

How to Build Modeling Agents to Support Web Searchers

Paul P. Maglio and Rob Barrett

IBM Almaden Research Center, San Jose, CA, USA

Abstract. In this paper, we sketch a model of what people do when they search for information on the web. From a theoretical perspective, our interest lies in the cognitive processes and internal representations that are both used in and affected by the search for information. From a practical perspective, our aim is to provide personal support for information-searching and to effectively transfer knowledge gained by one person to another. Toward these ends, we first collected behavioral data from people searching for information on the web; we next analyzed these data to learn what the searchers were doing and thinking; and we then constructed specific web agents to support searching behaviors we identified.

1 Introduction

The World Wide Web connects tens of millions of people with hundreds of millions of pages of information. The web's explosive growth, its simple means for authoring, and its simple means of access have combined to make it a place many people now rely on to find information on almost any topic. Yet people trying to use the vast resources of the web to answer particular questions often face substantial problems in locating information. For example, one question we encountered recently was, "What percentage of calories from fat do French fries contain?" We observed a person search the web for this information for more than 30 minutes before giving up. What influences a searcher's success or failure? Such questions are only now starting to be explored for the web and for other online information sources (Marchionini, 1995).

In this paper, we consider the problem of building agents to facilitate a person's search for information on the web. From a theoretical perspective, our interest lies in the cognitive processes and internal representations that are both used in and affected by the search for information. From a practical perspective, our aim is to construct personal supports for information-searching and to enable transfer of knowledge gained by one person to another. In the end, we describe mechanisms for building user models to support web searching. We rely on data gathered from observing the behavior of experienced web users searching for specific information on the web. Our observations suggest that (a) individuals repeat the same search patterns, and that they recall their searches in terms of their standard patterns—almost regardless of what they actually did; and (b) people focus on key nodes when recalling their searches, and that these structure memory for the searches. To assist searchers, we built two personal web agents: the first, to identify repeated search patterns and to suggest similar patterns for new searches; and the second, to identify key nodes in finding a piece of information and to maintain personal trails in terms of these. The agents were constructed using the Web Browser Intelligence toolkit (WBI, pronounced "WEB-ee"; see Barrett et al., 1997a). WBI provides a way to tap into the data flowing between a web browser and the web, enabling construction of agents that monitor user behavior, model user behavior, modify what the user sees, and add new user functions to the web.

This paper is organized in four parts. First, we sketch our data collection method. Second, we present data collected from several people searching the web for specific information. Third, we discuss the construction of user modeling agents to assist web searchers, outline the WBI architecture for building web agents, and detail the specific agents we constructed. Finally, we summarize our results and discuss future directions.

2 How Data Were Collected

Seven experienced web users, five males and two females, each reporting more than two year years of almost daily web use, participated in this study. They were instructed to find the answers to three questions:

1. Does the University of Western Ontario offer a Master's degree in psychology?
2. What are three drugs currently being tested to help Alzheimer's patients?
3. In how many U.S. states was Ralph Nader on the ballot for president in 1996?

These questions were chosen to represent three kinds of searches. The first one has a reasonably well defined target location: web page about the psychology program at the University of Western Ontario. In this case, it is merely a matter of finding that location. The second question is less well defined; answers might be found in recent news, in medical information, or in Alzheimer's specific sites. Moreover, a full answer might require finding several sites. The third question could be answered using U. S. election results, state by state results, federal election commission information, or Nader-specific web sites.

Questions were presented one at a time. The participant was then asked how he or she was going to obtain the information from the web, that is, to verbally provide a rough plan of attack. Next, the participant used the web to try to find the information (for up to 15 minutes). Each participant returned the following day and was presented again with the same three questions in the same order. In this case, however, the task was to first verbally recall what he or she had done the previous day in searching for the answers to each question, and then to retrace the steps by performing the same search using the web. Note that participants were not told on the first day that recall of the details of their searches would be required on the second day.

We analyzed the data by comparing each search path generated by an individual participant on the first day with the one generated on second day. In addition, we examined how the verbal reports of search plans on the first day and the verbal recall of the searches on the second day corresponded to what was actually done.

3 How People Search the Web

Each of the seven participants completed at least one of the searches on both days, but only two completed all three searches. Of the possible 21 (7×3) searches, 15 were completed on both days. Only three of these 15 were repeated identically on the second day. In what follows, we consider mainly the 15 completed searches. We sketch our data and analyses to argue that (a) people conceptualize their searches in terms of standard routines for searching, and (b) they remember only key nodes when recalling their searches.

