

On the Trail of Information Searchers

Paul P. Maglio Rob Barrett

IBM Almaden Research Center
650 Harry Road, NWED-B2
San Jose, CA 95120

{pmaglio,barrett}@almaden.ibm.com
Phone: 408-927-2857 FAX: 408-927-4322

Abstract

In this paper, we sketch a model of how people search for information on the World Wide Web. Our interest lies in the cognitive properties and internal representations used in the search for information. We first collected behavioral data from individuals searching for answers to specific questions on the web, and we then analyzed these data to learn what searchers were doing and thinking. One finding was that individuals focus on key nodes when recalling their searches, and that these key nodes help structure memory. A second finding was that people tend to use the same search patterns over and over, and that they recall their searches in terms of their standard patterns—regardless of what they actually did. Overall, our results suggest that people form cognitive maps of web space in much the same way that they form cognitive maps of physical space.

Introduction

When trying to use the vast resources of the World Wide Web to answer specific questions, people often face problems locating information. One possible reason is that the web is organized idiosyncratically rather than uniformly, and so information seekers cannot rely on a single strategy or tactic. On the web, there might not be any one correct query or path for finding a particular piece of information—by its very nature, web searching requires browsing. In any event, idiosyncratic organization makes the web different from information-seeking tasks that have been previously studied (e.g., Marchionini, 1995). But to what extent is searching the web like searching a highly structured information resource such as an online library or encyclopedia? Can we apply what is known about information-seeking behavior to the web? What influences a searcher's success or failure? Or more basically, how do people conceive of their searches? For instance, if people navigate hypertexts like they navigate in the physical world, what role do cognitive maps and spatial abilities play?

We do not attempt to answer all these questions here. Our work is only an initial step in this direction. We start with the assumption that people conceive of web space in much the same way they conceive of physical space. There are several reasons to believe that spatial abilities and concepts play a role in the way people understand and use online information. First, spatial ability has been shown to influence the use of both hierarchically organized computational environments such as file systems (Vicente & Williges, 1988), and

hypertextually organized online help systems (Campagnoni & Ehrlich, 1989). Recently, Dahlback, Hook and Sjolinder (1996) found that for people seeking answers in hypertext-based online documentation, spatial ability for manipulating mental images—rather than spatial ability for navigating in the real world—correlates with ability to effectively find answers online. Second, the familiar metaphor of information seekers *navigating* an information space (e.g., Nielsen, 1990) provides prima facie evidence that people conceptualize hypertext, such as the web, similar to how they conceptualize physical space. Moreover, even naive web users talk about using the web as if it is a kind of landscape over which they move to obtain information (Matlock & Maglio, 1996, 1997).

In this paper, we sketch an account of the cognitive processes and representations people use when searching for information on the web. We base our account on *cognitive maps* of information space. Cognitive maps of physical space are generally believed (e.g., Anderson, 1980) to incorporate landmark knowledge, route knowledge (procedures for how to get from one landmark to another), and survey knowledge (map-like representations). By observing the behavior of experienced web users looking for specific information, we found: (a) that people focus on key nodes (landmarks) when recalling searches, which suggests that such key nodes structure memory for searches; and (b) that they repeat search patterns (routes) and recall searches in terms of standard patterns, almost regardless of what was actually done.

This paper is organized in three parts. First, we detail the method we used to gather data. Next, we describe the results. Finally, we discuss implications of our findings.

How Data Were Collected

To learn how people search for information, we watched experienced users search the World Wide Web for answers to specific questions. To identify key cognitive aspects of their activities, we first asked them about their plans, and then tracked their behavior while they searched. Then a day later they were asked to recall the steps they had taken in each of their searches the previous day, and finally to retrace their steps. Participants were not warned on the first day that recall would be required on the second day. This method enabled us both to chart behavior to uncover search tactics (using the behavioral traces) and to extract some of the structure of their internal representations (using the recall data). In this way,

we hoped to discover what people do when they search for information.

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.

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. Three of these 15 were repeated identically on the second day. In what follows, we consider only the 15 completed searches. We sketch our data and analyses to argue that (a) people remember only key nodes when recalling their searches, and (b) they conceptualize their searches in terms of standard routines.

Searchers Rely on Anchor Points

Our data showed that participants recalled only a few of the sites they visited. Specifically, they remembered key nodes that led to the target information. We call these nodes *anchor points* by analogy to the notion of anchor points in the cognitive map literature (Couclelis, Golledge, Gale & Tobler, 1987; Ferguson & Hegarty, 1994). A cognitive map is an internal representation of physical space that can be used to make inferences about the space. When an individual constructs a cognitive map by traversing a physical space, certain relationships among landmarks are represented more prominently than others. In particular, such a map is organized in terms of regions, each of which is identified with a special landmark called an anchor point. Judgments about distance and orientation of landmarks within a region tends to be made relative to the region's anchor point, and judgments of landmarks between regions tends to be made in terms of anchor points rather than in terms of the landmarks themselves.

