

Improved Search Engines and Navigation Preference in Personal Information Management

OFER BERGMAN

Sheffield University

RUTH BEYTH-MAROM

The Open University of Israel

RAFI NACHMIAS AND NOA GRADOVITCH

Tel Aviv University

and

STEVE WHITTAKER

Sheffield University

Traditionally users access their personal files mainly by using folder navigation. We evaluate whether recent improvements in desktop search have changed this fundamental aspect of Personal Information Management (PIM). We tested this in two studies using the same questionnaire: (a) The Windows Study - a longitudinal comparison of *Google Desktop* and *Windows XP Search Companion*, and (b) The Mac Study - a large scale comparison of Mac *Spotlight* and *Sherlock*. There were few effects for improved search. First, regardless of search engine, there was a strong navigation preference: on average, users estimated that they used navigation for 56-68% of file retrieval events but searched for only 4-15% of events. Second, the effect of improving the quality of the search engine on search usage was limited and inconsistent. Third, search was used mainly as a last resort when users could not remember file location. Finally, there was no evidence that using improved desktop search engines leads people to change their filing habits to become less reliant on hierarchical file organization. We conclude by offering theoretical explanations for navigation preference, relating to differences between PIM and Internet retrieval, and suggest alternative design directions for PIM systems.

Categories and Subject Descriptors: H.5.2 User Interfaces (D.2.2, H.1.2, I.3.6): Evaluation/methodology; E.5 Files (D.4.3): sorting/searching, organization/structure

General Terms: Design, Human Factors, Measurement

Additional Key Words and Phrases: Personal information management, user study, files retrieval, search preference, navigation preference, personal search engines

1. INTRODUCTION

Personal Information Management (PIM) is an activity in which an individual stores his/her personal information items in order to retrieve and use them later. Such information items include files, emails, Web favourites, contacts and notes. In today's personal computers, the two main ways by which items can be retrieved are through Hierarchical Navigation and Search. Hierarchical Navigation (*Navigation* for short) is a two-phase process. First, users manually traverse their organizational hierarchy until they

Authors' addresses: Ofer Bergman & Steve Whittaker - Information Studies Department, Sheffield University, Regent Court, 211 Portobello Street, Sheffield, S1 4DP, UK, {s.whittaker; O.Bergman}@sheffield.ac.uk; Ruth Beyth-Marom: Department of Education and Psychology, The Open University of Israel, The Dorothy de Rothschild Campus, 108 Ravutski Street, P. O. Box 808, Raanana 43107, Israel, ruthbm@openu.ac.il; Rafi Nachmias, Noa Gradovitch, School of Education, Tel-Aviv University, Tel-Aviv 69978, Israel {nachmias; noahgrad}@post.tau.ac.il.

© ACM, 2008. This is the author's version of the work. It is posted here by permission of ACM for your personal use. Not for redistribution. The definitive version was published in ACM Transactions on Information Systems (TOIS), Volume 26, Issue 4 (September 2008) <http://doi.acm.org/10.1145/1402256.1402259>

reach the location in which the target item is stored (be it a directory or a folder). Second, they locate it within the directory or folder (either actively by sorting the items by attribute or by using the system default). *Search* is a process in which users first generate a query specifying some property of the target item, including at least one word related to the name of the information item, and/or the text that it contains (full text search) and/or any attribute relating to that item (e.g. the date the item was modified). The search engine returns a set of results from which the user selects the relevant item.

The main purpose of our research was to examine whether personal computer users prefer navigation or search when retrieving their files, and to test the effect of improved search engines on these retrieval preferences. We also looked at why people used search or navigation, and the effects of improved desktop search on organisational behaviours. We conducted two studies addressing these research questions, using variants of the same questionnaire, but testing different search engines and using different research designs. The Windows Study was a longitudinal study in which we installed *Google Desktop* search on the personal computers of 47 Windows XP users and examined changes in their retrieval habits over a period of 7 months. In the Mac Study, we compared the retrieval preferences of 519 Mac OS X 10.4 users who use the more advanced *Spotlight* search engine, to those of 70 Mac OS 10.0-10.3 users who use the older *Sherlock* search engine. The next section presents the theoretical motivation for our research.

2. THEORETICAL BACKGROUND

2.1 The Hierarchical Method

Navigation, in contrast to search, generally requires hierarchical storage, i.e. users need to create folders or directories and to store the information items “inside” them in preparation for future retrieval and use. Although other navigation methods have been proposed such as faceted classification and hypertext, neither is in common usage in widely used operating systems, so we restrict our discussion here to common hierarchical methods. Hierarchical storage was first introduced to end-users in the *Multics* operating system in the mid 60s. Users were allocated a personal directory, in which they could create their own subdirectories, sub-subdirectories, etc., and store their files in any of these “locations.” This directory structure was later applied in the Unix and the Linux operating systems. The location metaphor became even clearer with the creation of digital folders first introduced in the *Xerox Star* in 1981. A folder is a visual metaphor for a location: users can see information items “inside” folders, as well as manipulate items and folders in various straightforward ways, e.g. drag and drop information items from one folder to another, etc.¹ This folder hierarchy metaphor was later applied by Apple in the Mac operating systems and then by Microsoft in their Windows operating systems. Thus, location-based storage has been used without significant modifications, continuously and almost exclusively for several decades.

2.2 Search as an Alternative to Navigation

Through most of its long history, the hierarchical method has met with criticism. One disadvantage is that classification of information can ‘hide’ it from the user, and therefore reduce the chances of quick retrieval or reminding [Kidd 1994; Malone 1983; Whittaker and Sidner 1996]. In addition, the act of categorisation is itself cognitively challenging; users may find it hard to categorise information that could be stored in more than one category [Dumais et al. 2003; Malone 1983; Whittaker and Hirschberg 2001; Whittaker and Sidner 1996]. Categorisation is also difficult because it requires that people anticipate future usage; moreover, that usage may change over time [Kidd 1994; Whittaker and

¹ Of course these directories and folders do not really contain information items from a computer science point of view, but they do for the users, thus for the sake of simplicity we refer to them as doing so.

Hirschberg 2001; Whittaker and Sidner 1996]. At retrieval time, users need to recall how information was classified, which can be difficult when there are multiple categorisation possibilities [Lansdale 1988]. These problems were illustrated in a study of email categorisation by Whittaker and Sidner [1996]. They found that users with many categories found it harder to file, and were more likely to create spurious unused folders.

These apparent problems with navigation caused many PIM researchers and software developers to turn to *Search* as an alternative. There are intuitive potential advantages of search for both retrieval and organization. Search promises to be more flexible and efficient at *retrieval*, it does not depend on remembering the correct storage location; instead, users can specify in their query any attribute they happen to remember [Lansdale 1988]. They can also retrieve information via a single query instead of using multiple operations to laboriously navigate to the relevant part of their folder hierarchy. Regarding storage, search potentially finesses the *organizational problem* - as users don't have to engage in complex organizational strategies that exhaustively anticipate their future retrieval requirements. These arguments against navigation have been bolstered by recent developments in web access, where initial use of navigational systems such as Yahoo categories has almost completely been superseded by search engines such as Google [Kobayashi and Takeda 2000].

The same logic has led to the development of experimental PIM search engines such as *Phlat* [Cutrell et al. 2006b], *SIS* [Dumais et al. 2003], *Haystack* [Adar et al. 1999], and *Raton Laveur* [Bellotti and Smith 2000], as well as commercial systems such as *Einfish Personal*, *Copernic Desktop Search*, *Yahoo! Desktop Search* and *Microsoft Desktop Search*. Some more radical systems such as *Lifestreams* [Freeman and Gelernter 1996], *Canon Cat* [Raskin 2000], *Presto* [Dourish et al. 1999], *Placeless Documents* [Dourish et al. 2000], *MyLifeBits* [Gemmell et al. 2002], and *Swiftware* explore alternatives to location-based hierarchies.