3.1 Searchers Rely on Routines

Our data suggest that individuals rely on personal routines when trying to find information. For instance, some participants routinely used a particular search engine, such as AltaVista, whereas others routinely used a particular hierarchical catalog, such as Yahoo! The point is not that our searchers merely preferred to use one approach over another; rather, we believe that they *conceptualized* their search tasks in terms of their favorite routines. We believe this because it often did not matter what was actually done on the first day, our searchers remembered searching *as if* their personal routines had been followed.

More precisely, our data show that (a) each individual has a standard pattern of search behavior; and (b) when an individual deviates from the standard pattern, he or she recalls the search as fitting the standard pattern. For example, participant T usually queried the AltaVista search engine to find likely starting points. She used AltaVista in all three of her searches the first day. For the Ralph Nader question, however, T began with the Yahoo! catalog instead. It turned out that Yahoo! did not provide easy access to good candidates, and so T wound up using AltaVista in the end anyway. The next day, when asked what steps she had followed for that search, she did not mention Yahoo!, and when retracing her steps, she did not go to Yahoo! (see Figure 1). Yet at other times, T was very concerned with following as many of the first day's dead ends on the second day as she could find. In this case, T's use of Yahoo! was forgotten, presumably because Yahoo! was not her standard pattern of behavior.

To take another example, participant D recalled one of his searches as fitting his standard routine when in fact it had not. For the Ralph Nader question, D carefully retraced his first 11 steps on the second day, including several that took him down a dead end path (see Figure 2). On the second day, however, D finished his search by using AltaVista—which he stated was his standard routine, and which he used for the other two searches—though he did not use AltaVista for this search on the first day. Even D's verbal recall of the first day's search was inaccurate:

I started at the Mercury News and I looked for election information, and it was a dead end because all of the links were not as current . . . they give results and not ballot information. So then I went to Yahoo! for election information . . . and then went to AltaVista to search for Ralph Nader and Green.

Thus, it was not merely that D could not find the same set of links from Yahoo! that he found the first day (see Figure 2), he remembered his search as fitting his standard routine. Unlike the case previously described for participant T, in which non-standard paths were omitted during recall, in this case participant D *added* his standard routine during recall.

All participants relied on their own standard routines, such as searching for starting points using AltaVista or using Yahoo! More importantly, on the second day, *five of the seven added a routine or deleted a non-routine pattern* in the ways we have just illustrated. Thus, because personal routines play an important role in how people remember their searches, we conjecture that such routines form the basis for how people conceptualize searching.

3.2 Searchers Rely on Waypoints

In addition to the use of personal routines, a second observation that emerged from the behavioral data is that participants recalled and relied on only a few of the sites they visited. For instance,

