



massachusetts institute of technology — artificial intelligence laboratory

Surviving the Information Explosion: How People Find Their Electronic Information

Christine Alvarado, Jaime Teevan,
Mark S. Ackerman and David Karger

AI Memo 2003-006

April 2003

Abstract: We report on a study of how people look for information within email, files, and the Web. When locating a document or searching for a specific answer, people relied on their contextual knowledge of their information target to help them find it, often associating the target with a specific document. They appeared to prefer to use this contextual information as a guide in navigating locally in small steps to the desired document rather than directly jumping to their target. We found this behavior was especially true for people with unstructured information organization. We discuss the implications of our findings for the design of personal information management tools.

Funding: This research was supported by NTT, the Packard Foundation, MIT's Project Oxygen, the Arthur P. Sloan Foundation and the National Science Foundation. The third author was also supported by NSF grant IRI-0124878.

1 Introduction

Every day computer users manage and retrieve information from among dozens of incoming emails, hundreds of files on their personal computers, and billions of Web pages. Electronic information management can be difficult and time-consuming. Recent research has focused on building better tools to help people manage their information (e.g. Fertig, Freeman and Gelertner, 1996; Huynh, Karger and Quan, 2002). The proposed tools attempt to provide a more natural way of both organizing and accessing personal information. However, in our opinion, we do not yet have a sufficient understanding of what people will find natural.

In this paper we present a qualitative study examining what people did when working with their email, their file system, and the World Wide Web. We did not set about to test specific hypotheses, but rather to understand our participants' behavior. We wanted to understand what people did with their electronic information, and we focused in particular on the situations in which people reported exerting effort in locating information. In brief, we found people often had an association between the information they looked for and related contextual information. For example, a participant might know she could find the phone number of a restaurant in a particular email from a colleague. In many of these cases, people were able to associate their information target with a particular source. People seemed to prefer to solve their need by using a form of local navigation to find that source, similar to the Micronesian islanders' situated navigation in Suchman (1987). This kind of search appears to be important but under-supported.

This paper proceeds as follows. We first discuss related observational studies that have focused on how people interact with their information. We then describe the particulars of the study we performed. We further examine the local navigation search strategy described above, and discuss the implications of those results on personal information management tools.

2 Previous Work

Previous observational studies have focused on users' interaction with various different subsets of their personal information, such as their email, their files, and the Web. We use the term *corpus* to refer to any of these major subsets that are treated as a group by the user's tools.

Early studies provided a theoretical motivation for electronic information management by examining the ways that people organized their paper documents. Lansdale (1988) noted that the difficulty people had in retrieving information from this corpus arose because they were forced to group their information into categories that were not necessarily relevant for retrieval. Malone (1983) also studied how people organized paper documents. More recently, Whittaker and Hirshberg (2001) investigated personal paper archives to understand the value of paper over digital documents.

Similarly, a number of studies have investigated the different uses that people have for their email (Mackay, 1988; Whittaker and Sidner, 1996). Researchers found that in addition to using email as a communication tool, people also used it to keep track of upcoming appointments and often used their Inboxes as to-do lists. More recent studies on email, scheduling tools and instant messaging further explored these ideas (Bellotti and Smith, 2000; Nardi, Whittaker and Bradner, 2000).

Another class of studies investigated how people organize their computer files. While file management systems rely largely on content hierarchy, two studies (Nardi and Barreau, 1997; Nardi, Whittaker and Bradner, 2000) found that users preferred to use location-based search in storing and retrieving information, placing all of their currently relevant documents on the desktop and associating a location with each document. Fertig, Freeman, and Gelernter (1996) argued that people rely on temporal cues to locate files.

Finally, there is a large body of literature reporting on how people use the World Wide Web. These studies address both how people manage their bookmarks (Abrams, Baecker and Chignell,

1998), the various tasks people perform using the Web (Sellen, Murphy and Shaw 2002), and how people keep track of information on the Web (Jones, Bruce and Dumais 2001). A body of literature on information seeking (Chi, Pirolli, Chen and Pitkow, 1996; O'Day and Jeffries, 1993; Pirolli and Card, 1999) discusses how people seek out information from large collections, often focusing on the Web.

