Footprints: History-Rich Tools for Information Foraging

Alan Wexelblat, Pattie Maes

MIT Media Lab

20 Ames St.

Cambridge, MA 02139 USA

{wex, maes}@media.mit.edu

ABSTRACT

Inspired by Hill and Hollan's original work [7], we have been developing a theory of interaction history and building tools to apply this theory to navigation in a complex information space. We have built a series of tools — map, paths, annotations and signposts — based on a physical-world navigation metaphor. These tools have been in use for over a year. Our user study involved a controlled browse task and showed that users were able to get the same amount of work done with significantly less effort.

Keywords

information navigation, information foraging, interaction history, Web browsing

INTRODUCTION

Digital information has no history. It comes to us devoid of the patina that forms on physical objects as they are used. In the physical world we make extensive use of these traces to guide our actions, to make choices, and to find things of importance or interest. We call these traces *interaction history*; that is, the records of the interactions of people and objects. Physical objects may be described as *history-rich* if they have associated with them historical traces that can be used by people in the current time. For example, if you are driving your car down an unfamiliar highway and approach a curve, you may notice that the guardrail has a number of black streaks on it. Realizing that these streaks have been formed from the "interaction" of the guardrail and the bumpers of other cars, you slow down. You are able to negotiate the curve safely because you can take advantage of the interaction history.

Interaction history is the difference between buying and borrowing a book. Conventional information retrieval theory would say they were the same object, given the same words, same pictures, same organization, etc. However, the borrowed book comes with additional information such as notes in the margins, highlights and underlines, and dog-eared pages. Even the physical object reflects its history: a book opens more easily to certain places once it has been used.

In Norman's terms [15], the history-rich object acquires new affordances and we can use these affordances to interact with the object in new ways. We make use of interaction history every day in dozens of different ways without conscious reflection; we think it is natural. In fact, car bumpers and guardrails are man-made artifacts that we have come to understand and read as a part of becoming adults in our society. The fact that we undergo such extensive learning suggests that interaction history is highly valuable. Our project, called "Footprints" by analogy with the footprints we leave in the world, is an attempt to understand what is valuable about interaction history in the physical world, and to find ways to capture history for use with digital information. We believe that the lack of interaction history information represents a significant loss. Work done by users to solve problems in information systems should leave traces. These traces should be accessible to future users who could take advantage of the work done in the past to make their own problem-solving easier.

For example, recently Maes found herself shopping for a new car on the Web. She visited a number of car manufacturer sites, car dealer sites, read reviews on-line, and looked at various independent reports and tests of a number of different vehicles. At the end of this process, she had not picked a particular car to buy — in fact, her list of possible choices was longer than when she began. But all the work done in this process was lost when she finished. If Wexelblat wanted to take advantage of her work, he might ask her, because he happens to know she has done this task, and she might remember some of what she had done and learned. But for anyone who did not know she had done this work there is no way to recover any of the things she found, nor to avoid any of the mistakes she made.

In the digital realm, problem-solvers must approach situations as though they were the first and only people ever to make use of the information. Maes' digital footprints are unavailable, so we all must become information foragers — in the sense of Pirolli and Card [17][18] — over and over again. The Footprints project tries to alleviate some of this kind of problem by allowing users to leave traces in the virtual environment, creating history-rich digital objects.

The term history-rich object and its association with records of the interaction of people and digital information derives from work by Will Hill and Jim Hollan [7][8]. We have taken their initial insight and expanded it into a theoretical framework that allows us to talk about a wide variety of history systems. The next section of this paper gives a basic introduction to the theoretical framework. We then describe the tools we have built to enable history-rich navigation in complex information spaces, particularly the World Wide Web. Finally, we describe our experiment in having people use these tools in a controlled task and discuss our ongoing work in expanding and improving the tools.

INTERACTION HISTORY FRAMEWORK

We have developed a framework for talking about interaction history. This framework presents six properties that characterize interaction history systems. The goal of the framework is to bound a space of all possible interaction history systems, and to give designers of such systems guidance as to what things are important in building history-rich interfaces. We use six properties to describe this space.