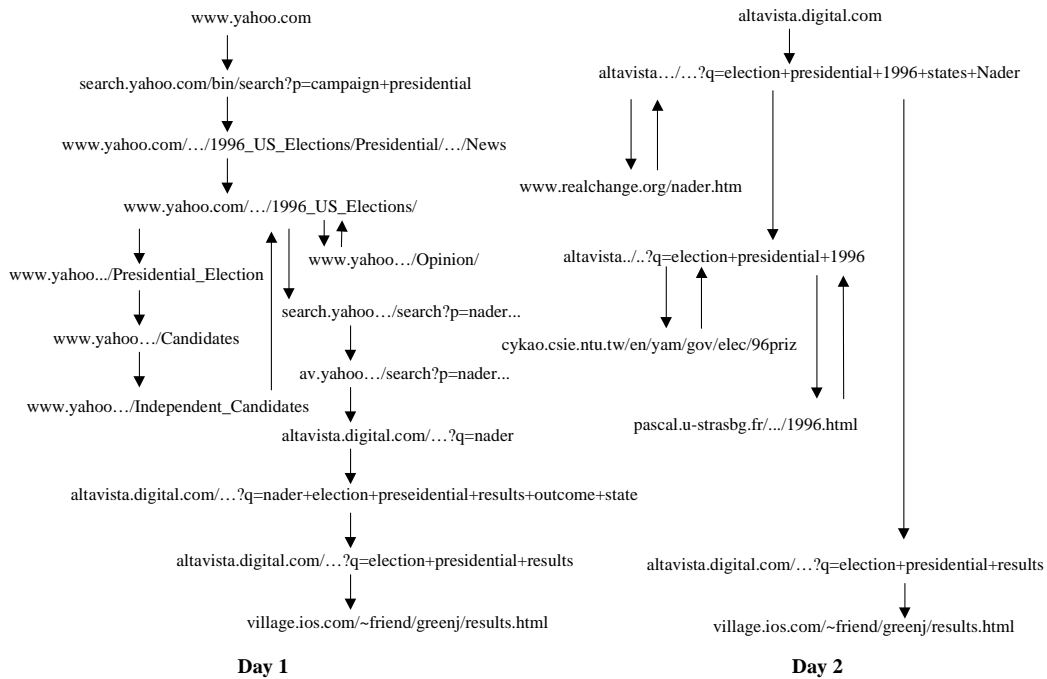


Figure 1. On the second day, participant T seemed to forget that she had used Yahoo! extensively on the first day. In this figure (and in those that follow), the nodes represent unique web pages, and the arcs represent transitions between pages.

some participants recalled mainly query terms to produce search results, one of which would often be recognized as one of the URLs that had lead to the goal. Others recalled a few of the specific URLs along the path they had followed, such as the University of Western Ontario home-page or the Alzheimer’s Association home page, and set out to find these. In short, participants remembered key nodes or *waypoints* that led to the goal. Formally, we define a waypoint as a node along a search path from which there is an unbroken sequence of links on successive pages that lead to the goal node (i.e., no URLs need to be typed in or explicitly recalled). Once traversed, waypoints are *recognized* as lying along the path to the goal—even if the same path is not followed to the goal in every case. For our participants, searching on the second day often meant finding waypoints encountered on the first day, rather than finding paths found on the first day.

Consider Figure 1 again. As shown, T’s search for the number of states in which Ralph Nader was on the U. S. presidential ballot ultimately relied on a specific AltaVista query: namely, one containing the keyword “results”. In fact, T explicitly mentioned this during verbal recall:

and I finally decided, oh yeah, I should just look under results or something, and then after, I went to a site that had the results, including how many states listed Ralph Nader on the ballot.

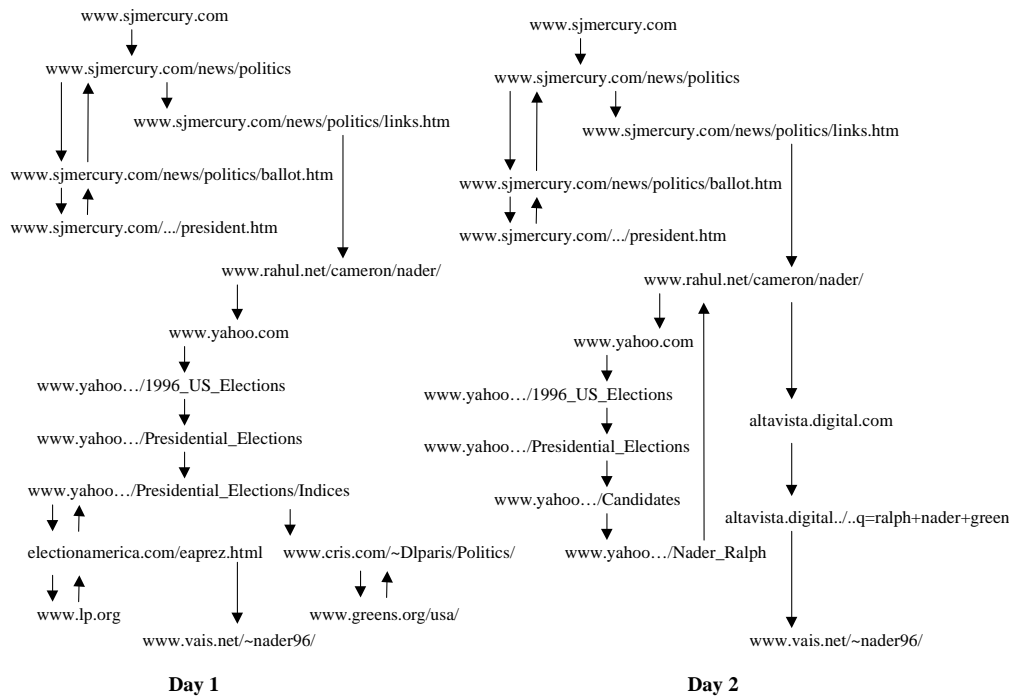


Figure 2. Participant D faithfully retraced the first 11 steps of this search, including dead ends, before switching to his usual routine (i.e., using AltaVista), even though he did not use AltaVista on the first day.