Some observational studies have been conducted across different corpora. For example, Jones, Bruce and Dumais (2001) have recently looked at how Web information is stored by observing user interaction with the same three electronic corpora (email, files and the Web) that we investigated. However, they focused specifically on how people used files and email to support their Web interactions. We took a broader focus and aimed to uncover how people interacted with electronic information in general.

3 Methods

We conducted semi-structured interviews in which participants reported their information activities twice a day over the course of a week. The interviewer would interrupt participants' work and prompt them for their recent information activities in the three different corpora. We felt this method would enable us to properly understand how people interact with their electronic information, since participants worked with their own information. Our method was similar to the diary studies used in many information seeking studies, as well as the Experimental Sampling Method (Palen, 2002).

Our participants consisted of 15 graduate students (5 women, 10 men) in Computer Science at *. All participants were experts at computer use. The length of time the participants had been at the university varied from one year to seven. By looking at a range we were able to observe both those who were in the process of developing their information organization, as well as those who had a long standing structure.

While Computer Science graduate students were convenient to study, more importantly, they were experienced users with complex digital information spaces. As such, our participants could discuss the issues that arise in current

information organization tools that occur despite users' considerable experience with those tools. In general, we believe that our participants are typical of many experienced users, but we will discuss the impact of * and Computer Science as cultures in Section 7.

We randomly interrupted each participant's work twice a day for five consecutive days. Each semi-structured interview lasted only five minutes in order not to unduly interrupt the participants' work. In the interviews, we asked the participants to describe what they most recently "looked at" and what they most recently "looked for" in each of the three corpora we studied (email, files and Web).

What precisely defined "looking for" versus merely "looking at", or *accessing*, was defined by the participants themselves based on what they considered effort. By allowing participants to self-categorize when they had to exert effort to find information (as in, for example, Bernard 1994), we were able to learn what types of information needs required effort and what techniques they relied on in those cases. We encouraged the participants to give as much detail as possible.

Overall, we obtained 151² interviews. In addition, we conducted longer semi-structured interviews (1 hour) with each participant about their information patterns and conducted some direct observations. The data were analyzed using standard qualitative techniques (e.g., Ackerman and Halverson, 1998).

4 Information Management

This section describes the participants' reported information activities and focuses on specific types of behavior we found to be important. In general, participants in our study managed incredibly complex information spaces. For example, Alex³ had two email accounts containing hundreds of email messages in his Inbox alone and many more in email directories (e.g., one email account was

² We inadvertently interviewed one participant 11 times. This participant is labeled "M" in future charts.

³ All names and details reported in this paper have been anonymized. Minor changes to the transcripts have been made for readability.

further divided into approximately 250 email directories, each containing anywhere from tens to hundreds of messages). His Web space and file system were similarly complex.

Despite their complex information spaces, participants were largely successful at managing and finding things. Most participants felt fairly in control of their information. For example, Alex, described above, when asked if he felt in control responded:

That's an interesting question. I think my email is the worst, because I have so much of it. And there are people on the other end who expect me to reply to it. My file system is pretty well organized. I have to go through it every once in a while, every couple of months and just kind of push things into the right folders and delete the old stuff. The Web just works usually.

In fact, we find that, despite people's large and complex information repositories, most information activity involved simply accessing information and did not involve the user exerting effort to find information. When we asked about their most recent information activity, only in 50 of the 453 instances (151 interviews times 3 corpora) did participants report "looking for" something. However, when we specifically asked participants what they had looked for most recently (following up, in most cases, upon their initial response), they reported 200 instances (of 453—not everyone remembered looking for something in each corpus prior to each interview). As might be expected, a number of these activities occurred on the Web (n=83, 42%). However, a significant number also occurred within email (n=65, 33%) and files (n=52, 26%).