However despite the rapid development of new such technologies, we know little about the effects of improved desktop search on user behaviour. In this study therefore we set out to test the following predictions about the effects of desktop search on both file retrieval and organization:

Retrieval: Search is more efficient and flexible for retrieval [Lansdale 1988, Fertig, Freeman and Gelernter 1996], thus improved quality of search engines should lead to a substantial increase in file search and eventually a preference for search over navigation.

File Organization: Users are known to have problems organizing files effectively for retrieval [Kidd 1994, Malone 1983, Whittaker and Sidner 1996, Whittaker and Hirschberg 2001]. Search allows retrieval without such manual organization and improved search should lead to a reduced use of filing strategies in preparation for later retrieval.

The goal of the present study was to test these two predictions.

2.3 Navigation or Search: Prior Evidence Pertaining to the Debate

Evidence concerning users' *search* preference comes from empirical studies that examine retrieval behaviour. An early paper concerning users' retrieval habits [Barreau and Nardi 1995], combined Barreau's interviews of novice personal computer users (using DOS, Windows 3.1 and OS/2) with Nardi's interviews of experienced Macintosh users. In both cases, users "overwhelmingly" preferred to navigate to their files than to search for them. Similar preferences for navigation were obtained in other more recent studies [Boardman and Sasse 2004; Capra and Pérez-Quiñones 2005; Kirk et al. 2006; Teevan et al. 2004]. These early findings raise a question— if search better suits users' requirements, why do they prefer navigation? One argument is that search technology is still immature. For example, Fertig and his colleagues [1996] argued that these navigation preferences result from limitations in search technology, and that improvements in search would inevitably

lead to the replacement of navigation. They noted that the PIM search engines of that time (the mid 90s) were “slow, difficult, or only operate on file names (not content)” and did not provide incremental indexing. Fertig et al. further speculated that “inclusion of these better search techniques into current systems could sway results” [p. 67]. However, their claim that the improvement of search engines would lead to an increased preference for search over navigation has not been tested empirically.

Other evidence challenging the effects of improved *search* concerns users’ organizational efforts to prepare for future retrieval. There is some evidence that users seem to want to preserve folders, even when improved search is possible. Jones, Phuwanartnurak, Gill, & Bruce [2005] asked 14 participants the following question: “Suppose you could find your personal information using a simple search rather than your current folders.... Can we take your folders away?” Only one participant responded positively. In contrast, Dumais et al.’s [2003] participants tended to mildly agree with the sentence “I would likely to put less effort into maintaining a detailed set of folders for my files if I could depend on SIS (i.e., the *Stuff I’ve Seen* search engine) to find what I am looking for”. Both studies asked whether the use of improved search engines would lead to less reliance on folders, but (perhaps because Jones et al. asked the question in a more extreme way) received different answers. Notice, however, that both researchers asked this as a hypothetical question.

2.4 Improvements in Desktop Search Engines

Today, more than a decade after Fertig et al.’s [1996] claims, commercial PIM search engines have improved considerably, newer search engines (such as *Google Desktop* and *Spotlight*) are better than the older ones (such as *Windows XP Search Companion* and *Mac Sherlock*) in the following ways:

Cross-format search: One limit of older search engines was that they allowed users to search only one format at a time. Following the SIS [Dumais et al. 2003] initiative, several improved search engines now support search across multiple datatypes – files, emails, instant messages and Web history within the same search query. This allows them to address the project fragmentation problem, where information items related to the same project but in different formats, are stored in different locations [Bergman et al. 2006].

Faster retrieval: Improved search engines are substantially faster than old ones. In some cases they have been demonstrated to be 1000 times faster [Farina 2005; Lowe 2006].

User-centred design – Choosing between formats was not the only step the user had to take in older search engines. In addition, users had to choose between file name search or full text search, and also optionally specify the time the file was recently modified. To achieve a reasonable retrieval time, the user needed to input more information in order for the computer to do less, a feature which reflects a machine-oriented design. Newer search engines’ retrieval speed allows them to reduce the query launching steps and complications to a minimum.

Incremental search: One advantage of newer search engines is that they support incremental search, so that the search begins as soon as the user types the first character of the query. This has the benefit of being interactive: allowing users to refine their query in light of the results returned, and truncate the query after typing just a few characters if the target item is already in view. Older search engines were less efficient: prompting the user via form filling to specify multiple attribute fields and hit carriage return before the query is sent off. Incrementality, according to Raskin [2000], has several advantages: (a) user and computer do not have to wait for each other, (b) users know they have typed enough to disambiguate their query because the desired file appears in the display, (c)

users receive constant feedback as to the results of the search – they can correct spelling mistakes or refine search words without interrupting the search.

Given these improvements in desktop search engines, it is now time to examine their implications: What are users' file retrieval preferences, what motivates retrieval by search, and what is the effect of improved desktop search engines on file retrieval preferences and file organization?

If the availability of these improved desktop search engines leads to a substantial increase in search, then it is reasonable to assume that this effect will continue to grow as search engines improve. If, on the other hand, no such effect is found, it raises questions regarding claims that improved search engines affect retrieval preferences and file organization, though it always can be claimed (following Fertig et al [1996]) that future improvements in search could change this. As search engines are consistently improving and will continue to do so, the examination of their implications on PIM should be a continuous effort.

In the present study, we focus on *file* retrieval as opposed to emails or web information, as files are considered by most users to be their main collection [Bergman 2006; Boardman 2004].

3. RESEARCH QUESTIONS

3.1 File Retrieval Preferences

Our first questions related to users' overall preferences for search versus navigation of personal files.

1.1 Which retrieval option is used more by personal computer users – navigation or search?

1.2 Are background variables (age, years of use of personal computers, evaluated experience and hours of daily use) related to search usage?

3.2 The Effect of Improved Desktop Search Engines on Retrieval Preferences

This question examines the assumption that improvements in search engines will lead to increased use of search and a reduced reliance on navigation.

2.1 Does the availability of an improved desktop search engine lead to an increase in search usage?

3.3 Motivation for Search

We next explored the motivation for search and in particular tested Barreau and Nardi's [1995] observation that search is typically used only as a 'last resort' when users cannot remember the location of their files.

3.1 How often do users fail to remember the location of a file they intend to retrieve?

3.2 Do users typically use search engines when they don't remember a file location?

3.3 Are background variables (age, years of use, evaluated experience and hours of daily use) related to remembering file location?

3.4 Improved Desktop Search Engines and Organizational Strategies

Finally, we explored the expected effects of improved desktop search engines on file organization; namely, that improved file retrieval via search removes some of the need for filing, so that users expend less effort in creating and maintaining systematic file organization.

4.1 Do improved search engines lead to reduced efforts to construct systematic hierarchical storage?

4. METHOD

Our research consists of a preliminary study and two main studies. All used the same basic questionnaire (described in section 4.3),² with variations according to the system being evaluated. The two main studies tested the research questions using different search engines (and different operating systems) to allow for wider generalization of the results. We used different research designs, and each was run on a different population in order to increase the convergent validity of our findings.

We wanted to test real users' retrieval behaviours with systems they employed on a daily basis. We therefore conducted a 'natural experiment' by investigating usage of readily available commercial software rather than research prototypes. This has three advantages: (a) it allows for relatively larger samples of participants, (b) testing can be done in their natural working environment, and (c) users are less prone during testing to biases arising from their inferences about the purpose of the experiment and their desire to please the experimenter.

In the next section (4.1), the three studies are briefly described. Table I shows the key design features of each. Section 4.2 describes the different search engines used. Section 4.3 describes our questionnaire, and Section 4.4 addresses the validity of our measurements.

