

TopicShop: Enhanced Support for Evaluating and Organizing Collections of Web Sites

Brian Amento^{1,2}, Loren Terveen¹, Will Hill¹, and Deborah Hix²

¹AT&T Labs - Research
180 Park Avenue, P.O. Box 971
Florham Park, NJ 07932, USA

E-mail: {brian, terveen, willhill}@research.att.com

²Department of Computer Science
Virginia Tech

Blacksburg, VA 24061 USA

E-mail: hix@cs.vt.edu

ABSTRACT

TopicShop is an interface that helps users evaluate and organize collections of web sites. The main interface components are *site profiles*, which contain information that helps users select high-quality items, and a *work area*, which offers thumbnail images, annotation, and lightweight grouping techniques to help users organize selected sites. The two components are linked to allow task integration.

Previous work [2] demonstrated that subjects who used *TopicShop* were able to select significantly more high-quality sites, in less time and with less effort. We report here on studies that confirm and extend these results. We also show that *TopicShop* subjects spent just half the time organizing sites, yet still created more groups and more annotations, and agreed more in how they grouped sites. Finally, *TopicShop* subjects tightly integrated the tasks of evaluating and organizing sites.

INTRODUCTION

In previous work [2], we motivated an important task for web users – gathering, evaluating, and organizing information resources for a given topic. Current web tools do not support this task well; specifically, they do not make it easy to evaluate collections of web sites to select the best ones or to organize sites for future reuse and sharing. Users have to browse and view sites 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, the most common form of keeping track of web sites, are a fairly primitive organizational technique.

We designed and implemented the *TopicShop* system to provide comprehensive, integrated support for this task. *TopicShop* aids users in finding a set of relevant sites, in narrowing down the set into a smaller set of high quality sites, and in organizing sites for future use. *TopicShop* has evolved through a number of design iterations, driven by extensive user testing. We report here on lessons we learned from a pilot study, how these lessons re-shaped our

understanding of the task and led to a significant re-design, and the results of a second, larger user study.

RELATED WORK

Our research program investigates the major information problems faced by users of the World Wide Web:

- finding collections of items *relevant* to their interests;
- identifying *high-quality* items within a collection;
- finding items that contain a certain *category* of information, e.g., episode guides (for a television show) or song lyrics (for a musician);
- *organizing* personalized subsets of items.

We have addressed these problems by developing algorithms, implementing them in web crawling and analysis tools, creating interfaces to support users in exploring, comprehending, and organizing collections of web documents, and performing user studies [2, 3, 4, 15]. The work reported here focuses on understanding the user tasks of evaluating and organizing collections of web sites, as illuminated by the design, evaluation, and re-design of interfaces to support these tasks.

Other researchers have investigated these issues. Much recent work has been devoted to algorithms for adding meta-information to collections of web sites to enhance user comprehension, typically by analyzing the structure of links between sites. This approach builds on the intuition that when the author of one site chooses to link to another, this often implies both that the sites have similar content and that the author is endorsing the content of the linked-to site. Pirolli, Pitkow and colleagues [12, 13] experimented with link-based algorithms for clustering and categorizing web pages. Kleinberg's HITS algorithm [8] defines *authoritative* and *hub* pages within a hypertext collection. 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.

After evaluating items and selecting the interesting ones, users must organize the items for future use. Card, Robertson, and Mackinlay [5] introduced the concept of information workspaces to refer to environments in which information items can be stored and manipulated. A departure point for most such systems is the file manager popularized by the Apple Macintosh and then in Microsoft Windows. Such systems typically include a list view, which shows various properties of items, and an icon view,

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

UIST '00, San Diego, CA USA

© 2000 ACM 1-58113-212-3/00/11... \$5.00

which lets users organize icons representing the items in a 2D space. Mander, Salomon, and Wong [10] enhanced the basic metaphor with the addition of “piles”. Users could create and manipulate piles of items. Interesting interaction techniques for displaying, browsing, and searching piles were designed and tested.

Bookmarks are the most popular way to create personal information workspaces of web resources. Bookmarks consist of lists of URLs; typically the title of the web page is used as the label for the URL. Users may organize their bookmarks into a hierarchical category structure. Abrams, Baecker, and Chignell [1] carried out an extensive study of how several hundred web users used bookmarks. 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.