The following subsection discusses a prevalent search strategy that relies on contextual information and that differs from the traditional concept of keyword search. In order to learn more about this under-explored area of search we then look at the different information targets people had when they searched.

4.1 More than Just Keyword Search

In the interviews, people often reported looking for things without resorting to traditional keyword search, as the following example illustrates. Jim is looking for the office address of a professor named Connie Monroe:

Interviewer: Have you looked for anything on the Web today?

Jim: I had to look for the office number of the Harvard professor.

I: So how did you go about doing that?

J: I went to the home page of Math Department at Harvard.

This participant then goes on to explain that he knows there is a specific Web page with the address:

I: Did you know it would be there [on a page] or you just hoped it would be there?

J: I knew that she had a very small Web page saying, I'm here at Harvard. Here's my contact information.

[...]

I: So you went to the Math department, and then what did you do over there?

J: It had a place where you can find people, a link to the page where you can find people and I went to that page and they had a dropdown list of visiting faculty, and so I went to that link and I looked for her name and there it was.

While Jim only wants Connie's office number, he first goes to the Harvard Web page, and then navigates in toward the information he is after. This search by localized or situated navigation is an example of what we call *orienteering*. Orienteering involves using contextual information to narrow in on the actual information target, often in a series of steps. This definition is similar to that of O'Day and Jeffries (1993), as we use orienteering to imply using information from the current location to decide where to go next. However, unlike their definition, our participants did not generally have an evolving information need. Orienteering, then, is an extension of situated activity reported in the Computer Supported Cooperative Work literature (e.g., Suchman, 1987; Ackerman and Halverson, 1998).

Participants also reported what we call *teleporting*. When a person teleports, they try to take themselves directly to the information they're looking for. For example, if instead of orienteering, Jim had tried to teleport, he could have typed "Connie Monroe office number" into a search engine hoping to find it directly.

It is important to note Bates' distinction between search strategies and search tactics here (Bates, 1979). Orienteering and teleporting are strategies; participants can use the same search tactic, in some cases, to achieve either strategy. For example, participants sometimes reported using keyword search in orienteering as illustrated by Carla when she had to look for a page she lost when her network connection died. Although she

performs a keyword search, her behavior is clearly orienteering because she takes small steps to narrow in on a goal.

I did a re-search for it on Google and then I clicked into it. [...] I actually wanted not the main page for Bon Jour Quebec, but I kept clicking on the links, from the main page. It was very easy to get what I wanted to get.

In Tables 1 and 2 we use a very conservative measure of orienteering, classifying an instance as teleporting if it involved a keyword search at any stage in the process. Even with this conservative measure, 60% of the time people looked for something, they orienteered.

4.2 Information Target

People used contextual information when orienteering to their information target. Because of the difference between the information people used in orienteering and the information they were actually looking for, it is fundamental in understanding orienteering to understand a person's information target. We found that when people looked for information, they looked for qualitatively different things across instances.

We investigated people's information targets by first labeling the participant's particular goal in each seeking activity they reported and then clustering those goals into general categories. To help us do this, in our interviews we asked people to not only tell us what they were looking for, but also why they wanted to it. We identified 21 different goals. Two of the authors independently categorized each seeking activity as one of these 21 goals. In the 3% of the cases where the researchers disagreed and in the 7% of the cases where the goal could not be determined, the activity was classified as "unknown".

When we further collapsed the 21 goals three main categories of information needs emerged: specific information, general information, and specific documents (emails, files or Web pages). A person has a specific information need when he is looking for a small fact, as typified by research on question answering (e.g., find the time of a meeting). When a user looks for general information, the user is interested in a broader set of information (e.g., determine a good pair of sneakers to buy). In the case when a specific document is the target, it is the actual document that is desired (e.g., a file to edit), and that

document is not replaceable by the information contained within it.

The existence of these three categories revealed unexpected search strategy patterns. Table 1 shows how often people used the different search strategies for each of the different targets. The large variety in search strategy in relation to information target implies that the two are closely related.

