kleinberg, J.M. Authoritative Sources in a Hyperlinked Environment, in *Proceedings of 1998 ACM-SIAM Symposium on Discrete Algorithms*.
8. Mackinlay, J.D., Rao, R., and Card, S.K. An Organic User Interface for Searching Citation Links, in *Proceedings of CHI'95* (Denver CO, May 1995), ACM Press, 67-73.
9. Nardi, B. and Barreau D. Finding and Reminding: File Organization from the Desktop. *ACM SIGCHI Bulletin*, 27, 3, July 1995.
10. Pirolli, P., Pitkow, J., and Rao, R. Silk from a Sow's Ear: Extracting Usable Structures from the Web, in *Proceedings of CHI'96* (Vancouver BC, April 1996), ACM Press, 118-125.
11. Pirolli, P., Schank, P., Hearst, M., and Diehl, Scatter/Gather Browsing Communicates the Topic Structure of a Very Large Text Collection, in *Proceedings of CHI'96* (Vancouver BC, April 1996), ACM Press, 213-220.
12. Pitkow, J., and Pirolli, P. Life, Death, and Lawfulness on the Electronic Frontier, in *Proceedings of CHI'97* (Atlanta GA, March 1997), ACM Press, 383-390.
13. Terveen, L.G., and Hill, W.C. Finding and Visualizing Inter-site Clan Graphs, in *Proceedings of CHI'98* (Los Angeles CA, April 1998), ACM Press, 448-455.
14. Terveen, L.G., and Hill, W.C. Evaluating Emergent Collaboration on the Web, in *Proceedings of CSCW'98* (Seattle WA, November 1998), ACM Press.
15. *What is twURL?* <http://www.roir.com/whatis.htm>