Many researchers have created experimental information workspace interfaces, often designed expressly for web documents. Card, Robertson, and York [5] 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 [9] developed a novel user interface for accessing articles from a citation database. The central UI object is a “Butterfly”, which represents an article, its references, and its citers. The interface makes it easy for users to browse among related articles, group articles, and generate queries to retrieve articles that stand in a particular relationship to the current article. The Data Mountain of Robertson et al [14] represents documents as thumbnail images in a 3D virtual space. Users can move and group the images freely, with various interesting visual and audio cues used to help users arrange the documents. In a study comparing the use of Data Mountain to Internet Explorer Favorites, Data Mountain users retrieved items more quickly, with fewer incorrect or failed retrievals.

Our research shares goals with much of the previous work. We focus on designing interfaces that make automatically extracted information about web sites readily accessible to users. We show that this increases users’ ability to select high-quality sites. Through ongoing user studies and re-design, we developed easy-to-use annotation and grouping techniques that let users organize items better and more quickly. Finally, we learned how users interleave work on various tasks and re-designed our interface to support such task interleaving.

TOPICSHOPEXPLORER, VERSION 1

The TopicShop Explorer is implemented in C++ and runs on Microsoft Windows platforms. Version 1 was based directly on the Macintosh file manager / MS Windows Explorer metaphor (see [3] for detail of TopicShop Version 1 and the pilot study). Accordingly, users could view collections in either a details (Figure 1) or icons (Figure 2) view. The details view showed *site profile* information (see below) to help users evaluate sites, and the icons view let users organize sites spatially.

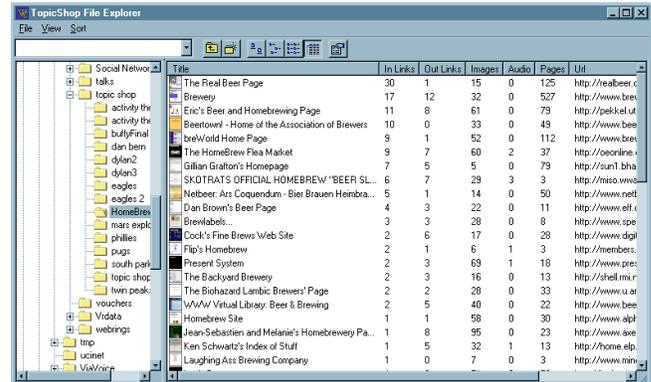


Figure 1: TopicShop Explorer (version 1), details view. Each web site is represented by a small thumbnail image, the site title, and profile data including the links to/from other sites in the collection, and the number of pages, images, and audio files on the site. Users can sort by a property by clicking on the appropriate column.

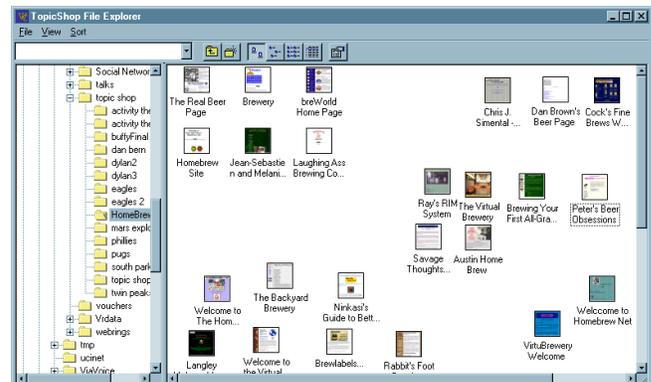


Figure 2: TopicShop Explorer (version 1), icons view. Each site is represented by a large thumbnail image and the site title. Users can organize sites by arranging them spatially, a technique especially useful in the early stages of exploration.

The collections of sites and site profile data used in TopicShop are obtained by running a webcrawler/analyzer. The crawler takes a user-specified set of seed sites as input, and follows links from the seeds to construct a graph of the seed sites, pages contained on these sites, and, optionally, sites determined to be related based on their textual and hyperlink connections to the seeds.

Site profiles are built by fetching a large number of pages from each site. Profiles contain content data, including the page title, an estimate of the page count, and a roster of audio files, movie files, and images; they also record links between sites in the collection. In addition, a thumbnail image of each site's root page is constructed.

The first goal of TopicShop is to help users evaluate and identify high quality sites. We sought to achieve this goal by providing site profile data and interface mechanisms for viewing and exploring the data. Making this data visible means that users no longer had to select sites to browse 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 may 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). 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 appropriate column. Users can “drill down” to investigate the profile data in detail, for example, to see a list of all the audio files on a site and all the other sites that it links to or that link to it. And users can browse the site in their default web browser just by double-clicking it.