4.1 Design, Procedure and Participants

Our preliminary study provided a baseline to test the retrieval preferences of current personal computer users, most of whom do not use state of the art search engines. The two major studies test the effect of the *improvement* in desktop search engines on preferences for search or navigation.

Preliminary Study - Retrieval Preference in Current Systems: Participants were 78 Israeli personal computer users using the Windows XP operating system with an old search engine – *Windows Search Companion* (WSC). Their ages ranged from 21 to 57 ($M = 30.44$, $SD = 7.86$), they had used a personal computer for between 0.5 to 22 years ($M = 8.83$, $SD = 4.99$), and they indicated a varied degree of computer experience ($M = 3.54$, $SD = 2.01$ on a 1-5 Likert scale).

Windows Study - Retrieval Preferences with old and improved search engines using Windows: The old search engine was Windows XP Search Companion (WSC), while the new one was Google Desktop (GD). This was a longitudinal within-subject design which manipulated the kind of search engine: The improved search engine was installed in participants' computers. The goal of the study was to investigate the changes, if any, in retrieval preferences (search vs. navigation) as a result of using an improved search engine.

The participants were 47 Israeli Windows XP users who used to work with the conventional WSC engine. Their age ranged from 15 to 55 ($M = 34.65$, $SD = 10.88$), they had used personal computers for between 1 and 26 years ($M = 11.24$, $SD = 6.08$), and their average self-evaluated computer experience was 3.62 ($SD = 0.76$) on a 1-5 Likert scale. The questionnaire asked them to assess their WSC usage, after which GD was installed on their computers. After ensuring that the participants knew how to use it, they were instructed to use it for the next three weeks instead of their old search engine. At the end of that period, the questionnaire was re-administered, this time focusing on the new search engine. After a period of 7 months, participants answered a third, follow-up, shorter³ questionnaire to test long term effects on search preference ($N = 43$).

² Available on request from the first author.

³ To ensure a high response rate, after 7 months, we decided to repeat only the most important questions. See section 4.3.

Mac Study – Retrieval Preferences with old and improved search engines for Mac: The old search engine was *Sherlock* (which is part of Mac OS X versions 10.0-10.3) while the new one was Mac *Spotlight* (which is part of Mac OS X version 10.4). The goal of the study was to investigate the same question as that of the Windows Study, this time using Mac search engines. This study had a between-subjects design. We used a Web survey answered by 589 Mac user group members in the US, UK and Australia. Respondents were divided into two groups depending on their specified Operating System: *Spotlight* search engine users (for Mac OS X 10.4) and *Sherlock* users (Mac OS 10.0-10.3)⁴. The *Spotlight* user group consisted of 519 participants; the *Sherlock* group, 70 participants.⁵

The *Spotlight* group's age ranged from 15 to 87 ($M = 58.39$, $SD = 14.81$), they had used a Mac for between 0.5 and 24 years ($M = 14.93$, $SD = 7.34$) and on a scale of 1-5 they evaluated their computer experience on average to be 4.3 ($SD = 0.88$). The *Sherlock* group's age ranged from 15 to 93 ($M = 54.9$, $SD = 16.43$), they had used a Mac for between 0.5 and 24 years ($M = 13.73$, $SD = 6.52$) and on a scale of 1-5 they evaluated their computer experience on average to be 3.85 ($SD = 1.05$). As Mac OS X 10.4 was released two years before the study took place it can reasonably be assumed that most of those who used the improved search engine had some experience with it at the time of the study.

Table I: Key features of the three studies

<i>Study</i>	<i>Preliminary study: PC baseline</i>	<i>Windows Study: Effects of using GD for PCs</i>	<i>Mac Study: Effects of using Spotlight for Mac</i>
<i>Features</i>			
<i>Population</i>	Israeli personal computer users	Israeli personal computer users	US, UK and Australian Mac users
<i>Size of sample</i>	78	47	589
<i>Old Search Engine</i>	WSC	WSC	Sherlock
<i>Improved Search Engine</i>	-	GD	Spotlight
<i>Research design</i>	One group	One group, within-subject longitudinal design	Two groups, between-subject design
<i>Questionnaire administration: Number of times to which group/s</i>	One administration to one group	Three administrations to the same group: before manipulation, after it, and follow-up	One administration to each one of two groups
<i>Questionnaire Administration procedure</i>	Paper and pencil	Paper and pencil	Web survey

4.2 Search Engines Used in the Studies

This section briefly describes the search engines used in our studies. WSC and *Sherlock* are older search engines, while GD and *Spotlight* belong to the newer improved generation. The main differences between the two generations were described in section

⁴ Because of the between-subject design, participants could not guess the aim of the study and respond accordingly.

⁵ The study excluded 34 participants who did not specify their operating system, and 5 participants who used older operating systems (now termed Mac Classic), as these may be substantially different from the Mac OS X 10.4.

2.4, namely cross format search, speed, user-centric design and incremental querying as presented in Table II.

Table II: Features of the Search Engines Used in the Study

<i>Feature</i>	<i>Cross-format</i>	<i>Fast</i>	<i>User-centric</i>	<i>Incremental</i>
<i>Search Engine</i>				
<i>WSC</i>	-	-	-	-
<i>Sherlock</i>	-	-	-	-
<i>GD</i>	+	+	+	+
<i>Spotlight</i>	+	+	+	+

In addition, each search engine also has its own characteristics as described below.

Windows XP Search Companion (WSC): Windows XP was released in 2001 and its search engine is only slightly different from the WSC in previous Windows versions. In addition to the lack of features listed in Table II, WSC has two further problems – invisibility and inaccessibility. WSC is hidden within the *Start* menu and it takes at least five steps before the user can start the query (*Start->Search->For Files and Folders->File format (e.g. Documents) -> Click in query box*) and there is no keyboard shortcut for it. In contrast, in the other search engines tested, the search option is visible on the users’ screen so they are constantly reminded of it and can access it with a single mouse click.

Sherlock: Apple introduced the *Sherlock* search engine in 1997 in the Mac OS 8.5 and its architecture was not greatly changed in the OS 10.0 - OS 10.3 version used in our study. The *Sherlock* query box is located at the top of the *Finder* – Mac’s file navigation apparatus – and therefore is highly visible (in addition it can be reached through the *File* menu, or by using a keyboard shortcut), and it requires fewer steps than WSC to begin a search.

Google Desktop (GD): *Google Desktop* was released in 2005. The GD query box appears on a sidebar or a desktop bar and is always available to users. GD presents its results either in a list or within the browser in a similar way to the Google Web search engine. It also allows cross format search. GD also profits from the popular Google Web search engine: when launching a Web query using Google, the query is also sent to GD and its local desktop results are blended with the Web results (Russell and Lawrence 2007).

Spotlight: *Spotlight* was introduced in the Mac OS X version 10.4 (also called Tiger), which was released in 2005. The *Spotlight* dialog box replaced *Sherlock* in the *Finder*, menu and keystroke shortcut. In addition to different information item formats, its cross-format search also searches for software applications.

4.3 The Questionnaire

Participants were first asked about their retrieval habits: “You can retrieve a file which you have saved on your computer in different ways: by navigating to the folder in which it is located and selecting it from the files there, by clicking on a shortcut on your Desktop if you created one, by using the *Search* option (located in the *Start menu* and in the *Folders menu*), by selecting it from a recent documents list (either in the *Start menu* or from the applications in the *File Menu*), or in other ways. Think of all the occasions when you retrieved your files (e.g., in the last week). Please try to assess the percentage of times you used each of these methods.” Participants estimated the percentage of use of each of the specified options. The estimated percentages of all options (including “others”) had to amount to 100%. For the Mac questionnaire, we used Mac terminology and a sixth retrieval option was added – *Smart Folders*. Two measures were extracted from the responses: the *estimated search percentage* and the *estimated navigation percentage*.