	Spec	Gen	Doc	Total
Orient	47	19	41	120
Teleport	34	23	17	80
Total	81	42	58	200

Table 1: Information need by search strategy (19 unknowns removed).

We were not surprised to find that people orienteered to specific documents. We expected people to maintain a large amount of context about documents. However, we expected to find that people would try to go directly (with a keyword search) to a specific information target and were surprised that this was not the case. Instead, people were much more likely to orienteer when looking for specific information than general information.

5 Finding the Information Source

We observed that people not only used orienteering techniques to look for documents, but also when looking for specific information. This behavior implies that people maintained a large amount of contextual information about the specific piece of information they are looking for. In particular, part of this contextual information involved associating the piece of information with a source, such as an email or specific Web page containing the desired information. We observed documents and specific information were often conflated in participants' descriptions. Due to this tendency, in many instances, their strategies for finding a specific piece of information reflect the best way of locating the source of the information. Marchionini (1995) observed the same information seeking behavior for large information collections, and Hearst (2000) suggested this behavior holds on the Web.

To better understand how participants associated information with an information source, consider the following activity, in which Dan is searching within his email:

Dan: Earlier today I was doing a search for a message that Kristi sent me that had this guy's address.

Interviewer: So you needed the address and then you remembered Kristi had sent it?

D: Yeah.

In this example, Dan tries to find someone's address, but instead of seeking the address directly (through a keyword search in his inbox or on the Web), Dan associates the address with an email he received from a colleague and looks instead for that email.

Users also made similar associations in the other corpora. For example, one participant had difficulty finding the location of a city in Switzerland. He did not know exactly where to find that information, but he had four map sites bookmarked. Rather than relying on a keyword search directly to locate the city, he used the bookmarks to access the map sites and then clicked around to see if he could find a map with the information he was looking for.

Associating a specific information need with a source is one piece of contextual information that participants maintain about this information. In addition, they maintain a lot of contextual information about the information source itself that helps them in their orienteering process. In the Harvard office example (Section 4.1), Jim knew the approximate URL for the page that contained Connie's office number, as well as the size of the containing page and what additional content would appear there.

It should be noted that the contextual information used is not always definitive, such as in the following example of a participant looking to find a specific piece of information in her email. In this example, the participant has only the memory that a piece of email was in a specific directory:

The last email I read was an email from Bill describing where to find the documentation on [a project]. I had searched for this email. It was an old email that I had. And I had to look for it. And I looked for it in the research directory which was where I put things that are sort of done for a research. [...] So anyway, so, in my research directory I found this. But I actually had to flip through all emails. I went and tried to look for the email that looked familiar for being the correct one. The only thing I had to go by was that it was probably

from Bill. But I wasn't exactly positive on that. And I wasn't sure where it would be anyway. So it took a long time to find. But I found it.

Orienteering and teleporting were not used uniformly across the corpora. Table 2 shows the number of searches in each corpus by each search strategy. Recall that in the table below, any instance in which a keyword search was used was classified as teleporting, which is a conservative measure of orienteering. As we shall see below, many of those cases are actually instances of orienteering.

	Orienteering	Teleporting
Email	59	6
Files	42	10
Web	19	64

Table 2: The number of times participants used each search strategy for each corpus

In general, participants preferred to orienteer toward their target. Although they occasionally used keyword search to help them narrow in on their target, they rarely teleported directly to what they were looking for. The next sections examine this behavior further for each of the three corpora.

5.1 Email

In email, almost none of the seeking activities involved teleporting, even though most participants' email programs facilitated it by supporting keyword search. Participants often looked in their email for specific information that they knew to be contained within a particular message. Most often people went directly to the correct folder (approximately half of the time, their Inbox) that contained the message. Then they would browse to the appropriate email, using either the date (the default ordering in most mail systems) or the sender as a guide. In the following example, Carla looks in her email to find the location of a meeting:

Carla: The last email I looked at was a email for Mischa's reading group I wanted to find out where we were meeting. [...] It was in my inbox because I would have left it in my Inbox. I knew that it was in my inbox and I searched for it knowing that it was by Mischa and I only had two emails from him.