The second goal is to make it easy for users to organize collections of sites for their own future use and to share with others. TopicShop let users organize sites both spatially (in the icons view) and by creating subfolders and moving resources into the subfolders. Thumbnail images also serve as effective memory aids to help users identify sites they already have visited.

PILOT STUDY

We needed a suitable yardstick of comparison for the user studies. For the task of exploring and evaluating web sites, we chose Yahoo, the most widely used search tool on the web. For the task of organizing web sites, we chose Netscape Communicator bookmarks, since bookmarks and the equivalents in other browsers are the primary means by which users organize web sites.

We chose two topics for the study: home brewing and the TV program “Buffy the Vampire Slayer” – each contained about 60 sites in their corresponding Yahoo category.

Design and Methodology

The experiment was a 2x2, between subjects design, with topic and user interface as factors. Sixteen members of our lab participated, resulting in four subjects in each condition.

The key metrics we wanted to measure were the quality of sites that users selected and the amount of effort required. To give a quality baseline, four experts for each topic were presented a list of the sites (in random order) on that topic. Experts had to browse each site, evaluate it based on its content and layout, and select the 20 “best” sites. For our

studies, we define “best” as a set of sites that collectively provide a useful and comprehensive overview for someone wanting to learn about the topic. During analysis, we used the “expert intersection”, the set of sites that all experts for each topic selected, as the yardstick for measuring the quality of sites selected by the subjects.

Subjects for a given topic were presented with the same set of sites to evaluate. The sites were obtained from the Yahoo category. 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 titles, thumbnail images, and profile data, as shown in Figures 1 and 2.

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. For the main task, subjects investigated the sites for their assigned topic by using the interface (TopicShop or Yahoo) and browsing sites. They were asked to choose the 15 best sites (as defined previously). Subjects were given 45 minutes to complete the task and were kept informed of the time, although they could take more time if they chose. 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 selected sites) between the two interfaces. And people do not spend unlimited amounts of time browsing, so we wanted to see whether users could find high-quality sites in a limited amount of time.

When subjects completed their task, they filled out a short questionnaire, and an informal interview was conducted.

Results

TopicShop subjects performed significantly better than did Yahoo subjects (see [3] for details). TopicShop subjects found over 80% more high-quality sites, i.e., sites in the expert intersection ($p < 0.05$) while browsing fewer sites and completing their task in less time. TopicShop's site profile data were the key to these results. The questionnaire and the informal interviews confirmed this; users emphasized the particular importance of the number of pages on a site and the number of other sites that link to it.

Lessons Learned

Despite these positive results, interviews and observations revealed some major shortcomings with TopicShop and thus important lessons for us. Subsequent reflection led to a major system redesign.

Like all artifacts, the TopicShop Explorer embodied claims about how users will conceive and carry out their tasks [7]. With its two separate windows for exploring site details and for organizing icons into groups, only one of which could be visible at a time, it embodied a claim that the tasks of evaluation and organization must be carried out separately.

Further, it assumed a single data set (the collection of all topic-relevant items), which could be manipulated in two ways (exploring site profiles or organizing by spatial grouping). The pilot study revealed problems with both implicit claims.

We found much evidence that users wanted to integrate work on the evaluation and organization tasks. First, they wanted to be able to organize items without losing sight of the detailed information contained in the site profiles. One subject commented:

I really want to organize the large icons, but don't want to lose the detailed information. Switching all the time is too painful, so I have to settle for the details view only.

Second, we realized that most items in a collection never would need to be organized, because users would not select them as worthy of further attention. Thus, rather than supporting a single collection, a better design would support two data sets. Users can evaluate the initial, machine-generated collection and select promising items. Organization will only be done for the selected items.

This also has implications for the nature of task integration. Users must be able to explore within groups they have created; for example, some users selected fairly large sets of similar sites, say ones that contained multimedia information, then wanted to keep only the best of these sites and throw the rest away. In order to do this, the interface should make it easy to sort within a user-defined group, e.g., to find multimedia sites with the most in-links or largest number of pages.

Third, the status of the user's task must be manifest. Most important, it had to be clear which items in the initial collection users had already evaluated and which they had not. Unevaluated items are a kind of agenda of pending work. Subject comments made this clear:

An indication of whether or not I visited the site would be useful. I can't tell what I've already seen.

It's hard to know what you've looked at and what you haven't...

Fourth, while the interface let users group sites by spatial organization or by creating explicit folders, users preferred the former technique. This is consistent with Nardi & Barreau [11], who found that users of graphical file systems preferred to organize their files by spatial organization. This is particularly useful early in the task while users are still discovering important distinctions among sites and explicit categories have not yet emerged. While the icons view supported spatial organization, the groups were not

first class objects. We wanted to explore spatial techniques to make it easy to create and manipulate groups.