Property 1 — Proxemic versus Distemic

Urban planning and social anthropology use the words *proxemic* and *distemic* to describe the closeness relationship of people and spaces. We consider proximity to be a function of both the physical distance and the cognitive distance between the person and the space. A proxemic space is one that is felt by users to be transparent, in that the signs and structures can be easily understood. People feel close to, or part of, the space. Conversely, distemic spaces are opaque to users. Signals go unseen, usually because the people in the space lack the required background or knowledge to translate or comprehend what they experience. We feel "close" to our bedroom even when far away from it and experience a certain "distance" when we sleep in someone else's guest bedroom.

Interaction history systems may be more or less proxemic based on how well they relate to their users and how well they take advantage of users' past experiences and knowledge. For example, the personal computer desktop interface pioneered with the Xerox Star was intended to be proxemic in that it attempted to recreate a space with which the user would be familiar.

Property 2 — Active versus Passive

Most interaction history is passive; it is recorded and made available without conscious effort, usually as a by-product of everyday use of objects. Conversely, when we stop to think about leaving a record, we are creating an active history element. The active/passive distinction is concerned with the user's mental state and relationship to history-rich objects.

The most common example of this distinction is in Web browser software, e.g. Netscape Navigator. The "history" or "go to" list is passive history because it is recorded for the user as she browses; the "bookmarks" or "favorites" list is active history because the user must stop to think that she may want to return to this location in the future. The challenge for history-rich computer systems is to find ways to allow interaction history to be passively collected when necessary so that users are not constantly thrown out of the cognitive state necessary to getting their tasks done.

Property 3 — Rate/Form of Change

History moves forward, building as more interactions take place. This "accretion" process is how history builds up. However, interaction history does not only accrete, it also fades out. One of the challenges for history-rich interfaces is deciding how to deal with this accretion. Just as a complete video playback of a meeting is usually not as useful as a summary, the total accumulation of history must be summarized so that it can be observed and used quickly. A good real-world example of this are patient charts in hospitals. These charts are annotated and added to by many different personnel under different situations over time, yet a physician must be able to come into the room, pick up the chart, and understand essential facts of the patient's current state at a glance.

In the digital realm, Hill and Hollan's "Editwear" tool [7] used a modified scrollbar to show areas within a source file which had been more or less heavily modified. Dozens or hundreds of accesses were summarized by an unobtrusive thickening of the "thumb" component of the scrollbar.

Property 4 — Degree of Permeation

Permeation is the degree to which interaction history is a part of the history-rich objects. History may be inseperable from the object, as in a flight of worn stairs, or it may be completely separate, as in records of stolen art. In a history-rich interface, we must decide how closely to link the objects of interaction and the history information. Digital data will only retain that history information that we choose to keep; therefore, any record of this information must be captured and displayed by tools that we create explicitly for that purpose, or by display systems built into existing tools; for example, the mode-line modification to Emacs described above. The tools we have built to display interaction history information are described in the next section of this paper.

Property 5 — Personal versus Social

History can be intimate to a person: what have I done? Or it can be social: what has been done here? Many tools focus on personal histories; for example, bookmarks in Web browsers that allow users to revisit sites they have noted. Group histories, such as knowledge repositories and shared digital libraries are more rare but, we believe, more valuable because most problem-solving tasks are collaborative in nature. One of the primary benefits of interaction history is to give newcomers the benefit of work done in the past. In fact, the slogan for the Footprints project is:

We all benefit from experience,

preferably someone else s

Property 6 — Kind of Information

There are an infinite variety of kinds of interaction history information that can be captured. What kinds of information are important are, to a large degree, dependent on the task that the observer is trying to accomplish. Since we cannot possibly characterize all the kinds of information available, we focus on the uses to which interaction history might be put. We categorize the kind of information available loosely into *what, who, why,* and *how.*

Knowing *what* was done can be useful if users are *searching for value,* particularly among clutter, or if they are in need of reassurance. This is particularly helpful for novices who lack the kind of practice that helps them know what is reasonable to do with a given computer system. Knowing what was done can also *give guidance;* that is, the process of directing someone in a task or journey.