In this case, T's search depended on a waypoint created by querying AltaVista. Participant W used the same approach, recalling only the top level query used in his search for the answer to the same question:

I went to AltaVista and looked for the California Green Party and I wandered around and eventually I found it.

For these two participants, the query terms used as part of their standard routines were most important because the answer did not lie far from the query.

In another case, T's search depended on getting to a web page at the library of the University of Western Ontario (UWO). As shown in Figure 3, T retraced her steps to a web page at UWO's library, but not to the same page she had visited the previous day. This method of retracing steps by waypoints is also evident in how T verbally recalled the search:

...once I got there [AltaVista] ... I just looked for anything that had ... University of Western Ontario ... and one was the library and I thought ... there's probably a link back once I get to that point ... some sort of a directory ... so I got into the library and ... I in fact did find a link that took me back ... [to] this meta level for the University of Western Ontario, and then I went down to academic units ... and under there they have psychology, and under psychology they had programs or degrees ...

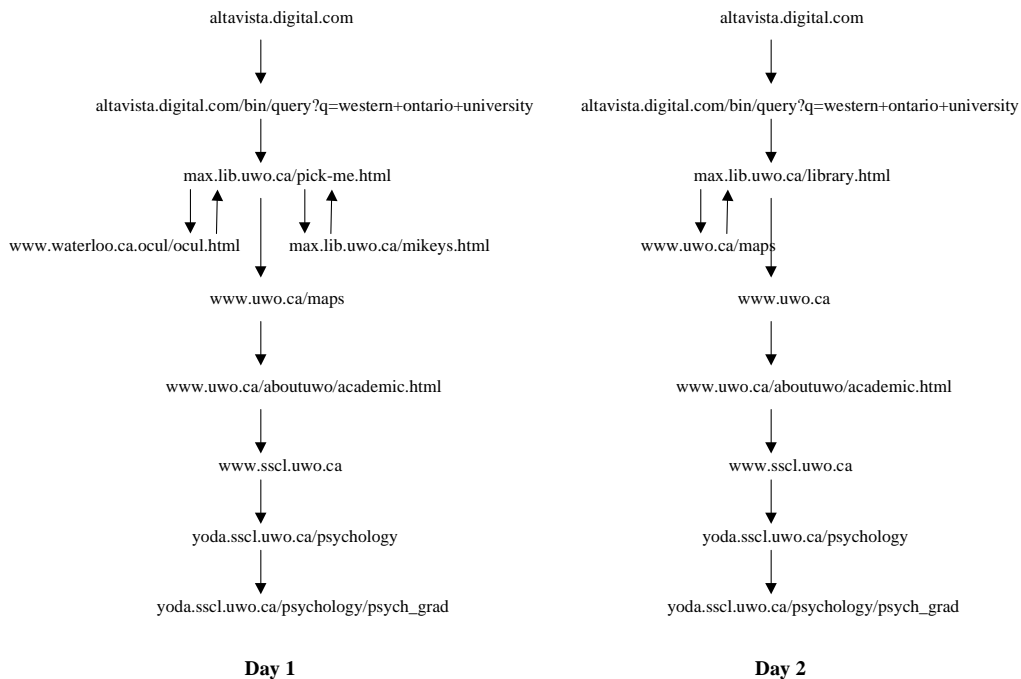


Figure 3. T's search for the psychology department at the University of Western Ontario depended on the UWO's library, but not on a specific location at the library.

This time, T recalled more than query terms, presumably because more than an AltaVista search was needed to find the answer. That is, AltaVista helped find a starting point at the University of Western Ontario, but the search proceeded independently from there.

Participant D had a similar experience finding the psychology department at UWO. As shown in Figure 4, D's waypoint was the homepage for the university, but his paths after finding this point were completely different on the two days. His verbal report suggests that finding UWO was the priority:

I'll go to AltaVista again. I searched on University of Western Ontario . . . and then take it from there.

Obviously, D believed the path would be clear once the UWO homepage was located. In a sense, the path followed the first day did not matter because some path could be found from the waypoint.