Finally, site recall could be improved by including more graphical and textual information. Many subjects asked for the ability to annotate both individual sites and groups of sites. (Note that annotations also make collections more informative for others.) And other subjects asked for a larger thumbnail image to provide a better visual cue:

A larger thumbnail would be nice... It can be used to refresh your memory ... and would be more effective if it looked more like the site.

TOPICSHOP EXPLORER, VERSION 2

We created a new version of TopicShop (see Figure 3) based on the above lessons. We describe the new features and discuss how they respond to these lessons.

Two always visible, linked views support task integration and a cleaner definition of each task.

The site profile data and a work area for organizing sites are visible at all times. Items in the initial collection are displayed in the Site Profiles window, and the Work Area is initially empty (unlike Figure 3, which shows the results of a subject from the user study). As users evaluate items and find good ones, they select them simply by dragging them and dropping them in the work area. Since icons are created just for selected items, the work area is uncluttered, and provides a clear picture of the sites users care about.

"Piling" icons makes it easy to create first-class groups by spatial arrangement.

Groups can be formed in the work area by simply dragging icons. When a user positions one icon "close enough" to another, a group is automatically formed. (How close two icons must be before a pile is formed is a system parameter, set by default to occur just when their bounding boxes touch.) Each group is assigned a color. As the views are linked, both the group of icons in the work area and the features for sites in that group in the Site Profiles window are displayed using the color as a background. To help users better organize their groups, they can perform operations on piles (i.e. move, name/annotate, arrange, and select) as well as the normal operations on single sites.

Multi-level sorting is a useful operation that can be applied to a pile; it also illustrates how the linked views support task integration. In the site profiles view, users can reorder the sites based on primary and secondary sort keys. Users commonly sorted first by the groups they defined and then by some additional feature, such as in-links or number of pages. This lets users evaluate and compare sites within a single group. Figure 3 shows just such a sort.