Knowing *who* has done something is important for reasons of *companionability* (doing things with friends), *sociability* (doing things with people who are similar to me), and for establishing *authority* and possibly *authenticity.*

Knowing *why* something was done can be important for reasons of *similarity of purpose.* I may care a great deal about something that was done by people with a goal similar to mine. A related reason is *goal discovery,* the process of starting off on one task and realizing that it relates to, or can be co-accomplished with, another task. Finally, knowing why something happened is crucial for *explanation and learning.*

Knowing *how* some bit of interaction history was hone can be important for issues of *naturalness.* For example, Microsoft Office's assistant has a "show me" mode in which it will show the user how to select the correct options from menus, how to fill in dialog boxes, and so forth.

APPLICATION TO THE WEB

To validate the theoretical framework, we built a series of tools applying interaction history to the problem of navigation in a complex information space. Earlier versions of these tools have been described in [20] [23]. The Footprints tools assume that people know what they want but may need help finding their way to the information and may need help understanding what they have found. Therefore, we do not use history to make recommendations. Instead we provide tools that use history information to **contextualize** Web pages that the user is seeing. This is information foraging: exploration combined with exploitation.

Our architecture is based on a proxy server (front end) and a database (back end). Both parts are written in Java and work on any platform with standard Web browsers. The front end controls the user interface tools, and records logs that are sent to the back end once per user session and incorporated into the database overnight. Interaction history information seen by users changes as they move from Web page to Web page, but the database itself changes only slowly. The one exception to this is user comments, as noted below.

Our tools are based on a metaphor of navigation — maps, paths and signposts — familiar from the physical world that we have implemented in the digital realm. There are, of course, many other tools that could have been implemented, but these both fit our metaphor and allowed us to explore interesting points in the space of possible interaction history systems described above.
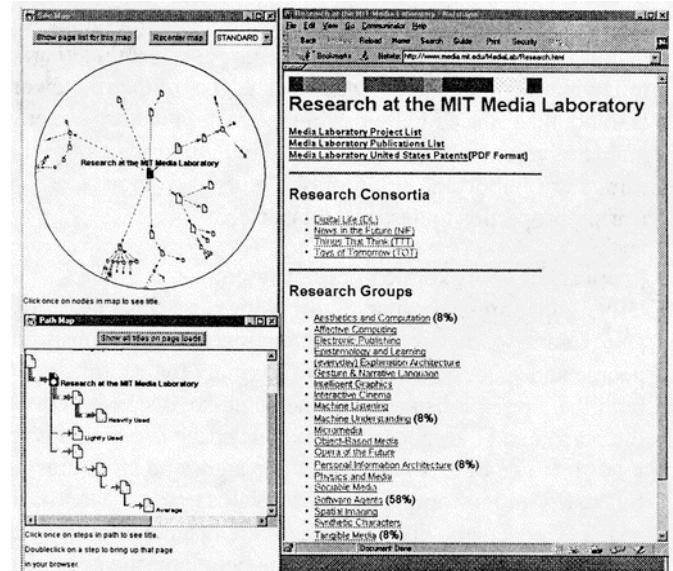

Figure 1 — Screen Shot of Footprints in Use

Each tool visualizes interaction history information in a different way, but they are active aids to navigation rather than static visualizations. The tools act in coordination. Selecting a document in one tool highlights it everywhere; focus is also coordinated. Tools also have some control buttons for manipulating document titles and helping users who get lost; these are explained below.

On start-up Footprints provides a control panel window that allows the user to show or hide each of the tools separately. Users can also shut down Footprints from the control panel. The Map and Path tools appear in separate windows alongside the Web browser. Figure 1 shows a screen shot of a user visiting the Media Lab Research Web page with all three tools turned on.