We found evidence in the behavioral data that each of the seven participants used waypoints in the ways just described. Specifically, for one of the completed searches, each participant either found a *similar* node on the second day from which the same path to the goal followed, or found the same node on the second day from which a *different* path to the goal followed. Overall, the

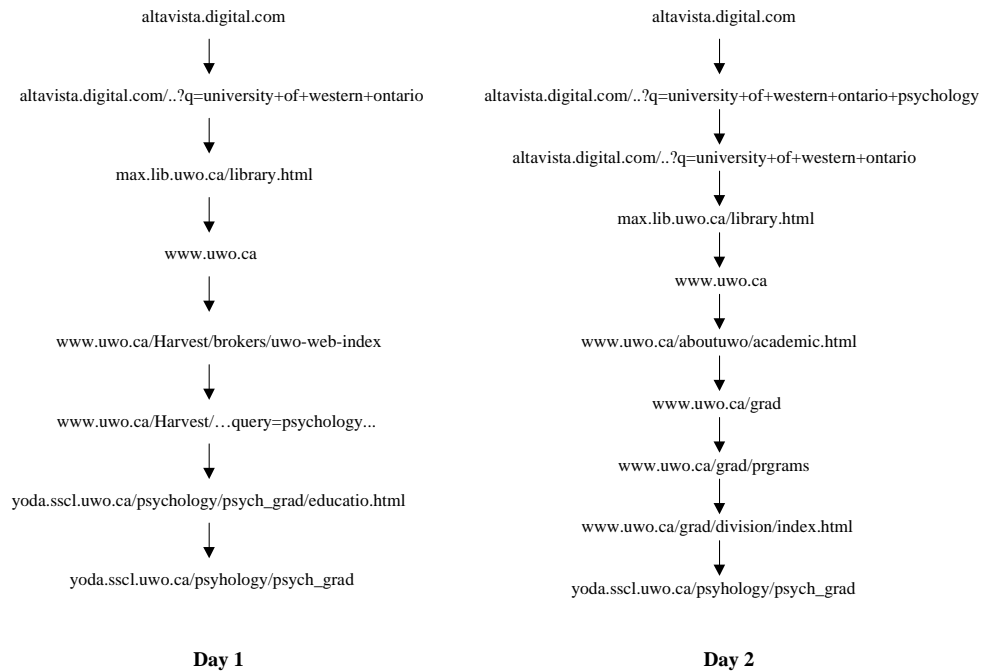


Figure 4. D’s search for the psychology department at the University of Western Ontario depended on the UWO’s homepage, but took different paths from there on the two days.

data suggest that memory for prior searches was structured around important nodes that led to the goal.

In summary, when searching the web, people use standard routines to find key points that are close to the desired information. To make this case, we have shown that our searchers: (a) generated only a few of the important nodes they visited when recalling their searches (both verbally and behaviorally); and (b) relied on personal routines both to find important nodes and to find specific information starting from these nodes. It follows that searchers do not fully *plan* in advance, but rely instead on heuristics (routines) and local context (waypoints) to find information (cf. Georgeff and Lansky, 1987). Put another way, searchers interleave directed and structured behavior—searching—with opportunistic and unstructured behavior—browsing—to find information (Bates, 1989; Marchionini, 1995).

3.3 Related Studies of Web Use

Our findings accord with other recent studies of how people use the web. Catledge and Pitkow (1995) and Tauscher and Greenberg (1996) analyzed several weeks worth of normal web usage gathered from dozens of individuals. Both studies found that web users do not often traverse the

same long sequence of nodes more than once. As we have shown in the present study, participants were unable to recall—and therefore repeat—specific sequences of URLs even when explicitly asked to do so. Tauscher and Greenberg also found that web users have about a 60% chance of revisiting web pages they have previously visited; although sequences of URLs were rarely repeated, specific URLs were often repeated. As we have shown, participants were more apt to refind waypoints than to refind sequences. Thus, although we focused on details of specific search behaviors observed in a few individuals, the patterns we found are consistent with a large body of quantitative data.

4 Web Agents Based on User Models

Computer agents are programs that *collaborate* with human users, mimicking behavior expected of a human assistant or advisor (Maes, 1994). To collaborate on a web searching task, an assistant or advisor might suggest links to follow, remind a user of paths taken before, or inform a user of related paths others have taken. Effective collaboration requires that agent and user share the same understanding of the domain (Rich and Sidner, 1997). Our behavioral data suggest that people think of the web in terms of routines and waypoints. Therefore, to best assist or advise web users, agents ought to model users' interactions with the web in terms of routines and waypoints. In this way, a particular agent would learn about a particular user's routines and waypoints, and then communicate with the user in terms of these.