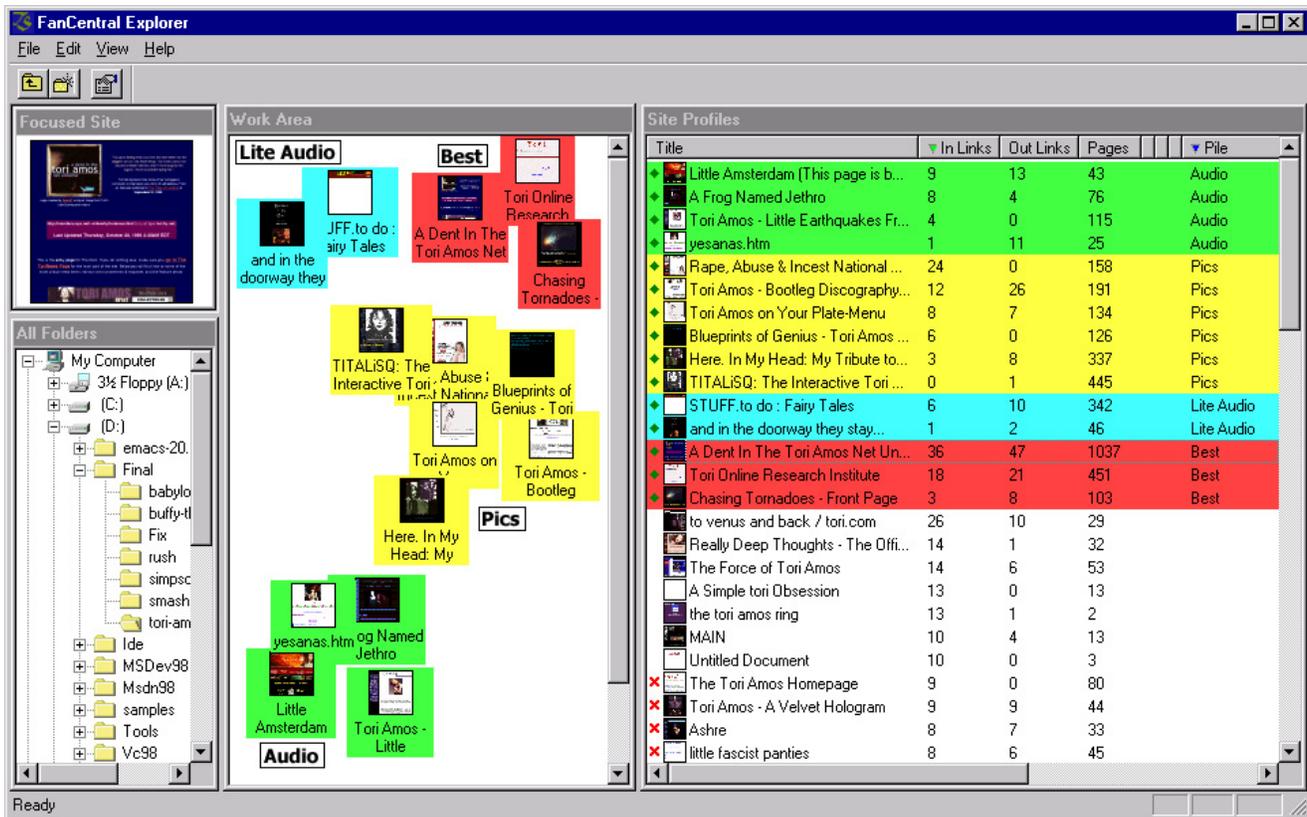


Figure 3: TopicShop Explorer (version 2)

Visual indicators make the task state apparent.

Any site included in the work area is marked with a green diamond in the site profile view and kept at the top for easy reference. Users can mark irrelevant or low-quality sites for deletion; this marks the sites with a red X and moves them to the bottom of the list. Thus, users quickly see which sites they have already processed (selected or deleted) and which need additional evaluation.

Annotations and large thumbnails support reuse and sharing.

The Focused Site window (upper left of Figure 3) displays a large thumbnail of the most recently clicked-on site. Users can create textual annotations for piles or individual sites in the work area. Annotations become visible as “pop ups” when the user lets the cursor linger over an object (pile or individual thumbnail) for a second or two.

USER STUDY

To test the advantages of the new design, we carried out a large empirical investigation of how web users evaluate and organize collections of web sites. In most respects, this study was similar to the pilot study. In describing the design and methodology, we highlight the differences.

Design and Methodology

We selected 5 popular entertainment topics, the television shows Babylon 5, Buffy The Vampire Slayer, and The Simpsons, and the musicians Tori Amos and the Smashing

Pumpkins. We again compared TopicShop to Yahoo+ Bookmarks, obtaining collections from Yahoo and applying our webcrawler to obtain site profiles and thumbnail images for use in TopicShop.

The experiment was a 2x5, between subjects design. We recruited 40 subjects from a local university. Subjects were assigned a topic and interface at random. The task still began with subjects selecting the 15 best sites. However, we also instructed subjects to organize their selected sites into groups and annotate the groups with descriptive labels. All subject actions were recorded and stored in log files.

We again used topic experts to rate site quality. We obtained 4 experts for The Simpsons, and 3 for all other topics. The expert task was a bit different, too. We decided it would be easier for them and more informative for us if experts rated site quality on a scale of 1 (worst) to 7 (best). A further change from the pilot study was due to the fact that the topic collections were much larger, ranging from about 90 to over 250 sites. Since we wanted to limit the number of sites experts rated to about 40, it was impossible for experts to rate all the sites. It wasn't even possible to rate all the sites that any subject selected. Instead, experts rated all the sites selected by multiple subjects and a sample of sites selected by one or no subjects.

Results

Our results showed the benefits of TopicShop in supporting both the evaluation and organization tasks and in enabling task integration. We present quantitative data and subject comments to illustrate these benefits.

Supporting The Evaluation Task

Once again, the key metrics we studied were the quality of sites that users selected and the amount of effort required.

TopicShop subjects again selected significantly higher quality sites than did Yahoo subjects. We generated the set of high-quality sites for each topic by sorting sites by their average expert score and selecting the top 15 (since subjects selected 15 sites). The quality of each subject's sites is measured by counting how many were among the top 15 expert sites. Table 1 shows the results for each topic and interface condition. On average, TopicShop subjects selected 76% more high quality sites ($p < 0.05$) – 7.4 of the expert sites vs. 4.5 for Yahoo subjects.

<i>Topic</i>	<i>TopicShop</i>	<i>Yahoo</i>	<i>% Increase</i>
Babylon 5	7.00	5.75	22
Buffy	7.25	3.50	107
Simpsons	6.50	5.25	24
Smash. Pumpkins	8.50	5.00	70
Tori Amos	7.75	3.00	158
Average	7.40	4.50	76

Table 1: Intersection between users selections and top 15 expert rated sites

Second, we wanted to be sure that users didn't gain quality by putting in more effort, so we measured the amount of time subjects spent on their task and the total number of sites they browsed. Again, TopicShop subjects had the advantage. They took about 72% of the time of Yahoo subjects (38 vs. 53 minutes), and they browsed about 67% as many sites (27 vs. 40).