Over the course of the project, we have designed, tested and implemented several different versions of these tools. Our goals have always been to test our theories and to make systems that people will actually use. As a result, our designs have changed significantly over time, though our basic meta-

phor has stayed the same. All our tools use Web navigation transitions as their basic information — the "what kind" of information from our framework. Footprints validates that the pages it displays are accessible, so the rate of change of all Footprints is the rate of change of the Web itself. Footprints does not have a notion of user identity; all user data is anonymized and merged with the data of other users. This has the advantage of protecting users' privacy — no one can tell what Web sites you have visited — but it has the disadvantage of not allowing users to see each others' paths. This is a deliberate trade-off; other, equally valid trade-offs could be made but the focus of our research is on the interaction history itself and not on mechanisms for personal privacy.

The first tool is the map, pictured in Figure 2. This map shows the traffic through a Web site. Nodes are documents and links are transitions between them. Note that this is not all the documents and transitions, only the ones that people have actually visited or used. This is, typically, only a fraction of the actual site content. Additionally, we track all transitions made by the user, whether they come from selecting a link on the page, typing in a URL, selecting a bookmark, etc. The result of this is that links on the map often do not correspond directly to links embedded in the Web page.


Figure 2 — Footprints Site Map

In Footprints these user-created transitions are considered to be as important as the transitions (i.e. links) provided by Web-page designers. In some sense they are more important, since they reveal user's models of how information *should* be connected. As we described in earlier publications ([20] [23]) the patterns shown in the maps and paths are an externalization of users' mental models. This theory is reinforced by our experiment, described below.

The map visualization we use is derived from [13]. Users can drag the display in any direction to bring nodes from the edge towards the center. Individual nodes can be single-clicked to show their titles, or double-clicked to bring that document up in the Web browser. Titles might overlap, so the user may right-drag to rotate the map.

Popularity of documents is shown by shades of red — the hottest documents are in red, then shades of pink down to white (shown here as shades of grey). The document currently displayed in the browser is shown in black. Because users can get lost while viewing the map, there is a "Recenter map" button that redraws the map centered around the node in which the user expressed the most recent interest, either by single-clicking it or double-clicking it.

The titles of all nodes in the map can be viewed by clicking on the "Show Page List for this Map" button. Since many titles would clutter the display, the titles are shown in a separate window. This window presents the titles alphabetized, with the current document highlighted. Clicking on any title shows the title in the map view; double-clicking on a title brings up the document in the Web browser.

In the terms of our framework, the map view is social, combining data from all users. It is passive in that the data are added without requiring user intervention. It is distemic in that it requires users to learn new rules for interaction, and unpermeated in that the data are kept and displayed separately from the Web documents to which they refer.

The second tool is the path view, shown in Figure 3. If we think of the map as the high-level view, the path view is "lower" level in that it shows the user what paths have been followed by other people.
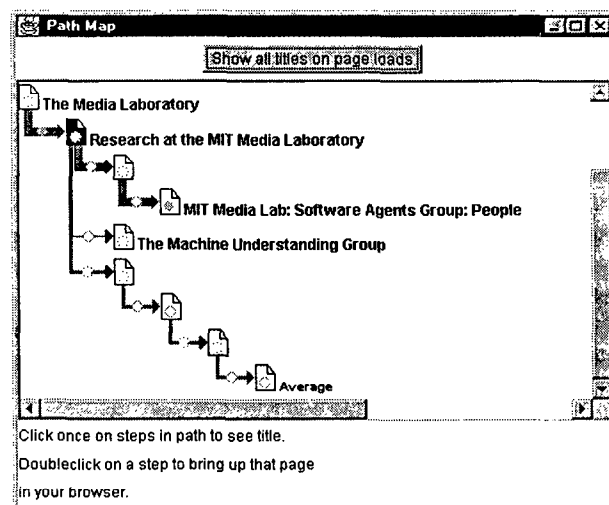

Figure 3 — Footprints Paths

Paths are coherent sequences of nodes followed by an individual. The map is much like a real-world map, with each document appearing once just as a city would appear once. The Path view is like a list of routes that go through these cities. A city appears at least once on each highway; likewise, a document in the Path view appears at least once on each path.