For present purposes, we define an anchor point 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, anchor points 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 anchors encountered on the first day, rather than finding paths found on the first day.

Consider Figure 1. 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 ... I should just look under results ... and then after, I went to a site that had the results, including how many states listed Ralph Nader on the ballot.

In this case, T's search depended on an anchor point 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, AltaVista queries resulted in anchor points that defined information regions that contained the target information.

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 2, 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 anchors 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

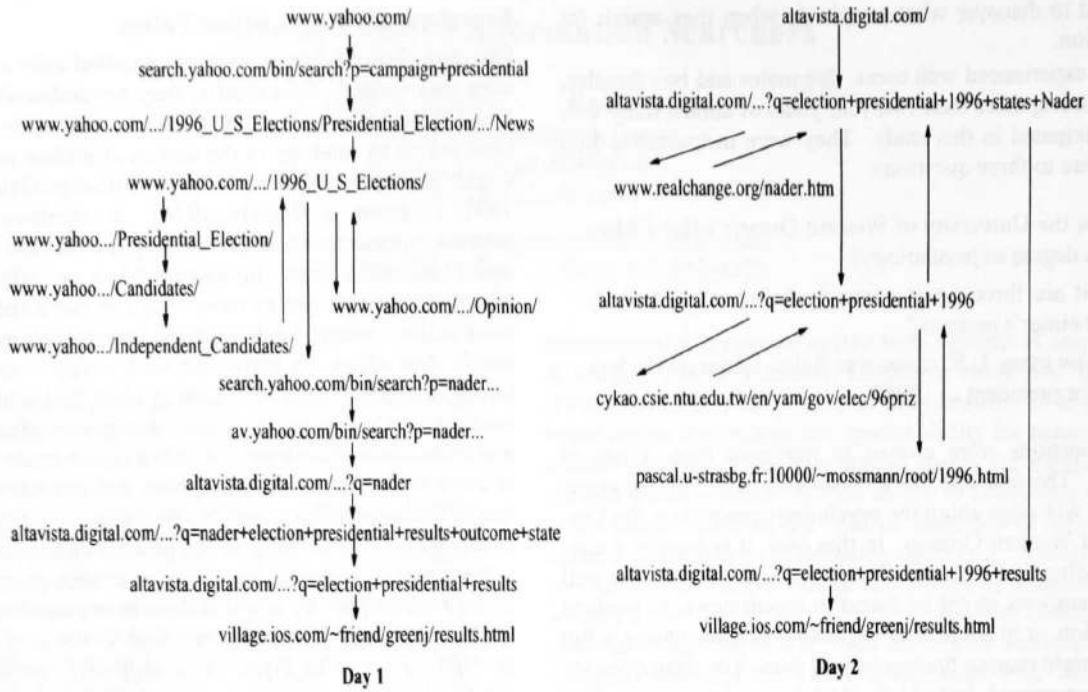


Figure 1: On both days, participant T's search relied on a specific AltaVista query. In this figure (and in those that follow), the nodes represent unique web pages, and the arcs represent transitions between pages.



Figure 2: 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.

one was the library ... so I got into the library and ... I in fact did find a link that took me back ... [to] 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 ...

This time, T recalled that an AltaVista query enabled her to find the library at UWO, from which there was an unbroken series of links to the target. In this case, the library was the anchor, though the specific URL at the library was different on the two days.

Participant D had a similar experience finding the psychology department at UWO. As shown in Figure 3, D's anchor point 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 that was followed the first day did not matter because some path could be found once he was in the region defined by his anchor point.

We found evidence in the behavioral data that each of the seven participants used anchor points in the ways we just described. That is, 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 data suggest that memory for prior searches was structured around important nodes that led to the goal.

Searchers Rely on Familiar Routes

In addition to the use of anchor points, a second observation that emerged from the behavioral data is 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. On the analogy to cognitive maps of physical space, personal routines correspond to the *familiar routes* that an individual uses to get from one landmark (or anchor point) to another.

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 AltaVista 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! (refer back to 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 4). 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 ... 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 4), 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; that is, routines structure internal representations.