Next, in order to identify the motivation for search, in the Windows Study (stages 1 and 2) and in the Mac Study we asked two questions: a qualitative open question followed by a quantitative open one. The qualitative question was: "Please specify when you usually use the search option to retrieve your files." The quantitative question was: "Let us now focus on the events during which you retrieved your files by using the Search option. These are now 100% of your events. Divide the 100% between the following two categories: In ___% of Search events, I didn't know exactly where the files were located (prior to the search), and in ___% of Search events, I knew where the files were located (prior to the search)." In our pilot we experimented with various phrasings but settled on the final wording. In particular the question was conservative in assessing knowledge by requesting whether users knew the file's *exact* location. In the Mac Study, this quantitative question was preceded by the following question: "Do you use the search option for your files when you know where they are located?", where they could select "never", "sometimes" or "always". Participants who answered "never" were not asked the quantitative question and their answer was set to 100 and 0, respectively, for the first and second parts of the question.

We also wanted to know how often participants forgot where a file was located. To get an estimate of the percentage of *retrieval occasions* in which participants could not remember where a file was located, in the Windows Study (stages 1 and 2) we asked: "Think of all the occasions on which you retrieved a file. Try to estimate the percentage of events in which you didn't know exactly where the file was located." We labelled this variable *estimated unknown locations percentage*. Note that this question is independent of retrieval method (i.e. whether the user navigated or searched).

In the questionnaire given in stage 3 of the Windows Study and in the Mac Study, participants were asked two additional questions regarding changes in their file organization habits using the improved search engine. The first one was: "Do you think that you changed the way you organize your files after changing to your current operating system? (If this is your first Mac operating system, skip this question.)". Those who answered "yes" were asked to specify what those changes were.

The last questions collected data about participants' age, years of use of their personal computer, how experienced they felt with that computer (on a scale of 1 to 5), and how many hours a day they used the computer.

4.4 Validity of Estimates

Our research clearly relies on participants' ability to estimate their retrieval preferences. One issue is the extent to which such estimates reflect actual behaviours. To validate these estimates, we compared them to *actual retrieval behaviours* in two ways: (a) to user reports of how they retrieved specific files, and (b) to analyses of logfiles.

We asked a sample of 16 participants who had been using GD for 3 weeks to print out their Recent Documents. Each Recent Documents printout contained 15 files making a total of 240 files. Then the interviewer asked the participants how they had retrieved each of the files. Note that although the Recent Documents list contains files opened in the recent retrieval activities, when these activities took place the users did not necessarily retrieve recently opened files. The files in the list could have been retrieved at any time in the past. Therefore the procedure is not biased in terms of temporal selection. We first excluded various items: items that were not retrieved by that person for people with shared systems (30 files); items that were not stored as files e.g. an attachment to an email (21 files); and finally items where users did not remember the way they were retrieved (21 files). The percentage of the recent retrievals which people reported were retrieved using search was calculated for each participant. We examined how this data (the *actual search percentage*) corresponded with the percentage estimated in the questionnaire, allowing for within-subject validation. The average *actual search*

percentage was 11% (SD = 26%), which is slightly lower than the average *estimated search percentage* for these participants, which was 15.5% (SD = 25%). The Pearson correlation between user estimates and actual search behaviour was extremely high $r=0.94, p<0.01$.

As an independent validation, we also examined the relation between users' estimations and logfile usage data. We sampled the command lines of 11 Linux users, who we also asked the same estimation questions. We employed Linux users because it is straightforward to collect logs of retrieval behaviours. The fact that we used a different OS (with potentially different levels of search) to test the validity of our method was irrelevant because we were not directly interested in absolute levels of search but rather *the relation between reported estimates and actual behaviours*. The command lines included 190 file retrieval commands. We calculated actual search usage as a percentage of total file retrievals using other navigational methods such as listing files. The results from the sampled command lines show an average *actual search percentage* of 15% (SD = 13%), which is again slightly lower than the average Linux participants' *estimated search percentage* – 20% (SD = 14%).

Together these results indicate that users' estimations of their search percentage are accurate and valid.

5. RESULTS

Table III presents the average percentages given to each of the retrieval options in Question 1 for all studies and conditions. These data address the first two research questions: current baseline retrieval preferences and the effect of improved search engines on retrieval preferences.

Table III: Overall Estimated Percentages for Different Retrieval Options – *Mean and (Standard Deviations)*

	<i>Search</i>	<i>Navigation</i>	<i>Shortcuts</i>	<i>Recent Documents</i>	<i>Smart Folders</i>	<i>Others</i>	<i>Total</i>
<i>Prel. study WSC (N=78)</i>	4% (9%)	64% (32%)	19% (24%)	12% (19%)	n/a*	1% (4%)	100%
<i>Windows Study WSC (N=47)</i>	7% (14%)	68% (27%)	19% (22%)	6% (8%)	n/a	0% (1%)	100%
<i>Windows Study GD (N=47)</i>	15% (24%)	63% (30%)	17% (23%)	5% (10%)	n/a	0% (0%)	100%
<i>Windows Study GD follow-up (N=43)</i>	10% (16%)	63.5% (32%)	20% (26%)	5.5% (9%)	n/a	1% (3%)	100%
<i>Mac Study Sherlock (N=70)</i>	15% (22%)	58% (31%)	18% (24%)	6% (11%)	1% (5%)	2% (9%)	100%
<i>Mac Study Spotlight (N=519)</i>	13% (17%)	56% (30%)	16% (22%)	11% (16%)	1% (5%)	3% (11%)	100%

* n/a = not applicable

5.1 File Retrieval Preferences

Which retrieval option is used more by personal computer users – navigation or search? To evaluate retrieval preferences, the estimated search percentage and the estimated navigation percentage were compared in all studies and conditions (six altogether). Most retrieval is done using navigation (ranging from 56 to 68% of the retrieval events). Mean estimated search percentages were much lower: 4-7 % for WSC users and 11-15 % for GD, Sherlock and Spotlight users. Paired t-tests comparing these two estimates revealed a significant preference for navigation over search for all studies and conditions (see Table IV).

Table IV: Paired t-tests between estimated search percentage and estimated navigation percentage

	<i>Df</i>	<i>t</i>	<i>p</i>
Prel. study WSC	77	10.87	<0.01
Windows Study WSC	46	11.83	<0.01
Windows Study GD	46	6.62	<0.01
Windows Study GD follow-up	42	8.55	<0.01
Mac Study Sherlock	69	7.60	<0.01
Mac Study Spotlight	518	24.95	<0.01

Are background variables related to search usage? In all studies and all conditions, we calculated the correlation between each of the four background variables (age, years of use, evaluated experience and hours of daily use) and the *estimated search percentage*. Only one of 24 Pearson correlations was significant: between *estimated search percentage* and *age* for *Spotlight* users ($r=-0.12$, $p<0.01$). Notice that with large samples even a relatively low correlation might be significant. Overall, therefore, the *estimated search percentage* was found to be rather stable and generally unaffected by background variables.

5.2 The Effect of Improved Desktop Search Engines on Retrieval Preferences

Does the availability of an improved desktop search engine lead to an increase in search usage? This question was tested by comparing the estimated search percentage in the improved search engine condition (GD/Spotlight) with the estimated search percentage in the older search engine condition baseline (WSC/Sherlock) in Studies 1 and 2. The estimated percentage of all search options is presented in Table III. We will now report separately on the results of the Windows Study and the Mac Study.