The number of paths formed this way is very large, of course, so we only show the paths that are relevant to (include) the current document. Note that paths with common starting points are merged, so users can see branching — forks in the road — more easily. For example, imagine that the following sets of paths are in the database.

$$A \to B \to C \to D$$
$$A \to B \to C \to E$$
$$A \to B \to D \to F$$

Then if the user was looking at page B, the representation would be:

```
A
 ↳ B
    ↦ C
       ↦ E
       ↳ D
    ↦ D
       ↳ F
```

The two representations of 'D' are not collapsed since they represent different paths through the space. The multiple representations of A and B are collapsed, however, since their positions in the path space are identical. If the user now selects page C, the representation would change to:

```
A
 ↳ B
    ↳ C
       ↦ E
       ↳ D
```

The A→B→D→F path is removed from the display because it is no longer relevant. Other paths containing page C (that do not also contain page B) would be shown if they were in the database. If the user backtracks, the previous path is redisplayed.

Paths also respond to single clicks — by showing titles — and double-clicks, by taking the user to the new document. There is also a button that allows users to see or hide all titles at once. Since the path view is arranged in a stair-step fashion, titles can be shown in the same window. Paths are coded for degree of use. The thickness of the line representing the path increases as the path is more heavily used; we also use a text string to give an approximate level of use. Paths are social, passive and unpermeated in the same way as maps; however, they are intermediate between distemic and proxemic because they take advantage of users' familiarity with tools such as outline listings and hierarchical file browsers.

The next tool is annotations, seen on the Web page in Figure 1. These are our only inactive aids. Annotations are marks — in our case numbers — inserted in the Web page that show what percentage of users have followed each link on the page. Footprints parses the HTML of the page in order to insert the annotations; therefore, we cannot annotate links that are "inside" imagemaps, applets, etc. Annotations are social and passive as with maps and paths. However, they are proxemic and permeated as they represent the "wear" directly in the page.

The final tool is signposts, or comments. These are the means by which users can enter feedback on the interaction history they have seen. Figure 3 shows a path view both with comments (filled circle, upper left) and without comments (open circles). Unlike other systems that only allow comments on pages, Footprints allows users to comment on both pages and paths. This can be useful, for example, in marking forks in the road. One of our beta test users provided an annotation that said "Go this way for software agents; go that way for artificial life." Users can click on the circles to bring up a simple text window. If comments already exist for that path, they are shown and the user has the option to add a comment. Clicking on an "Add Comment" button takes the user to a text input widget. Comments are social, active, proxemic and permeated.

Unlike the passive history information, comments are entered into the database immediately. Once the user clicks "OK" on the add comment window, the path view updates so that the circle is filled if it was not before. Clicking on the filled circle brings up the comments, including the new one, sorted so that the most recent comment is at the top. Older comments appear below. We do not delete comments; users can read the entire history and can converse or exchange ideas. With our small user population this has worked; as Footprints is used by larger groups, we will investigate whether any editorial policies are necessary.

USE OF THE TOOLS
Earlier versions of our tools have been in use for over a year. The first version was used only by alpha testers in-house, and at two Media Lab sponsor companies. The first public beta version of Footprints occurred in October 1997 via our Web site (footprints.media.mit.edu) and improved versions were released to sponsors for internal use. The third major release, described in this paper, occurred in August 1998. Each release has been used by a wider audience.

Although, as noted above, we have been guided both by a navigation artifact metaphor, and by a desire to explore the space of possible interaction history systems, we have also been fortunate to have continual feedback from our users. Many of the features found in the current implementation are a result of requests from users; for example, the recenter button.