Discussion

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

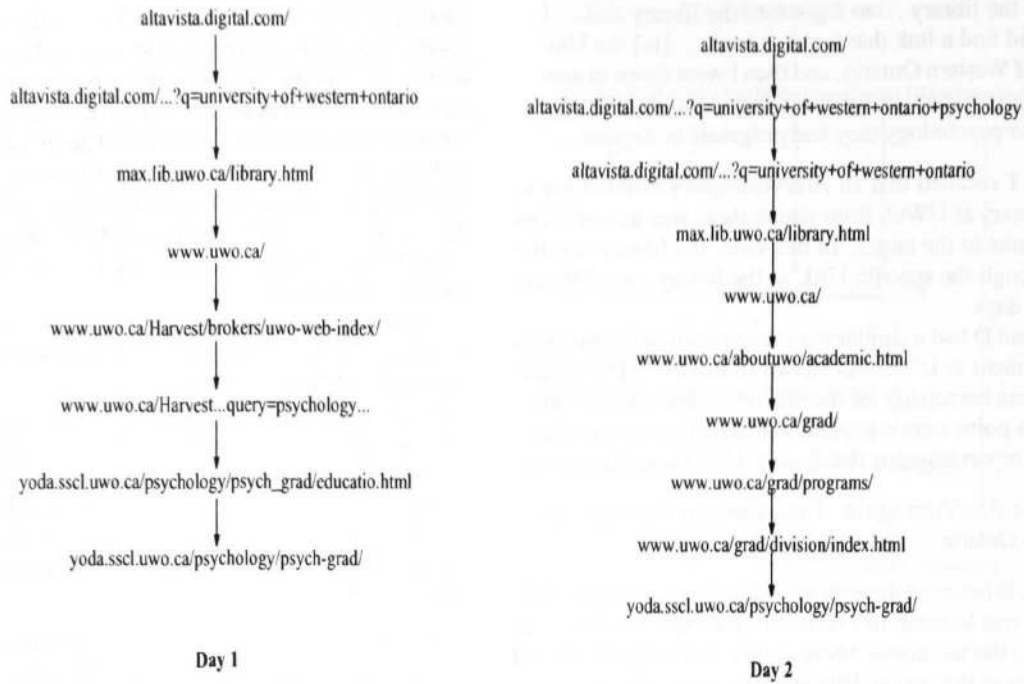


Figure 3: D's search for the psychology department at the University of Western Ontario depended on the UWO's homepage, but took different paths from there on the two days.

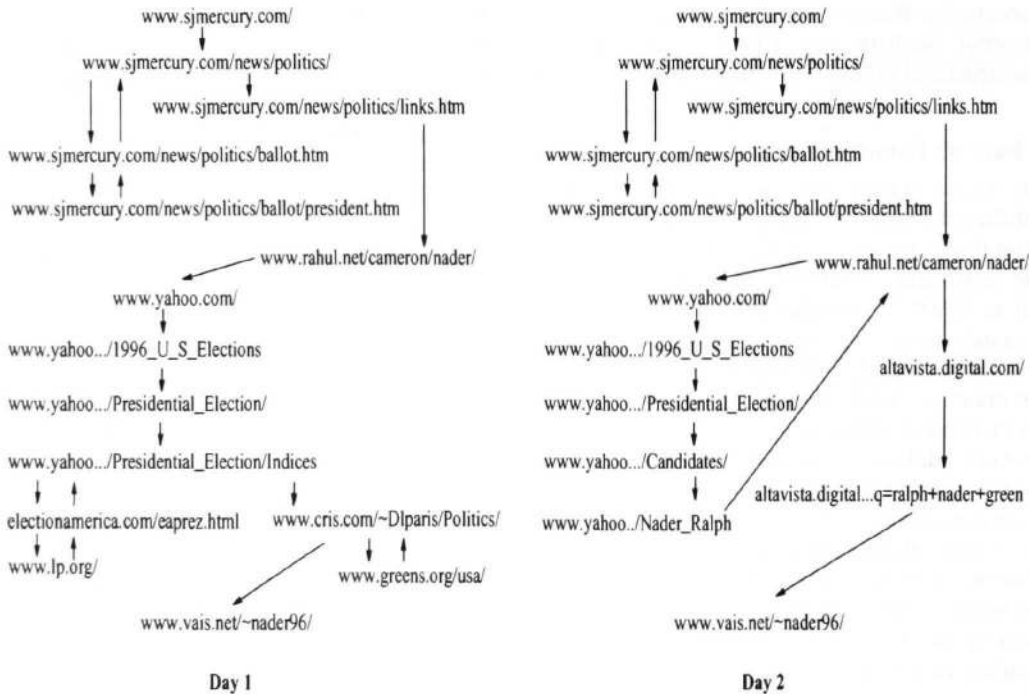


Figure 4: Participant D faithfully retraced the first 11 steps of this search, including dead ends, before switching to his usual routine (i.e., using AltaVista), even though he did not use AltaVista on the first day.