A specific user model can be based on explicit conversations with an agent (Etzioni and Weld, 1994), or as a mixture of explicit interactions with an agent and interactions with the web itself (Armstrong et al., 1995; Thomas and Fischer, 1996). By contrast, our agents were designed to operate in a style we believe is better suited to assisting users in their *normal use* of the web. That is, rather than radically changing the web experience, our aim has been to incorporate new functions without substantially altering a user's ordinary interaction (Norman, 1994). Thus, to build a user model, our web agents monitor user-web interactions implicitly (Selker, 1994), and to communicate with the user, our agents annotate retrieved web pages and respond to specific user requests (see also Barrett et al., 1997a).

In particular, we built two web agents. One identifies repeated link-following patterns and suggests similar patterns in similar contexts. The other identifies waypoints used in finding information to help visualize user trails. These agents were constructed using the WBI toolkit.

4.1 WBI Agent Architecture

WBI provides an architecture that taps into the communication stream between a user's web browser and the web itself. Using WBI, we can attach agents to this stream to observe the data flowing along the stream, or to alter the data as it flows past. Such agents can learn about the user, influence what the user sees by marking-up pages before passing them on, and provide entirely new functions through the web browser. WBI is implemented as a proxy server that can run either on the user's client workstation or on a server that many users access (and can be downloaded for Windows 95/NT or OS/2, see IBM).

In WBI, agents are defined by a set of monitors, editors, generators, and rules. A monitor is given a copy of a web request and a copy of the resulting page so that it can record user actions.

For example, a *history* agent would use monitors to record the series of URLs visited and the contents of the pages viewed. An editor is allowed to modify the resulting page before it is returned to the browser. For example, an editor might be used to add extra buttons for accessing WBI functions. A generator is used to handle web requests. For example, a generator can be used to search a user's history and return a list of pages that contain a given keyword. Rules define which specific monitors, editors, and generators should be instantiated for a given request. For example, certain editors might only want to modify HTML documents, and certain monitors might only want to observe requests made to servers in a certain domain. A more complete description of the WBI architecture can be found elsewhere (Barrett et al., 1997a, 1997b).

4.2 Shortcut Agent Identifies Repeated Patterns

As we have argued, people follow standard routines when searching for information. To support this behavior, we constructed a *shortcut* agent to extract repeated patterns from a user's history of interactions with the web. For instance, suppose a user routinely searches for documentation on the Java programming language by going to the Java home page, then to programming information, then to the reference manual. The shortcut agent identifies this pattern of repeated actions in the user's history and adds a link on the Java home page to the Java reference manual page. Thus, the user can jump directly from the Java home page to the reference manual, skipping intermediate pages (see Figure 5).

The shortcut agent relies both on a monitor that records an ordered list of all URLs visited, and on an editor to add links to web pages. When a user requests a URL, the editor scans the visited list for previous occurrences of it. All URLs that were visited within a certain neighborhood of each previous occurrence of the current URL (e.g., within 5 steps) are collected and sorted by frequency. Any such neighboring URLs that were visited more than a certain number of times (e.g., 4) and more than a certain proportion of the time (e.g., 25%) are added to the top of the page by the shortcut editor.

4.3 Waypoint Agent Identifies Key Nodes

It seems clear users would benefit from being able to see their histories of interaction with the web—so that they can return to previously found locations, or so that they can share their histories with others, to list two examples. Our behavioral data suggest that people rely on waypoints rather than on sequences of URLs to reconstruct their searches. This implies that displaying a complete history of URLs is not only redundant but potentially confusing for users (see also Tauscher and Greenberg, 1996, and Cockburn and Jones, 1996). We built the *waypoint* agent to extract likely waypoints from a user's history of interactions with the web to enable appropriate visualization of user trails. Recall that a waypoint is a node from which there is an unbroken sequence of links on successive pages that lead to some goal. As a first pass at this, our waypoint agent partitions the user's history of interactions into segments that contain no backtracking and that are connected by a sequence of links. The user can request that the waypoint agent display his or her interaction history in terms of such node groups (see Figure 6).