The results of the Windows Study are shown in rows 2 - 4 of Table III. Three weeks after the initial installation, the *estimated search percentage* when using GD ($M = 15.5\%$) was significantly higher than the initial WSC baseline before GD was installed ($M = 7\%$), $t(46)=2.76$, $p<0.01$). However, in a follow-up questionnaire conducted 7 months later, the estimated search percentage had dropped back ($M = 10\%$). This final estimate was not significantly higher than the baseline ($t(42)=0.81$, $p>0.05$). This may be partially explained by the fact that 16 of the 43 participants (37%) who answered the follow-up questionnaire had removed GD from their computer. The *estimated search percentage* of the remaining 27 participants was 14.5% ($SD = 19\%$) which is significantly higher than the baseline using the WSC search engine ($t(26)=2.13$, $p<0.05$).

When these 16 people who uninstalled GD were asked for their reason for doing so, 7 answered that they had no use for it, 4 claimed that it slowed down their computer, 3 uninstalled it “accidentally” and did not reinstall, one participant reported technical problems and another replaced it with an alternative search engine. With the exception of

the final user, these data suggest that for these people there was no compelling reason to use desktop search.

Thus, for those participants who persisted with the new search engine, there was an increase in search frequency. However these adopters are counterbalanced by the other set of users who (with one exception) did not see strong benefits for search, and uninstalled the system. The fact that retrieval by search did not increase over time, refutes possible claims that after 3 weeks of use the increase in *estimated search percentage* was small due to inexperience with the new engine.

The results of the Mac Study are presented in the last two rows of Table III. This study compares the retrieval preferences of Mac users using *Spotlight* (improved search engine) to the retrieval preferences of Mac users who use *Sherlock* (old search engine). Contrary to expectations, the results show that the average *search percentage* of the improved search engine (13%) was lower than that of the old search engine users (15%), although this difference was not statistically significant: $t(587)=0.95, p>0.05$.

The Mac also offers a retrieval option called *Smart Folders*. *Smart Folders* is a hybrid that combines the search and the navigation options: the user saves the search query results in a dynamic folder and navigates to that folder in order to retrieve the query. The search results are automatically updated in the folder. The results presented in Table III indicate that Mac users make only marginal use of this option.

Overall, the data presented in this section do not provide strong evidence that improved search engines change participants' retrieval behaviour significantly.

5.3 Motivation for Search

Table V summarizes all data relevant to Search Motivation (questions 3.1 and 3.2).

Table V: Search and Memory for File Location

	<i>Average estimated unknown file location as a percent of all retrievals (%) M (SD)</i>	<i>% of subjects indicating using search when they "can't remember location of file"</i>	<i>Of all search retrievals, average estimated percent of events when location of file was known M (SD)</i>	<i>Of all search retrievals, average estimated percent of events when location of file was unknown M (SD)</i>	<i>Of all retrievals, average estimated percent of search when location of file was known M (SD)</i>
<i>Windows Study – WSC</i>	10% (15%) (N=47)	82% (N= 35)	2% (4%) (N=30)	98% (4%) (N=30)	0% (0%) (N=30)
<i>Windows Study – GD⁶</i>	15% (23%) (N=47)	79% (N=34)	9% (20%) (N=39)	91% (20%) (N=39)	4% (11%) (N=39)
<i>Mac Study – Sherlock</i>	24% (29%) (N=67)	73% (N=55)	14% (26%) (N=70)	86% (26%) (N=70)	3% (11%) (N=70)
<i>Mac Study - Spotlight</i>	26% (26%) (N=515)	75% (N=443)	17% (27%) (N=519)	83% (27%) (N=519)	3% (9%) (N=70)

⁶ Second questionnaire, after three weeks of GD use.

How often do users fail to remember the location of a file they intend to retrieve? We asked participants to estimate the percentage of retrieval events in which they did not remember exactly where the file was located. Column 1 of Table V presents these percentages. The results show an average of 10-26%. Search engine quality does not affect these estimates; there were no significant differences in *evaluated unknown location percentage* between the improved search engine and the older search engine conditions: $t(46)=1.81, p>0.05$ and $t(580)=0.52, p>0.05$ for Studies 1 and 2, respectively.

Do users typically use search engines when they don't remember a file's location? To pursue this research issue, we used an open question and a quantitative estimation question.

The open question was: "Please specify when you usually use the search option to retrieve your files". Column 2 of Table V presents the percentage of responses when participants spontaneously answered that they search when they don't remember where their files are located (also phrased in a minority of cases "when I can't find it", "when I have no choice" and "when I cannot remember its path structure, Duh?"). Inability to remember the file location was the dominant reason for search – accounting for over 75% of responses, with no statistical differences between older and improved search engines.

The third and fourth columns in Table V present the data (from Studies 1 and 2) regarding the quantitative estimation question in which participants were asked to think about all search events and estimate the percent of events when location of file was known (in Column 3) and the percent of events when its location was unknown (in Column 4). Users estimated that in a large majority of search events, they did not know where their files were located.

In the last column, for each participant, we multiplied the "percentage of retrievals when the location of the file was known, as a percentage of all search retrievals" (the mean is presented in Table V, Column 3) by the estimated *search percentage* (the mean is presented in Table III, Column 1). The results show that only a tiny percentage of retrievals were searches for files at a known location (0-4%).

These results show that personal computer users tend to use search mainly when they cannot remember a file's location.

Are background variables related to remembering file location? For all groups and conditions, we calculated the correlation between each of the four background variables (age, years of use, evaluated experience and hours of daily use) and the estimated percent of retrieval events in which the location of the file was unknown. None of the 20 correlations was found to be significant. Thus the *estimated unknown locations percentage* was found to be rather stable and unaffected by background variables.

5.4 Improved Desktop Search Engines and Organizational Strategies

Our final research question was "*Do improved search engines lead to reduced efforts to construct systematic hierarchical storage?*" This question was tested using a direct question (presented in the follow-up questionnaire of the Windows Study and to the *Spotlight* group in the Mac Study) and by measuring the effect of improved search engines on *the estimated unknown file locations percentages*.

Of the 43 participants who answered the follow-up questionnaire in the Windows Study, none reported to have changed the way they organize their files due to their use of GD. Of the 481 improved search engine users in the Mac Study, a majority (386, 80%) responded that they didn't change the way they organized their files after starting to use improved search engines. Of the 95 participants who said they were organizing files differently (20%), only 12 (i.e. 2.5% of all participants) reported to be less organized or to use fewer, more general folders. Indeed, another 12 participants reported being *more* organized or using deeper hierarchies with more subfolders with improved search.

Improved search engines should cause users to rely less on their file organization, allowing them to file more casually, because they are confident that they can retrieve files via search. This expected reduction in organization should also lead to an increase in the estimated *unknown location percentage* both because the actual filing can be more general and less consistent and because filing actively is likely to improve memory [Kidd 1994]. However our results show no significant difference in these estimates between users of old and improved search engines: $t(46)=1.81$, $p>0.05$ for the Windows study, and $t(580)=0.52$, $p>0.05$ for the Mac study.

To conclude this section, our studies show little evidence for the claim that improved search engines lead to less use of hierarchical storage.

6. DISCUSSION

6.1 Main Findings

Our overall results suggest that improvements in commercial search engines don't seem to affect retrieval and storage. Regarding retrieval, more than a decade after Fertig et al. [1996] predicted that the prevalence of improved search engines would lead to an increased preference for search over navigation, our results indicate that they do not. Regarding storage, two decades after Lansdale [1988] suggested search as an alternative to the hierarchical method, our results show that improved search engines do not lead to reduced reliance on folders.

Our results show a strong overall preference for navigation. Retrieval preferences were tested with *Windows Search Companion* (WSC), *Google Desktop* (GD), *Sherlock* and *Spotlight* search engines users. Users estimated that they used navigation for a majority – between 56% and 68% – of their retrieval events. The average *estimated search percentages* were much lower: 4-7 % for WSC users and 11-15 % for GD, *Sherlock* and *Spotlight* users. These results confirm the observations made by Barreau and Nardi [1995], and more recently by others [Boardman and Sasse 2004; Capra and Pérez-Quiñones 2005; Kirk et al. 2006; Teevan et al. 2004]. In addition, the evaluated search percentage seems to be quite stable and unaffected by age, years of use, evaluated experience, and hours of daily use (with the exception of a single low correlation of $r=-0.12$, $p<0.01$ between age and *estimated search percentage*, only for *Spotlight* users).