Related Studies of Web Use

Our results accord with other recent studies of how people use the web. Catledge and Pitkow (1995) and Tauscher and Greenberg (1997a, 1997b) analyzed several weeks worth of normal web usage gathered from dozens of college students. 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 anchor points than to refind sequences. People do not follow the same trails because they do not remember their trails: they remember generic procedures and particular anchor points. Thus, although we focused on details of specific search behaviors observed in a few individuals, the patterns we found are consistent with large body of quantitative data.

Cognitive Maps of Web Space

As mentioned, cognitive maps incorporate landmark knowledge, route knowledge, and survey knowledge. As shown, the web searchers in our study recalled their searches in terms of landmarks (anchor points) and routes (personal routines) for moving between them. Thorndyke and Hayes-Roth (1982) found that route knowledge gained by traversing a space becomes transformed by repeated experience into survey knowledge. Thus, one question ripe for investigation is whether web searchers' route knowledge becomes more like survey knowledge with increased experience. One potential problem with this suggestion is that survey maps of such a highly interconnected information space might be too complicated for people to remember and use. Just as route information provides better guidance than map information for a person driving a car in an unfamiliar region (Streeter, Vitello & Wonsiewicz, 1985), anchors and routines might be the most efficient way of representing the web's structure.

In summary, our study of how people find information on the World Wide Web suggests that they conceive of the web's information space like they conceive of physical space. Cognitive maps of the web are comprised of (a) regions of information defined by landmarks and anchor points, (b) and familiar routes for moving among landmarks and anchors.

Acknowledgments

Thanks to Ted Selker and Shumin Zhai for helpful conversations. Thanks to Teenie Matlock for editorial comments.

References

- AltaVista. AltaVista Search. Available as <http://altavista.digital.com>.
- Anderson, J. R. (1980). *Cognitive psychology and its implications*. San Francisco, CA: Freeman.
- Bates, M. J. (1989). The design of browsing and berrypicking techniques for the on-line search interface. *Online Review*, 13, 407–431.
- Campagnoni, F. R. & Ehrlich, K. (1989). Information retrieval using a hypertext-based help system. *ACM Transactions on Information Systems*, 7, 271–291.
- Catledge, L. & Pitkow, J. (1995). Characterizing browsing in the World-Wide Web. In *Proceedings of the Third International World Wide Web Conference*.
- Couclelis, H., Gollidge, G., Gale, N., & Tobler, W. (1987). Exploring the anchor-point hypothesis of spatial cognition. *Journal of Environmental Psychology*, 7, 99–122.
- Dahlback, N., Hook, K., & Sjolinder, M. (1996). Spatial cognition in the mind and in the world—the case of hypermedia navigation. In *Proceedings of the Eighteenth Annual Conference of the Cognitive Science Society*, pages 195–200. Mahwah, NJ: Lawrence Erlbaum.
- Ferguson, E. L. & Hegarty, M. (1994). Properties of cognitive maps constructed from texts. *Memory & Cognition*, 22, 455–473.
- Georgeff, M. & Lansky, A. (1987). Reactive reasoning and planning. In *Proceedings of the Sixth National Conference on Artificial Intelligence*, pages 677–682.
- Marchionini, G. (1995). *Information Seeking in Electronic Environments*. Cambridge, England: Cambridge University Press.
- Matlock, T. & Maglio, P. P. (1996). Apparent motion on the World Wide Web. In *Proceedings of the Eighteenth Annual Conference of the Cognitive Science Society*, page 810. Mahwah, NJ: Lawrence Erlbaum.
- Matlock, T. & Maglio, P. P. (1997). Untangling talk about the World Wide Web. Manuscript submitted for publication.
- Nielsen, J. (1990). The art of navigating hypertext. *Communications of the ACM*, 33, 297–310.
- Streeter, L. A., Vitello, D., & Wonsiewicz, S. (1985). How to tell people where to go: Comparing navigational aids. *International Journal of Man-Machine Studies*, 22, 549–562.
- Tauscher, L. & Greenberg, S. (1997a). How people revisit web pages: Empirical findings and implications for the design of history systems. *International Journal of Human-Computer Studies*.
- Tauscher, L. & Greenberg, S. (1997b). Revisitation patterns in World Wide Web navigation. In *Proceedings of the Conference on Human Factors in Computer Systems (CHI '97)*, pages 399–406. New York, NY: ACM Press.
- Thorndyke, P. W. & Hayes-Roth, B. (1982). Differences in spatial knowledge acquired from maps and navigation. *Cognitive Psychology*, 14, 560–589.
- Vicente, K. J. & Williges, R. C. (1988). Accommodating individual differences in searching a hierarchical file system. *International Journal of Man-Machine Studies*, 29, 647–668.
- Yahoo! Yahoo! Available as <http://www.yahoo.com>.