EXPERIMENTAL RESULTS

We also performed a controlled experiment to evaluate both the subjective and objective usefulness of the tools. Subjects performed a timed (20 minute) browsing task, one group unaided and one group with the Footprints tools. Subjects were told that they had approximately $20,000 to spend on a car and were to find cars that might be interesting to them. They were encouraged to use their normal Web browsing patterns and tools. The second group had available the interaction history generated by the first group, and received a 5-minute instruction on how to use the Footprints tools based on a data set we created around the Media Lab Web site.

In testing a system that is designed to help people with an imprecise task such as browsing, it is hard to find useful measures. We settled on two objective and two subjective measures. Objective measures were the number of alternatives generated (that is, how many car makes/models they found), and the number of pages visited to generate those alternatives. Subjective measures were the users' sense of satisfaction and judgement of how easy the task was.

Twenty subjects participated in each of the two conditions. Subjects were volunteers given a token reward for participation. All subjects were expected to be familiar with Netscape Navigator before the experiment. Full details of the experimental conditions and evaluation can be found in [21]. Our pre-test hypotheses were that Footprints tools would increase the number of alternatives generated and reduce the number of pages visited. We also hoped that users would find it easier to find and understand relevant information, and would have a greater sense of satisfaction. The available interaction history information could lead people to a greater sense of having explored the problem.

The experiment partially supported our first hypothesis and gave a surprising result on our second. The number of alternatives generated by the two subject groups was not significantly different; however, the mean number of pages required to reach the same alternative level was significantly less for the Footprints group: 24.8 pages for the unaided group versus 18.75 pages for the Footprints group ($p < .05$).

In measuring the subjective responses, no significant differences were observed, with one exception. There was a significant interaction effect across conditions for those subjects who had, prior to the experiment, looked for car information on the Web, as shown in the ANOVA below. In this table, we test the interaction of user's previous activity (*looked* means that they have looked for car information on the Web before this experiment) with their response to a question about their satisfaction with the experience. Satisfaction was measured on a scale of 1-5, with 1 representing "Totally Satisfied" and 5 representing "Totally Dissatisfied." The table shows that while there is no significant effect for either effect considered separately, the two-way interaction of experimental condition (unaided vs. Footprints) and previous experience (looked vs. has not looked) was significant ($p < .01$).

TABLE 1. Interaction of Previous Experience with Satisfaction

Variation Source	Sum of Squares	DF	Mean Square	F	Sig of F
Main Effects	1.607	2	.803	1.5	.237
Expmt #	.207	1	.207	.385	.539
Looked	1.079	1	1.079	2.014	.164
2-Way Interacts					
Expmt # / Looked	4.872	1	4.872	9.093	.005

This was surprising as we had been assuming that interaction history models would help naive users; instead we seem to have found a situation in which past users' models are recognized and used by experienced browsers. In fact, naive users — in our case users who do not have experience with the domain — find themselves less satisfied when seeing these models. This most likely relates to our notion of proxemic/distemic, and reminds us that all naive Web users are not alike. Those who had a mental model of what car information on the Web was like found the Footprints representations much more proxemic and were able to make much better use of them. This conclusion was reinforced by other subject-reported experiences. In particular, subjects who rated their level of Web expertise lower reported having a harder time finding information that was relevant to their problem and less satisfaction with the solutions they found.

At this writing we are doing more detailed analyses of how subjects used the Footprints tools. Informal observation and post-test conversation suggested three patterns of use were in evidence. Some subjects simply took off in directions we had not seen before and so received little or no help from Footprints. This suggests that our test data set could be improved. Other subjects started off using Footprints information then went off in new directions since their tastes in vehicles differed from those of our first group of subjects. These subjects usually started with a popular site such as Yahoo! or Edmund's for which we had lots of history data from the first group. This variety of use patterns was expected; we cannot possibly cover all the possible car makes and models in which subjects might be interested.