The waypoint agent relies on a monitor that records each URL requested as well as hyperlinks on all pages retrieved. In addition, the waypoint agent uses a generator to construct a web page to display the user's history of interactions. To do this, the generator constructs a disconnected

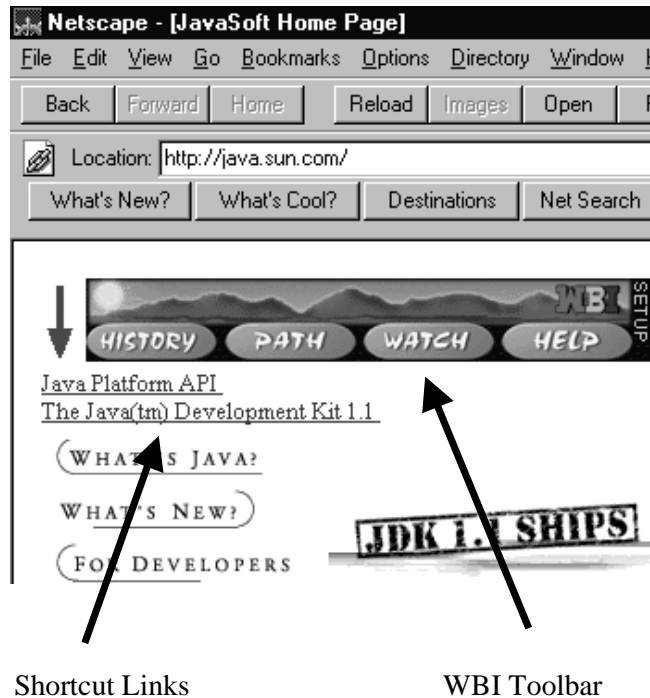


Figure 5. The shortcut agent has added two links to the top of the JavaSoft homepage, one to the Java APIs, the other to the development kit. These links were added because the user routinely follows a path from here to the development kit page and then to the APIs page.

graph of the user's path through the web using the link information that the monitor has recorded, and removes backtracking from these each of the connected subgraphs. Discontinuities arise when there are no links connecting successive pages. Finally, this generator is invoked when the user visits a specific URL (such as `http://wbi/mytrails`), and each of the subgraphs is displayed as an ordered list of titles with duplicates removed. Note that in this preliminary version, the waypoint agent does not distinguish between important and unimportant URLs in the clusters it generates, and so does it not in fact produce specific waypoints. The clusters it produces, however, distill a user's traversal of the web into self-contained regions.

4.4 Related Work

In general, WBI's waypoint and shortcut agents are unique in that they are based on an analysis of *how people search the web*. For instance, Thomas and Fischer (1996) describe modeling agents that assist web searchers. One difference between Thomas and Fischer's BASAR system and WBI is that users explicitly configure and interact with BASAR's agents through a special interface, whereas users interact with WBI's modeling agents by simply searching the

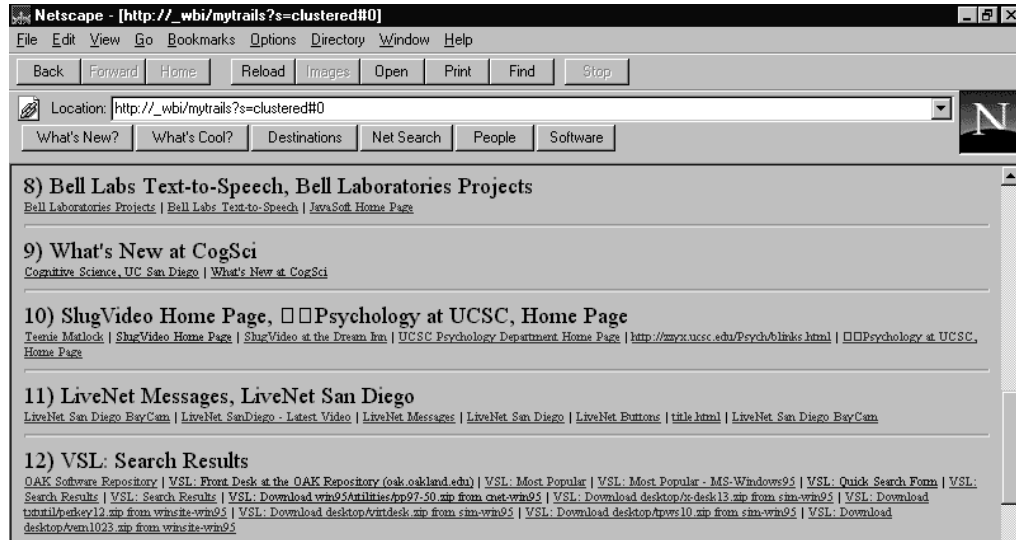


Figure 6. The waypoint agent displays a user's history of interactions with the web in term of locally connected clusters of URLs.

web. Similarly, Armstrong et al.'s (1995) WebWatcher is a web search assistant that requires explicit user feedback to incrementally learn the key words that distinguish one topic from another. WebWatcher uses a statistical machine learning algorithm to try to partition the universe of key words. Lieberman's (1995) Letizia uses simple heuristics, such as persistence-of-interest and limited lookahead, to determine which links a user might want to follow. WBI's modeling agents differ from all of these is that WBI's agents are empirically motivated.