In two different research designs for Windows and Mac, we tested the main research question: whether improved search engines lead to more search in absolute terms and also relative to navigation. In the Windows Study, the installation of GD significantly increased the *estimated search percentage* of WSC users from 7% to 15% on average after three weeks of use. However, seven months later, the *estimated search percentage* dropped back to 10% (which was not significantly higher than baseline). The Mac Study, which had a larger number of participants (589), found no increase in search at all when using an improved search engine. Altogether, the results show a limited and inconsistent effect of improved search engines on retrieval preferences, which do not affect the preference for navigation over search.

A possible interpretation of the results is that the visibility and accessibility of the search engines, and not their quality, affected the estimated search percentages. In WSC where the search option is not visible (thus users are not reminded of it) and can be accessed through a minimum of five steps, the *estimated search percentage* is very small (4-7%). Under this condition, participants might not be familiar with the search option or simply not “in the habit” of using it as “out of sight” is often “out of mind”. However, in the other search engines (old and improved) in which the search option is visible and thus more easily accessible, the estimated search percentage was higher, around 13-15%, regardless of the other features which differentiate old search engines and improved ones (presented in section 2.4). Regardless of this interpretation, it is clear that the use of

improved search engines did not challenge navigation preference. Possible reasons will be discussed in the next section.

These results are also consistent with navigation preferences in highly advanced experimental search engines. *Stuff I've Seen* (SIS) [Dumais et al. 2003] allows fast, incremental, cross format searches with a sophisticated UI that supports complex presentation of search results. When we consider files only (the focus of our current study), participants searched for files using SIS an average of about once every two days. The same research group tested an even more advanced search engine called *Phlat* that also supported file tagging [Cutrell et al. 2006b]. In their test, 225 participants used *Phlat* for 8 months and performed 5,144 query sessions – an average of less than one query session per week per participant. As the participants (Microsoft employees) were likely to retrieve many files each day, it would be safe to conclude that neither search engine led to preference for search over navigation, or to a dramatic increase in search.

Why, then, is search so infrequent? One reason is that search is only used when people *forget a file location*. Our results confirm Barreau and Nardi's [1995] observation that search was typically used as a 'last resort' when users couldn't remember a file's location. Most users (between 73% and 82% for the different studies) spontaneously noted that they usually searched when they did not know where their files were located, and estimated that the majority of their searches were performed in this case (an average range of 83-97% for the different studies). According to participants' reports, only a tiny percentage of retrievals were searches for files whose location was known (from an average of 0% for Windows XP to an average of 3% for Mac). One way of thinking about this might be that the percentage of files for which the users *don't* remember the exact location can be regarded as an approximate upper limit to search percentage growth. This estimate, of about 25%, was not affected by the use of improved search engines, or by background variables. Notice that for all conditions, the average *evaluated search percentage* is lower than the average *evaluated unknown locations percentage*. This means that often when participants didn't know the file location, they still preferred *not* to search for them, but to retrieve them in other ways.⁷ Therefore, a less ambitious, but perhaps more realistic, challenge for search engine developers could be that users will search for the 25% of files whose location they can't remember.

In addition, there was also little to indicate that improved search engines lead to less reliance on hierarchical storage. Only 12 of 481 *Spotlight* users reported that an improved search engine resulted in their being less organized or using fewer, more general, folders. Indeed, another 12 participants reported they became more organized or used deeper hierarchies with more subfolders. Furthermore, there was no significant difference in *evaluated unknown location percentage* between new and older search engines, thus there was no indication that the use of improved search engines led people to be more casual about how they organized information and hence have a poorer memory for where information was located. Altogether, these results confirm the results of Jones et al. [2005] and contrast with those of Dumais et al. [2003], as they show no evidence that improved search engines reduce reliance on hierarchical storage.

6.2 Possible Explanations for Navigation Preference over Search

One possible explanation for the observed preference for navigation could be that provided by Fertig et al. [1996] more than a decade ago. One might argue that although search engines have improved greatly in that time, (now having cross format, incremental, rapid search, and user centred design), nevertheless they are *still* not

⁷ An estimate of the percentage of search retrievals out of the unknown file retrievals can be calculated, for each participant, from other estimates they gave. This was done for both conditions in the Mac study. On average, when the location of files was unknown, for *Spotlight* and older Mac search engines only 35% and 42% of the retrievals used search engines, respectively.

sophisticated enough. When more advanced search engines are deployed, search will then be preferred to navigation. Although this remains a possibility, the burden of proof lies with those who make this argument. If our results had shown that the use of improved search engines resulted in a substantial increase in search it could be reasonably assumed that further improvements would eventually lead to the dominance of search. However, for the improved search engines we examined, there were few increases in search estimates; and search was generally restricted to cases where people could not remember file locations, which seemed to have a threshold of about 25%. The present results therefore provide some support for our view that in current PIM environments, there is an inherent preference for navigation, irrespective of the search engine's sophistication. Here are a number of possible theoretical explanations:

Consistency – It is widely agreed that consistency is a virtue in design and human-computer interaction as it confirms user expectations [Shneiderman 1997]. The hierarchical method is boringly consistent – the files are stored in a “location” and stay there until decided otherwise by the users, who can expect to find them in the same place at retrieval. The process of navigation may require more steps than search; but users consistently use those same steps. In contrast, the flexibility of search may compromise consistency as users are able to retrieve the same file using a different search term. Even when using the same search word, the same file might appear in different places in the query results list at different times. Consistency can also explain why users remembered the location of most of their retrieved files. It is true that they *could have* stored according to multiple different categories, however after choosing one, and retrieving it from its location a few times, they seem to remember “that it is there”.

Recognition vs. Recall – Recognition tasks are easier and require less cognitive effort than recall tasks [Tulving and Thomson 1973]. When users perform a search, they typically have to generate a set of relevant search terms. The choice of terms is sometimes challenging, as it requires users to *recall* file names or attributes. In contrast, navigation is based mainly on *recognition* – where each step going down the hierarchy provides incremental visual and contextual feedback about navigation success as well as clues about the next choice of folder [Teevan et al. 2004]. This concerns *information addressability*: Both search and navigation require that partial information be supplied but differ in how the information is provided (at once, or over the course of an interaction) and whether incremental feedback is supplied back to the user. It may therefore be that our results represent a preference for a style of interaction based around specifying partial information incrementally.

Procedural vs. Declarative Memory – Specifying a search term forces users to rely on declarative memory; i.e., users have to *know* that this word appears in that file. However, folder navigation can also rely on procedural memory – users have to know *how* to navigate to that file [Barreau 1995]. It may be, too, that users retain both ‘motor memory’ as well as visual recognition about a file’s location.

Cognitive Automation –When users browse in a familiar self-created and consistent setting, the task may require less cognitive attention and may become automated. In navigation (as opposed to search) users typically don’t need to take their mind off the work they are doing in order to retrieve the file. Navigation may take longer than search but users can continue to think of the project they are working on at the time.

The strength of the location metaphor – Location is only a metaphor for file organization. Files are not “really” inside the folders and directories. There is therefore no a-priori guarantee that the location metaphor is the best possible one [Cutrell, personal conversation based on Lakoff and Johnson 1980]. However, the metaphor seems very natural because of its long use in human development from childhood to adulthood: from a very early age, children store items in physical locations and expect to find them there when required [Piaget 1928]. Regardless of retrieval options, users seem to like to know

where their information items are, which may explain why they persist with locational methods.