The third group of subjects did not start out using the Footprints information. They had different search strategies. However, they ended up using Footprints information once their searches brought them near to popular car-related sites. At this point, the map proved particularly useful; as one subject put it: "As soon as I got there, the map had a bunch of alternatives and I used those."

RELATED WORK

Our theoretical work derives from two major influences. The first is ethnographic studies of how people work in teams in real situations, primarily research from Hutchins [11], Orr [16], and Suchman [19]. The second is urban studies, particularly the work of Lynch [14] and Brand [1]. From these we have developed our theory of how interaction history information can be used by people involved in their normal problem-solving tasks.

Hill and Hollan's original work [7] involved a series of tools, called "editware," "readware," and so on that were oriented toward helping people on a software development project keep track of which portions of the code and documentation were being the most heavily modified, most heavily read, etc.

Chalmers and his collaborators [3] have also been applying history (or activity-centrism as they call it) to tracing users' paths through the Web. Their tools are oriented towards providing recommendations for possible Web pages to visit, based on differences between the current user's paths and paths recorded in the system's history database.

The notion of paths through digital information and their use is at least as old as Bush's famous MEMEX essay [2]:

The owner of the memex... runs through an encyclopedia, finds an interesting but sketchy article, leaves it projected. Next, in a history, he finds another pertinent item, and ties the two together. Thus he goes, building a trail of many items. [H]is trails do not fade.

Hypertext systems have used map and path mechanisms for many years. However, these are typically top-down created artifacts put in the system by the designer for guidance or pedagogical purposes. Zellweger's "Scripted Documents" [24] are an excellent example of this. This notion is also being carried into the Web domain by projects such as CMU's *WebWatcher* [12], a tour-guide agent for the Web, and *Walden's Paths* [6], a K-12 educational application of scripted paths.

Some related work falls into the category of assisting social navigation. Dieberger [5] describes an enhanced MOO system, which keeps track of how many people use passages between rooms in the MOO and augments textual descriptions with information on how heavily used the passages appear. Dahlback, Hooks and others in the PERSONA project [4] have been exploring a number of different aspects of social and personal navigation, including the uses of artifacts in these processes and individual differences in the navigation process.

Other related work has been done in the area of community-created information sources. Hill and Terveen, particularly in their PHOAKS project [9][10], have been active in creating new techniques for mining existing information — on the Web and in Usenet newsgroups — for traces that can be collected and made available to future users. PHOAKS collects URLs that have been positively mentioned from postings and Frequently Asked Questions documents. These URLs are then provided as recommendations on a central server to people interested in the topic of the newsgroup from which they were extracted. Alexa (www.alexa.com) provides a real-time local Web-page recommendation system. They use history information as part of their input in determining what pages to recommend; however, it is unclear just how history is used or how it is integrated with the keyword matching that forms the basis for their recommendations.

CONCLUSIONS AND FUTURE WORK

We have built a set of tools to support undirected Web browsing. The tools are based around the concepts of interaction history and the notion that the work done by past users can be important to helping current users solve problems such as navigation in a complex information space. Our tools have been in use and available on the Web for over a year. The user community is small but growing. Our tools have been popular with Web information users and designers.

The experiment reported here showed that our tools are successful in two respects:

- they enable users to get the same work done with significantly less effort, and

- experienced users were able to recognize the information models left behind by other users and reported a significantly higher sense of satisfaction when working with these models.

More work remains to be done in testing the use of active history tools, as well as scaling up our user community. Applications of these ideas to areas such as electronic commerce and information management are also being investigated.

Finally, we set out to take something pervasive in the physical world, characterize it, and extract use from it for the digital realm. We have begun to show success in this endeavor; we have given history to digital information.

Thanks

prints tools was written by Felix Klock, Alex Lian, and James Matysczak. Jennifer Smith provided statistical help.