5 Conclusion

Using data gathered by observing experienced web users searching for specific information on the web, we argued that (a) individuals repeat the same search patterns and conceptualize their searches in terms of their standard patterns, and (b) people focus on key nodes or waypoints when recalling their searches. On the basis of these findings, we built two personal web agents to assist searchers: the shortcut agent identifies repeated search patterns and suggests similar patterns for new searches, and the waypoint agent identifies and displays nodes that go together in a user's history. Because these agents communicate with a user in terms close to the user's model of the domain, we believe these agents are well-adapted for user-web collaboration.

One direction for future work is to empirically determine whether our modeling agents make web searching easier. Do people repeat searches faster using these agents? Do they locate in-

formation more quickly in the first place? Informal observations suggest that shortcuts in fact enable users to skip steps—and therefore to save time and cut down on confusion—when repeating standard patterns of action. A second direction is to define a method to find the most important nodes in the clusters the waypoint agent generates. Do individual waypoints give users a better view of their past than clusters of URLs? In any case, final direction for future work is to determine whether a single user's waypoint history can be easily used by others. That is, do waypoints capture the most important structure of a user's trails?

References

- AltaVista. AltaVista Search. Available as <http://altavista.digital.com>.
- Armstrong, R., Freitag, D., Joachims, T., and Mitchell, T. (1995). WebWatcher: A learning apprentice for the World Wide Web. In *Proceedings of the 1995 AAAI Spring Symposium on Information Gathering from Heterogeneous, Distributed Environments*.
- Barrett, R., Maglio, P. P., and Kellem, D. C. (1997a). How to personalize the web. In *Proceedings of the Conference on Human Factors in Computer Systems (CHI '97)*. New York, NY: ACM Press.
- Barrett, R., Maglio, P. P., and Kellem, D. C. (1997b). WBI: A confederation of agents that personalize the web. In *Proceedings of the First International Conference on Autonomous Agents*, 496–499. New York, NY: ACM Press.
- Bates, M. J. (1989). The design of browsing and berrypicking techniques for the on-line search interface. *Online Review* 13:407–431.
- Catledge, L., and Pitkow, J. (1995). Characterizing browsing in the World-Wide Web. In *Proceedings of the Third International World Wide Web Conference*.
- Cockburn, A., and Jones, S. (1996). Which way now? Analysing and easing inadequacies in WWW navigation. *International Journal of Human-Computer Studies* 45:105–129.
- Etzioni, O., and Weld, D. (1994). A softbot-based interface to the internet. *Communications of the ACM* 37(7):72–76.
- Georgeff, M., and Lansky, A. (1987). Reactive reasoning and planning. In *Proceedings of the Sixth National Conference on Artificial Intelligence*, 677–682.
- IBM. Web Browser Intelligence. Available as <http://www.raleigh.ibm.com/wbi/wbisoft.htm>.
- Lieberman, H. (1995). Letizia: An agent that assists web browsing. In *International Joint Conference on Artificial Intelligence*, 924–929.
- Maes, P. (1994). Agents that reduce work and information overload. *Communications of the ACM* 37(7):31–40.
- Marchionini, G. (1995). *Information Seeking in Electronic Environments*. Cambridge, England: Cambridge University Press.
- Norman, D. A. (1994). How might people interact with agents. *Communications of the ACM* 37(7):68–71.
- Rich, C., and Sidner, C. L. (1997). COLLAGEN: When agents collaborate with people. In *Proceedings of the First International Conference on Autonomous Agents*, 284–291. New York: ACM Press.
- Selker, T. (1994). COACH: A teaching agent that learns. *Communications of the ACM* 37(7):92–99.
- Tauscher, L., and Greenberg, S. (1996). How people revisit web pages: Empirical findings and implications for the design of history systems. Unpublished manuscript. University of Calgary, Department of Computer Science.
- Thomas, C. G., and Fischer, G. (1996). Using agents to improve the usability and usefulness of the World-Wide Web. In *Fifth International Conference on User Modeling*, 5–12.
- Yahoo! Yahoo! Available as <http://www.yahoo.com>.