Search is a major method by which Internet users retrieve their information; however, in PIM, search seems to be used mainly as a last resort; that is, when there is no other alternative. There is a fundamental difference between the two information environments, which, in our mind, is the source of the difference in retrieval behaviour. In PIM, the same user both organizes and retrieves the information. Users are free to organize their information in a subjective way that suits their requirements. Their folders “represent an emerging, often hard-won, understanding of the information items contained within” [Jones et al. 2005]. Users are therefore familiar with their own folders/directories structure, and become even more familiar each time they navigate to it. This familiarity helps them find their way to the files they are seeking.

The function of folder organisation as serving a sense-making function may help explain why improved search did not change filing habits. Folders may serve as a useful way to conceptualize and understand information regardless of their use for retrieval [Cutrell et al. 2006a; Jones et al. 1995]. Future theoretical and empirical work needs to tease apart these different, but related, reasons for organization.

6.3 Hybrid Methods which Combine Navigation and Search

As there are clear advantages to both navigation and search, there have been constant attempts to combine the two. These efforts are not limited to PIM; hybrid systems that combine browsing and search have been suggested for hypertext environments [Belkin et al. 1993; Frisse and Cousins 1989], expert systems and knowledge bases [Belkin et al. 1995; Croft and Thompson 1987], picture archives [Ame 2001], and the World Wide Web [Hearst 2006; Olston and Chi 2003; Perugini and Ramakrishnan 2006].

In recent years, there have been several attempts to develop novel systems that combine search with navigation qualities (these navigations are not based on folders, but on alternative classification methods). Such systems include *Haystack* [Adar et al. 1999], *Lifestreams* [Freeman and Gelernter 1996], *Presto* [Dourish et al. 1999] and *Placeless Documents* [Dourish et al. 2000]. However, none of these systems has been evaluated to determine whether these hybrid methods are better than navigation and search in a PIM environment. Our data indicate that hybrids may not always be successful. In the Mac Study, we tested a hybrid method called *Smart Folders* and found its usage to be marginal in real settings. However, further tests should be conducted before any conclusions can be drawn.

6.4 Navigation and Orienteering

Teevan, Capra and Pérez-Quñones [2007] discussed folder navigation in the larger context of an information seeking strategy called *orienteering*, in which users take small steps using partial information and contextual knowledge as a guide to gradually reach the target information. This strategy is contrasted with *teleporting* where information seekers try to jump directly to their information target (e.g. in a single search query).

In addition to folder navigation (e.g. for files and emails) where the user opens nested folders, *orienteering* also includes Web browsing where the users progress via hypertext links. Web navigation is different from folder navigation in that users do not create and organize the information, and because the progression is typically nonlinear. Although generally the Web is less familiar to users than their own file organization, users can become familiar with a link trail when they use it repeatedly, or rely on Web design conventions. *Orienteering* can proceed across different collections and include a combination of navigation and browsing; for example, the user may navigate to a mailbox folder to find an email which contains a link to their target Web site, and then navigate to a particular page within that site. Our research confirms previous studies

[Ravasio et al. 2004; Teevan et al. 2004], showing preference for *orienteering* over *teleporting* for personal information retrieval.

7. CONCLUSIONS AND FUTURE DIRECTIONS

Our results show preference for navigation over search regardless of the use of improved desktop search engines. By no means does this imply that improving search engines is a waste of time. If search engines are the PIM ‘fire escapes’, then, in case of emergency, you want to have the best one available. On those occasions where users are frustrated because they fail to remember where they located their files, they should find improvement in search engines to be of great importance. This suggests that experimentation with new techniques of page rank, results presentation and in-place search need to be explored more. However, in light of our findings, there is still no evidence that what prevents users from using search as a preferred retrieval strategy is the “primitive” nature of current search engines; improving them did not change users’ clear preference for navigation in information retrieval.

These observations point to a need to develop more sophisticated models of the relation between navigation, search and retrieval. We have shown there is a strong relation between user *memory for file location* and the use of search, but can we develop more nuanced predictive models about *exactly which* search features help here? Earlier we speculated that that search visibility may induced changes in search preferences, but we also need to develop motivated accounts of how other advanced (presentation, page rank, metadata manipulation) can possibly influence retrieval. Other future research might involve looking at files delivered in email, and whether these are searched for more, as well as the effects of file system depth and complexity.

In closing we would like to suggest two promising directions, which do not rely on search but on navigation - the retrieval strategy which was found to be the dominant method in our study.

As navigation is the preferred retrieval method used by millions of users many times a day, it needs more research and development than is currently focused on it. Shortening navigation time, even by a second, could accumulate to several working years per day for the entire population. For example, Jones et al. [2005] observed that some of their participants, who were using Windows OS, wrote the letter “a” at the beginning of an important file name so that it would remain at the top of the folder file list, and be easily and quickly accessed when needed. Macintosh computers allow users to “place” files manually in their folder, so they can easily organize important files at the top of their folders. Future research may determine which interface features aid navigation.

The user-subjective approach [Bergman et al. 2003], a design approach developed specifically for PIM systems, offers a number of design schemes to assist navigation. It does not attempt to replace the hierarchical method, only to improve it. It doesn’t oppose search, and makes use of this option, however, it also doesn’t rely on it. Instead, the approach proposes a novel direction for PIM systems design. It takes advantage of the fact that in PIM, the person who stores the information items is the same person who retrieves it later on. It suggests that PIM systems should use subjective (user specific) attributes given to the information item during the user-information interaction in order to assist future retrieval. Recently, the approach has received empirical support, and been used to develop seven novel design schemes to improve navigation [Bergman et al. 2008]. Such design schemes are *ProjectFolders*, in which there is only one single folder hierarchy for all information items (e.g., files, emails and Web Favourites) [see also Bergman et al. 2006], and *GrayArea* which is an area at the bottom of the folder into which the users can drag and drop information items of low relevance. Future work will evaluate these designs.

We do not intend this paper to discourage experimental and commercial software developers in their diligent quest to improve search engines. Our results, however, suggest that the expectation that improved search engines will lead to preference for search and to less reliance on hierarchical methods were not realized. Perhaps it is time to explore alternative PIM design directions which focus on navigation and the improvement of human computer interface based on users' perceptions and preferences.

ACKNOWLEDGEMENTS

We thank our participants for their time and efforts, Lilach Geresh for her help with the Windows Study; David Mioducer, Jose San Pedro Wandelmer, Daniela Petrelli, Orly Lahav and Ephrat Beloosesky, our Mac evangelist friends, for their help with the Mac Study; and our anonymous reviewers. This research was partially funded by the European Union Marie Curie Grant, TOK 30008.