Interviewer: Did you sort?

C: I'm always sorted by date.

I: By date, okay, so you didn't sort by sender?

C: I didn't sort by sender. I don't have so many people in my inbox. I knew that it was rather recently. So, since I knew it was rather recently it was sort of easier

just to visually sort. I just have to scroll down one or two pages to find the email.

In finding the location of the meeting, Carla relies on a large amount of contextual information: it is in an email sent to her by Mischa, this email is in her inbox because she left it there, she only has two emails from Mischa, etc. Note that most of this contextual information has to do with the email containing the information about the meeting and nothing to do with the meeting in question.

In three cases classified as teleporting, the same participant searched for the sender of the message in question and then browsed through the returned messages. Because this participant's email client did not support sorting the messages by sender, this participant likely relied on keyword search to fulfill the same role.

5.2 Files

Participants used orienteering techniques much more often than teleporting within their files. Most often their goal was actually to find a particular file to read or edit. Although participants did not often look for specific information within their files, when they did, they again tended to orienteer to the file containing that information, as in the following example:

I actually was not looking for a file. I was looking for the contents of a particular file. ... I remembered that a Perl template that I had had those four lines of Perl in it. So I went to look at where my templates are.

In some cases the user did not appear to be able to associate the information with a particular source, and they used a keyword search to teleport to the information, as illustrated in this example:

I needed to search through some Emacs source files in order to find a particular line of code that I needed. I went to the directory above all and then did a recursive `grep` [keyword search] down the tree looking for anything that had the appropriate set of characters in it.

5.3 Web

On the Web, participants' search techniques were more varied and they tended to rely on teleporting more often, possibly because keyword search tools are more sophisticated for the Web than for files or email, because of the larger search space, or because the information on the Web continually changes its structure.

Even on the Web, participants did sometimes associate their information needs with a specific information source and seemed to prefer to

orienteer to that source when possible, as in the example in which Jim is looking for the Harvard professor's office number. Additionally, at least one fourth of the instances listed as teleporting in Table 2 were actually cases of orienteering where keyword search was used as a step, as in the example where Carla used a keyword search to get to a page on Quebec and then continued narrowing in on her target from this page. As well, 16 of the searches classified as teleporting consisted of searches where the user had already navigated to the neighborhood of their information target (e.g., a company's home page)—an orienteering activity—before performing the keyword search.

As with email and files, participants tended to teleport when they could not associate their information need with a specific source. In one case, a user spent a considerable amount of time looking for how much to tip hairdressers simply performing various keyword searches using the words "tip", "hairdresser", "percent", and "gratuities".

6 Individual Tactics

While everyone orienteered as well as teleported as strategies, some people used keyword search more as a tactic. Surprisingly, these same individuals tended to put more effort into organization.

These two groups emerged from our observations of people's email use. When people reported looking for information in their email, they either found that information in their Inbox a majority of the time, or they found it in a specific folder a majority of the time. Figure 1 shows the percentage of time the 13 participants who reported looking for something in their email found it in their Inbox⁴. Notice that the difference was not because one group searched more, as the number of instances was similar for the two groups.

⁴ Two participants never reported finding anything in their email.

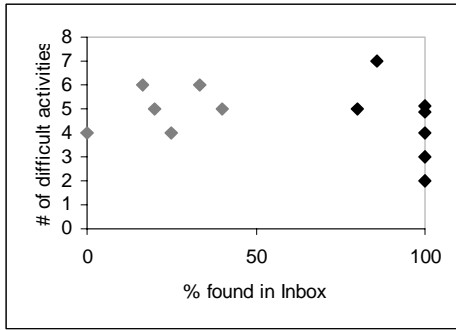


Figure 1: For each user, the percent of email targets found in the user's Inbox.