In summary, TopicShop subjects selected higher quality sites, in less time and with less effort. We believe these benefits are due to TopicShop's site profile data. User comments and survey responses support this belief.

TopicShop subjects commented on the utility of the information they saw:

It presented me with lots of information very quickly. I could get a feel for what the site had to offer before visiting it, saving time to find the info that interested me. I got more than a site description, I got site facts.

The different sorting methods make it very easy for you to find what you're looking for.

And Yahoo subjects asked were near unanimous in asking for more information to judge sites:

[Show] some sort of popularity information to evaluate the sites.

[Show] something like an indication of how popular [the sites] were. Some rating of content.

Add some sort of ranking, that would be nice.

[Show] number of web pages, top 10 most visited.

List the type of audio or video offered on the multimedia pages.

I would add the approximate graphic level [i.e., number of images on a site] (so as to be able to judge the worthiness).

Subjects also were given a survey. It included a question asking them to rate the utility of the site profile features. Number of in-links was first, and number of pages was second (responses were similar in the pilot study). Interestingly, we have shown that both in-links and number of pages are good predictors of site quality [4]. Thus, subjects proved accurate in their utility judgements.

Supporting The Organization Task

The second part of the subjects' task was to organize their selected sites into groups and to name the groups. In the TopicShop condition, subjects grouped items by piling them together, while Yahoo subjects created folders in the Netscape Communicator Bookmarks window and placed items in the folders.

We defined a number of metrics to measure performance on the organization task. The metrics characterize the effort involved, the level of detail of the organization, and the amount of agreement between subjects on how sites should be grouped.

First, we examined the log files to compute how much time subjects spent on the organization task. TopicShop subjects spent 18% of their total time, while Yahoo subjects spent 36% of theirs. Since TopicShop subjects spent less time organizing sites, they were able to devote more time to evaluating and understanding the content of sites and selecting the good ones. Yet, even while taking less time, TopicShop users still created finer grained and more informative organizations, as we discuss next.

Second, we computed the number of groups subjects created. TopicShop subjects created 4 groups on average, and Yahoo subjects created 3. Thus, TopicShop subjects articulated the structure of the topic somewhat more. In addition, TopicShop subjects grouped nearly all of their selected sites (3% were left ungrouped), while Yahoo subjects left more ungrouped (15% were left).

Third, TopicShop subjects created more site annotations, thus making their collections more informative for their own use or for sharing with others. The experiment didn't require subjects to annotate sites. Yet 10 of 20 TopicShop subjects did so, annotating a total of 15% of their selected sites. Two Yahoo subjects annotated a total of four sites.

Fourth, TopicShop subjects tended to agree more about how sites should be grouped. This is a difficult issue to investigate; in general, it requires interpreting the semantics of groups. We computed a simpler metric: for each pair of subjects within a topic and interface condition, for each pair of sites that they both selected, did they group the sites together or not? If both subjects grouped the pair of sites together, or both grouped them separately, we counted this as agreement; otherwise, we counted it as disagreement. Table 2 summarizes the results. It shows that TopicShop subjects agreed 68% of the time on average, while Yahoo subjects agreed 43% of the time; thus, TopicShop subjects agreed 61% more.

Topic	Avg. agreement		Avg. % difference
	TopicShop	Yahoo	
Babylon 5	0.78	0.39	100
Buffy	0.59	0.44	34
Simpsons	0.78	0.36	116
Smash. Pumpkins	0.75	0.53	40
Tori Amos	0.48	0.41	17
Total	<i>0.68</i>	<i>0.43</i>	<i>61%</i>

Table 2: Agreement in grouping items

Taken cumulatively, the results show that TopicShop subjects appear to do a better job of organizing the items they select – they create more groups, they annotate more sites, and they agree in how they group items more of the time – and achieve these results in half the time Yahoo subjects devote to the task. We believe these results are because TopicShop makes grouping and annotation very easy, because of the rich information about sites that is available and remains visible while users organize sites. Subject comments support these beliefs.

- *TopicShop subjects found it easy to group sites.*

Piling web sites and annotating them makes grouping easy. You can easily see an overview of the organization.

Easily viewing category annotations and colored groups in the work area helps when attempting to determine what the important areas within a topic are.