REFERENCES

- ADAR, E., KARGER, D. R. AND STEIN, L. A. 1999 Haystack: Per-user information environments In *Proceedings of the eighth international conference on information and knowledge management* ACM Press, Kansas City, Missouri, United States 413-422
- AME, E. 2001. Flamenco image browser: Using metadata to improve image search during architectural design. In *CHI '01 extended abstracts on Human factors in computing systems*, ACM, Seattle, Washington.
- BARREAU, D. K. 1995. Context as a factor in personal information management systems. *Journal of the American Society for Information Science* 46 327-339.
- BARREAU, D. K. AND NARDI, B. A. 1995. Finding and reminding: File organization from the desktop. *SIGCHI Bulletin* 27, 39-43.
- BELKIN, N. J., COOL, C., STEIN, A. AND THIEL, U. 1995. Cases, scripts, and information-seeking strategies: On the design of interactive information retrieval systems. *Expert Systems with Applications* 9, 379-395.
- BELKIN, N. J., MARCHETTI, P. G. AND COOL, C. 1993. Braque: Design of an interface to support user interaction in information retrieval. *Information Processing and Management, Special Issue on Hypertext* 29, 325-344.
- BELLOTTI, V. AND SMITH, I. 2000. Informing the design of an information management system with iterative fieldwork In *Proceedings of the conference on designing interactive systems: Processes, practices, methods, and techniques* ACM Press, New York City, New York, United States 227-237.
- BERGMAN, O. 2006. The use of subjective attributes in personal information management systems. In *School of Education*, Tel Aviv University, Tel Aviv, Israel, 135.
- BERGMAN, O., BEYTH-MAROM, R. AND NACHMIAS, R. 2003. The user-subjective approach to personal information management systems. *Journal of the American Society for Information Science and Technology* 54, 872-878.
- BERGMAN, O., BEYTH-MAROM, R. AND NACHMIAS, R. 2006. The project fragmentation problem in personal information management. In *CHI 2006 conference on human factors in computing systems*, ACM Press, Montreal, Canada, 271-274.
- BERGMAN, O., BEYTH-MAROM, R. AND NACHMIAS, R. 2008. The user-subjective approach to personal information management systems design - evidence and implementations. *Journal of the American Society for Information Science and Technology* 59, 235-246.
- BOARDMAN, R. 2004. Improving tool support for personal information management. In *Electrical and Electronic Engineering*, Imperial College, London, 285.
- BOARDMAN, R. AND SASSE, M. A. 2004. "Stuff goes into the computer and doesn't come out": A cross-tool study of personal information management In *Proceedings of the SIGCHI conference on human factors in computing systems*, ACM Press, Vienna, Austria, 583-590.
- CAPRA, R. G. AND PÉREZ-QUÍÑONES, M. A. 2005. Using web search engines to find and refind information *Computer* 38, 36-42.
- CROFT, W. B. AND THOMPSON, R. H. 1987. I3r: A new approach to the design of document retrieval systems. *Journal of the American Society for Information Science* 38, 389 - 404.
- CUTRELL, E., DUMAIS, S. T. AND TEEVAN, J. 2006a. Searching to eliminate personal information management. *Communications of the ACM* 49 58-64.
- CUTRELL, E., ROBBINS, D. C., DUMAIS, S. T. AND SARIN, R. 2006b. Fast, flexible filtering with Phlat: Personal search and organization made easy. In *CHI 2006 conference on human factors in computing systems*, ACM Press, Montreal, Canada, 261-270.
- DOURISH, P., EDWARDS, W. K., LAMARCA, A., LAMPING, J., PETERSEN, K., SALISBURY, M., TERRY, D. B. AND THORNTON, J. 2000. Extending document management systems with user-specific active properties *ACM Trans. Inf. Syst.* 18, 140-170.

DOURISH, P., EDWARDS, W. K., LAMARCA, A. AND SALISBURY, M. 1999. Presto: An experimental architecture for fluid interactive document spaces *ACM Transactions on Computer-Human Interactions* 6, 133-161.

DUMAIS, S. T., CUTRELL, E., CADIZ, J. J., JANCKE, G., SARIN, R. AND ROBBINS, D. C. 2003. Stuff I've seen: A system for personal information retrieval and re-use. In *Proceedings of the 26th annual international ACM SIGIR conference on research and development in information retrieval*, ACM Press, Toronto, Canada, 72-79.

FARINA, P. A. 2005. A comparison of two desktop search engines: Google Desktop Search (beta) vs. Windows XP Search Companion. In *21st Computer Science Seminar*, Hartford.

FERTIG, S., FREEMAN, E. AND GELERNTER, D. 1996. "Finding and reminding" Reconsidered *SIGCHI Bulletin* 28, 66-69.

FREEMAN, E. AND GELERNTER, D. 1996. Lifestreams: A storage model for personal data. *SIGMOD Record* 25, 80-86.

FRISSE, M. E. AND COUSINS, S. B. 1989. Information retrieval from hypertext: Update on the dynamic medical handbook project. In *Second Annual ACM Conference on Hypertext*, ACM, New York, NY, Pittsburgh, PA, 199-212.

GEMMELL, J., BELL, G., LUEDER, R., DRUCKER, S. AND WONG, C. 2002. Mylifebits: Fulfilling the memex vision. In *Proceedings of the tenth ACM international conference on Multimedia*, ACM, Juan-les-Pins, France.

HEARST, M. A. 2006. Clustering versus faceted categories for information exploration. *Communications of the ACM* 49, 59-61.

JONES, W., PHUWANARTNURAK, A. J., GILL, R. AND BRUCE, H. 2005. Don't take my folders away! Organizing personal information to get things done. In *CHI '05 extended abstracts on human factors in computing systems*, ACM Press, Portland, Oregon.

KIDD, A. 1994. The marks are on the knowledge worker In *Proceedings of the SIGCHI conference on human factors in computing systems: Celebrating interdependence*, ACM Press, Boston, MA, 186-191.

KIRK, D., SELLEN, A., ROTHER, C. AND WOOD, K. 2006. Understanding photowork. In *Proceedings of the SIGCHI conference on human factors in computing systems*, ACM, Montreal, Quebec, Canada.

KOBAYASHI, M. AND TAKEDA, K. 2000. Information retrieval on the web 32, 144-173.

LAKOFF, G. AND JOHNSON, M. 1980. *Metaphors we live by*. University of Chicago Press, Chicago.

LANSDALE, M. W. 1988. The psychology of personal information management. *Applied Ergonomics* 19, 55-66.

LOWE, M. 2006. Evaluation of desktop search applications. Kalio, Sydney, Australia.

MALONE, T. W. 1983. How do people organize their desks? Implications for the design of office information systems. *ACM Transactions on Office Information Systems* 1, 99-112.

OLSTON, C. AND CHI, E. H. 2003. Scentrails: Integrating browsing and searching on the web. *ACM Transactions on Computer-Human Interaction* 10, 177-197.

PERUGINI, S. AND RAMAKRISHNAN, N. 2006. Interacting with web hierarchies. *IT Professional* 8, 19-28.

PIAGET, J. 1928. *Judgement and reasoning in the child*. Routledge and Kegan Paul, London.

RASKIN, J. 2000. *The humane interface: New directions for designing interactive systems*. ACM Press/Addison-Wesley Publishing Co., Boston.

RAVASIO, P., SCH, S. G. AND KRUEGER, H. 2004. In pursuit of desktop evolution: User problems and practices with modern desktop systems. *ACM Transactions on Computer-Human Interaction* 11, 156-180.

RUSSELL, D. AND LAWRENCE, S. 2007. Search everything. In *Personal information management*, W. JONES AND J. TEEVAN, eds., University of Washington Press, Seattle and London, 153-166.

SHNEIDERMAN, B. 1997. *Designing the user interface: Strategies for effective human-computer interaction*. Addison-Wesley Longman Publishing Co., Inc.

TEEVAN, J., ALVARADO, C., ACKERMAN, M. S. AND KARGER, D. R. 2004. The perfect search engine is not enough: A study of orienteering behavior in directed search. In *Proceedings of the SIGCHI conference on human factors in computing systems*, ACM Press, Vienna, Austria, 415-422.

TEEVAN, J., CAPRA, R. G. AND PÉREZ-QUIÑONES, M. A. 2007. How people find personal information. In *Personal information management*, W. JONES AND J. TEEVAN, eds., University of Washington Press, Seattle and London, 22-34.

TULVING, E. AND THOMSON, D. M. 1973. Encoding specificity and retrieval processes in episodic memory. *Psychological Review* 80, 342-373.

WHITTAKER, S. AND HIRSCHBERG, J. 2001. The character, value, and management of personal paper archives. *ACM Transactions on Computer-Human Interaction* 8, 150-170.

WHITTAKER, S. AND SIDNER, C. 1996. Email overload: Exploring personal information management of email. In *Proceedings of the SIGCHI conference on human factors in computing systems: Common ground*, ACM Press, Vancouver, British Columbia, Canada, 276-283.