Those on the right-hand cluster of Figure 1 almost never spoke of interacting with emails that were not in their Inboxes and almost always expected to find messages in their Inboxes, implying they did not file their messages in general. Those in the left-hand cluster almost always went directly to folders and never expected to find messages in their Inboxes implying they regularly filed their messages. We can view the left-hand group in Figure 1 as *filers*, and the right-hand group as *pilers* (Malone 1983). A *filer* is a person who organizes information using a rigid structure, and a *piler* is someone who maintains a mostly unstructured information organization.

Filers and pilers tended to rely on different search tactics when looking for things within their files and on the Web. Figures 2 and 3 show the number of different search tactics used by participants in files and on the Web, respectively. The top 6 participants in each graph are those classified as filers based on their email search patterns. From Figure 2 it appears that filers reported having to look for files (or information within those files) more often than pilers. Furthermore, filers relied more on keyword search than pilers.

More work will be required to understand the nature of this association. It appears that both groups orienteered toward the information they were seeking, so we do not believe that filers were more likely to teleport. Rather, perhaps pilers associate a finer grain of contextual information with what they are looking for and then can take more local steps to get to their goal. Because they are more confident in their ability to rely on contextual information, they do not need to

maintain complicated organizational structures to keep track of their information. Moreover, they do not need to rely on keyword search because they are able to take local steps using this contextual information.




Figure 2: The number of times participants used each search tactic in their files.




Figure 3: The number of times participants used each tactic on the Web. Note that pilers appear to use specific search tools more often.

7 Conclusions and Future Work

In this paper, we reported on a study of how people looked for electronic information in their email, their file system, and on the Web. We found that a majority of the times that people looked for things they used an orienteering strategy, taking advantage of the large amount of contextual information they had about their information target.

As with any qualitatively-based study, there are limitations to this study. We could point here only to the existence of an interesting phenomenon, and discuss its potential associations and nuances. Further work will be required to assess how

general and pervasive orienteering might be as well as to assess its casual roots and effects. Nonetheless, we believe we have described an important search strategy to support. As well, our participants were members of the * and Computer Science cultures, both of which place an interest in and a value on information handling. However, in our opinion, the observation that this population valued orienteering only shows orienteering's importance in dealing with large amounts of electronic information.

Several possibilities exist to explain our findings. It could be that people orienteer rather than teleport because current search tools do not work well enough to find what the user requests. This problem may be alleviated by work in information retrieval. Or, current tools may not allow users to take advantage of their contextual information and fully specify the information they are looking for. New tools are attempting to address this by adding additional meta-data support, mostly within keyword search tools (e.g., Microsoft XP).

However, as noted in this paper, we observed that even when people could use their contextual information to teleport directly to their information target, they often preferred to orienteer to the information instead. We believe this occurs because fully specifying an information need and all of its meta-data would require considerable cognitive effort. Often our participants were not entirely sure of the contextual information they could use to retrieve their target, or they were not even aware of it until reminded. Orienteering also has the added benefit of helping the user not to over-specify her target, allowing her to backtrack more easily. It also gives her information about the source, which may be important in determining the validity of the information. These possible explanations suggest that future systems should deeply consider orienteering or other approaches to help people use contextual data, perhaps by prompting them with contextual information instead of requiring them to fully specify all of this information at query time.

To support the use of contextual information, we must understand exactly what information people know about their information target. People often orienteered to targets they had seen before,

and in these cases appeared to use different contextual information than in cases when they had not seen their target previously. In future work we will further examine the nature of the contextual information used in each case.

As well, what people remember appears to be corpus dependent. What was sometimes considered "looking for" in one corpus was not the same for the others. The intricacies of these boundaries, and their use in search and retrieval, are other interesting areas to pursue.

As the amount of information we interact with grows, electronic information management will increasingly become a problem we must deal with. Our study revealed behavioral patterns we can examine further in order to build tools to make this interaction more manageable in the future.