- *Thumbnail images and textual annotations were effective memory aids for identifying sites and recalling their distinctive properties; TopicShop users commented on their utility, and Yahoo users expressed a desire for these types of functionality.*

Treating a site as a graphical object that can be dragged and dropped like anything else in your normal windows environment was much easier to conceptualize than treating sites as text links that required cutting, pasting, editing [TopicShop subject].

A thumbnail of the site ... would help the user who has been using several sites remember the site by looking at its thumbnail [Yahoo subject].

I used annotations to remind me about a site so I could tell the difference from the many other sites that I looked at [TopicShop subject].

Some way to take notes while surfing would be useful [Yahoo subject].

Relationship between Evaluation and Organization Tasks

We also studied the relationship between the evaluation and organization tasks. The TopicShop Explorer allows the tasks to be integrated, but doesn't force it. On the other hand, in the Yahoo/bookmarks condition, browsing sites and organizing bookmarks are separate tasks.

The log files contain data that let us quantify the relationship between tasks. Each user action is timestamped, and we know whether it was an evaluation or organization action. Evaluation actions included visiting a page in a web browser and sorting data in the Site Profiles Window. For TopicShop, organization actions included moving or annotating icons or groups in the Work Area. In the Yahoo/bookmarks condition, organization actions included creating a bookmarks folder, naming a folder, naming a bookmarked item, and placing an item in a folder.

We computed how many actions of each type occurred in each quartile of the task, i.e. how many occurred in the first 25% of the total time a subject spent on task, how many in the second 25%, etc. Table 3 shows the results for organizational actions. First, it shows how much more organizational work TopicShop users did – 533 actions vs. 172. (And recall they did this in half the time.) Second, as expected, TopicShop users integrated organization and evaluation to a much greater extent than did Yahoo users. They did about a quarter of their total organizational work in each of the first two quartiles, dipped slightly in the third quartile, then increased a bit in the final quartile. Yahoo users, on the other hand, did virtually no organizational work in the first quartile of their task, then ended by doing more than 50% in the last quartile. We should emphasize that TopicShop does not force task integration; rather, it enables it. And when users had the choice, they overwhelmingly preferred integration.

Quartile	TopicShop		Yahoo	
	# of actions	% of total	# of actions	% of total
Quartile 1	125	23%	2	1%
Quartile 2	138	26%	31	18%
Quartile 3	110	21%	50	29%
Quartile 4	160	30%	89	52%
Total	533		172	

Table 3: Interleaving tasks

We also can construct detailed timelines of user activity. Figure 4 shows such timelines for two Yahoo and two TopicShop subjects. They provide vivid illustrations of the overall results. TopicShop users interleaved the two tasks throughout the course of their work and performed many more organization actions. On the other hand, Yahoo users began by focusing exclusively on evaluation; then, toward the end of the task, they shifted to focus mostly on organization. And they did much less organization.

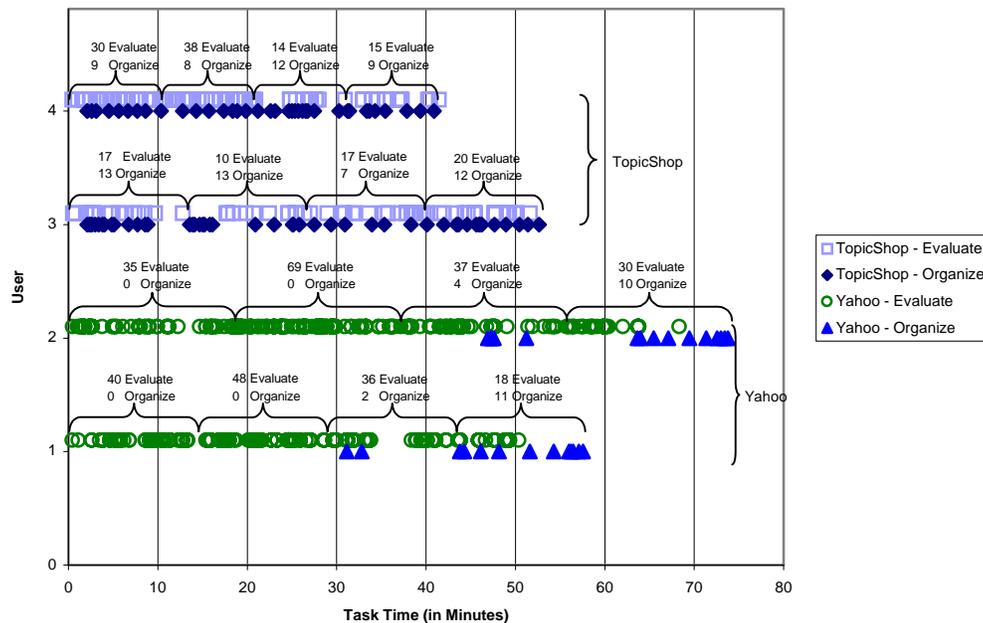


Figure 4: Timelines of user activity. TopicShop users do more organization actions and interleave organization with evaluation. Yahoo/bookmarks users did less organization, and did it at the end of their task.