REFERENCES

1. Brand, Stuart. *How Buildings Learn: What happens after they're built*, Penguin Books, 1994.

2. Bush, Vannevar. "As We May Think," *Atlantic Monthly*, July 1945.

3. Chalmers, Matthew, K. Rodden & D. Brodbeck. "The Order of Things: Activity-Centred Information Access," *Proceedings of the 7th International World Wide Web Conference* (WWW7), 1998.

4. Dahlback, Nils (ed.) *Exploring Navigation: Towards a Framework for Design and Evaluation of Navigation in Electronic Spaces*, Swedish Institute of Computer Science TR98:01, 1998.

5. Dieberger, Andreas. "Supporting Social Navigation on the World-Wide Web," *International Journal of Human-Computer Studies*, Vol. 46, 1997.

6. Furuta et al. Hypertext Paths and the World-Wide Web: Experiences with Walden's Paths, *Hypertext'97 Proceedings*, ACM Press, 1997.

7. Hill, Hollan, Wroblewski & McCandles. "Edit Wear and Read Wear," *Proceedings of CHI'92 Conference on Human Factors in Computing Systems*, ACM Press, 1992.

8. Hill, Will & Jim Hollan. "History-Enriched Digital Objects," *Proceedings of Computers, Freedom, and Privacy (CFP'93)*, available from http://www.cpsr.org/dox/conferences/cfp93/hill-hollan.html

9. Hill, Stead, Rosenstein & Furnas. "Recommending and Evaluating Choices in a Virtual Community of Use," *Proceedings of CHI'95 Conference on Human Factors in Computing Systems*, ACM Press, 1995.

10. Hill, Will & Loren Terveen. "Using Frequency-of-Mention in Public Conversations for Social Filtering," *Proceedings of CSCW'96 Conference on Computer-Supported Cooperative Work*, ACM Press, 1996.

11. Hutchins, Edwin. *Cognition in the Wild*, MIT Press, 1995.

12. Joachims, Freitag & Mitchell. "WebWatcher: A Tour Guide for the World Wide Web," *Proceedings of IJCAI97*, 1997.

13. Lamping, Rao & Pirolli. "A Focus+Context Technique Based on Hyperbolic Geometry for Visualizing Large Hierarchies," *Proceedings of CHI'95 Conference on Human Factors in Computing Systems*, ACM Press, 1995

14. Lynch, Kevin. *What Time is this Place?* MIT Press, 1972.

15. Norman, Don. *The Psychology of Everyday Things*, Basic Books, 1988.

16. Orr, Julian. *Talking About Machines: An Ethnography of a Modern Job*, Cornell University Press, 1996.

17. Pirolli, Peter & Stuart Card. "Information Foraging in Information Access Environments," *Proceedings of ACM CHI'95 Conference on Human Factors in Computing Systems*, ACM Press, 1995.

18. Pirolli, Peter. "Computational Models of Information Scent-Following in a Very Large Browsable Text Collection," *Proceedings of ACM CHI 97 Conference on Human Factors in Computing Systems*, ACM Press, 1997.

19. Suchman, Lucy A. *Plans and Situated Actions*, Cambridge University Press, 1987.

20. Wexelblat, Alan. "Communities through Time: Using History for Social Navigation," in *Community Computing and Support Systems*, Toru Ishida (ed.), Lecture Notes in Computer Science, Volume 1519, Springer Verlag, 1998.

21. Wexelblat, Alan. *Footprints: Interaction History for Digital Objects*, Ph.D. Thesis, MIT Program in Media Arts & Sciences, 1999.

22. Wexelblat, Alan. "History-Based Tools for Navigation," *Proceedings of the 32nd Hawaii International Conference on Systems Sciences (HICSS-32)*, IEEE Computer Society Press, 1999.

23. Wexelblat, Alan & Pattie Maes. "Visualizing Histories for Web Browsing," *RIAO'97: Computer-Assisted Information Retrieval on the Internet*, Montreal, 1997.

24. Zellweger, Polle. "Scripted Documents: A Hypermedia Path Mechanism," *Hypertext'89 Proceedings*, ACM Press, 1989.