References

- Abrams, D., Baecker, R. and Chignell, M. (1998). Information archiving with bookmarks: personal Web space construction and organization, Proceedings CHI '98, 141-148.
- Ackerman, M. S., and Halverson, C. (1998). considering an organization's memory. Proceedings of CSCW'98, 39-48.
- Bates, M. (1979). Information search tactics. Journal of the American Society for Information Science: 205-214.
- Bellotti, V. and Smith, I. (2000). Informing the design of an information management system with iterative fieldwork. In Proceedings of DIS '00, 227-238.
- Bernard, H. R. (1994). Research methods in anthropology: Qualitative and quantitative approaches. Landham, MD: Altamira Press.
- Buckland, M., & Gey, F. (1994). The relationship between recall and precision. Journal of the American Society for Information Science, 45, 12-19.
- Chi, E. H., Pirolli, P., Chen, K., and Pitkoww, J. (2001). Using information scent to model user information needs and actions on the Web. In Proceedings of CHI '01, 490-497.
- Fertig, S., Freeman, E. and Gelernter, D. (1996). Lifestreams: an alternative to the desktop metaphor. In Proceedings of CHI '96, 410-411.

- Hearst, M. (2000). Next generation Web search: setting our sites. *IEEE Data Engineering Bulletin*, Special issue on Next Generation Web Search, 23, 3, 38-48.
- Huynh, D., Karger, D., and Quan, D (2002). Haystack: A platform for creating, organizing and visualizing in-formation using RDF. In *Semantic Web Workshop 2002 Proceedings*.
- Jones, W., Bruce, H., and Dumais, S. (2001). Keeping found things found on the Web. In *Proceedings of CIKM '01*, 119-134.
- Lansdale, M. (1988). The psychology of personal information management. *Applied Ergonomics* 19, 1, 458-465.
- Mackay, W. E. (1988). More than just a communication system: diversity in the use of electronic mail. In *Proceedings of CSCW '88*, 344-353.
- Malone, T. E. (1983). How do people organize their desks? *ACM Transactions on Office Information Systems* 1, 1, 99-112.
- Mander, R., Salomon, G. and Wong, Y. Y. (1992). A "pile" metaphor for supporting casual organization of information. In *Proceedings of CHI '92*, 627-634.
- Marchionini, Gary. (1995). *Information Seeking in Electronic Environments*. New York: Cambridge University Press.
- Nardi, B. and Barreau, D. (1995). Finding and reminding: file organization from the desktop. *ACM SIGCHI Bulletin* 27, 3, 39-43.
- Nardi, B. and Barreau, D. (1997). "Finding and reminding" revisited: appropriate metaphors for file organization at the desktop. *ACM SIGCHI Bulletin* 29, 1, 76-78.
- Nardi, B., Whittaker, S. and Bradner, E. (2000). Interaction and outeraction: instant messaging in action. In *Proceedings of CSCW '00*, 79-88.
- O'Day, V., and Jeffries, R. (1993). Orienteering in an information landscape: How information seekers get from here to there. In *Proceedings of CHI '93*, 438-445.
- Palen, L. and Salzman M. (2002). Voice-mail diary studies for naturalistic data capture under mobile conditions. *Proceedings of the CSCW'02*, 87-95.
- Pirolli, P., and Card, S. (1999). Information foraging. *Psychological Review* 106, 4, 643-675.
- Sellen, A. J., Murphy, R. and Shaw, K. (2002). How knowledge workers use the Web. In *Proceedings of CHI '02*, 227-234.
- Suchman, L. A. (1987). *Plans and situated actions*. New York: Cambridge University Press.
- Whittaker, S. and Hirschberg, J. (2001). The character, value, and management of personal paper archives. *ACM Transactions of Computer-Human Interaction* 8, 2, 150-170.
- Whittaker, S. and Sidner, C. (1996). Email overload: exploring personal information management of email. In *Proceedings of CHI '96*, 276-283.