Several comments showed that subjects appreciated the ability to integrate tasks and having the state of their task made visible.

- *Linked views helped users integrate the evaluation and organization tasks. In particular, they could evaluate within groups they created.*

Coloring was nice, because it gives me the ability to quickly SEE what was in what pile. Sorting within a pile was helpful for picking things out of each pile.

- *TopicShop made the state of the task apparent, allowing users to treat the initial collection of sites as an agenda of items to be processed.*

The graphics indicators let you quickly see what's left, because they show what you've already picked and what you didn't like.

SUMMARY

The popularity of the World Wide Web has made the problems of information retrieval and management more acute. More people than ever before face the problems of identifying relevant and high quality information and organizing information for their own use and for sharing with others.

The TopicShop systems improve people's ability to solve these problems. It provides information and interaction techniques that help people select the best sites from large collections of web sites. Two user studies have demonstrated that users can select better sites, more quickly and with less effort. It also offers 2D spatial arrangement techniques for creating groups of sites, and thumbnail images and annotations that enhance site recall and make the collections more informative. A study showed that users found it easy and fast to create groups and annotate their work. Finally, TopicShop makes it possible to

integrate the two major tasks of evaluating and organizing web sites. A user study showed that users preferred to integrate these two tasks when permitted by the interface.

ACKNOWLEDGEMENTS

We thank all the subjects who participated in our studies. We appreciate design and implementation work by Peter Ju and Philip Isenhour, and many helpful comments from Steve Whittaker, Bonnie Nardi, and Julia Hirschberg.

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. Amento, B. User Interfaces for Topic Management of Web Sites. Ph.D. Thesis, Department of Computer Science, Virginia Tech, forthcoming
3. Amento, B., Hill, W., Terveen, L., Hix, D., and Ju, P. An Empirical Evaluation of User Interfaces for Topic Management of Web Sites, in *Proceedings of CHI'99* (Pittsburgh PA, May 1999), ACM Press, 552-559.
4. Amento, B., Terveen, L., and Hill, W. Does "Authority" Mean Quality? Predicting Expert Quality Ratings of Web Documents, in *Proceedings of SIGIR'2000* (Athens Greece, July 2000), ACM Press.
5. Card, S.K., Robertson, G.C., and Mackinlay, J.D. The Information Visualizer, an Information Workspace, in *Proceedings of CHI'91* (New Orleans LA, April 1991), ACM Press, 181-188.
6. 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.
7. Carroll, J.M. and Rosson, M.B. Getting Around the Task-Artifact Cycle: How to Make Claims and Design By Scenario. *ACM Transactions on Information Systems* 10(2), 181-212, 1992.
8. Kleinberg, J.M. Authoritative Sources in a Hyperlinked Environment, in *Proceedings of 1998 ACM-SIAM Symposium on Discrete Algorithms*.
9. 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.
10. Mander, R., Salomon, G., and Wong, Y.Y. A 'Pile' Metaphor for Supporting Casual Organization of Information, in *Proceedings of CHI'92* (Monterey CA, May 1992), ACM Press, 627-634.
11. Nardi, B. and Barreau D. Finding and Reminding: File Organization from the Desktop. *ACM SIGCHI Bulletin*, 27, 3, July 1995.
12. Pirollo, 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.
13. Pitkow, J., and Pirollo, P. Life, Death, and Lawfulness on the Electronic Frontier, in *Proceedings of CHI'97* (Atlanta GA, March 1997), ACM Press, 383-390.
14. Robertson, G., Czerwinski, M., Larson, K., Robbins, D.C., Thiel, C., van Dantzich, M. Data Mountain: Using Spatial Memory for Document Management, in *Proceedings of UIST'98* (San Francisco CA, November 1998), ACM Press, 153-162.
15. Terveen, L., Hill, W., and Amento, B. Constructing, Organizing, and Visualizing Collections of Topically Related Web Resources. *ACM Transactions on Computer-Human Interaction*, 6,1 (March 1999), 67-94.