

Where did I put that? Matching document retrieval to human memory



ABSTRACT

The discipline of human-computer interaction is often driven by attempts to resolve mismatches between the way computers operate and the way people do. The area we have investigated is memory and retrieval. Is it possible to access files on the computer in a way that fits with the way the owner will remember them? We focused on the file-opening dialog as a point of contact between the personal archive and the current activity. Two prototype file-opening dialogs were developed and evaluated. These allowed users to specify files by combinations of characteristics aligned with those which are important in people's semantic and episodic memory. In the second evaluation, participants achieved 100% scores in a retrieval exercise while reporting very low effort required.

Author Keywords

File, activity, open dialog, human memory

ACM Classification Keywords

H.5.2 Information Interfaces and Presentation: Miscellaneous. (D.2.2, H.1.2, I.3.6)

General Terms

Human Factors, Design, Experimentation

INTRODUCTION

In the early days of computer technology, there were those who insisted on using the word “storage”, not “memory” in relation to computers, because they saw memory as a purely human faculty. It may be that now it is accepted to speak of computer memory, but it is nevertheless true that differences in the way that computers store and structure data and the way people remember is still a fundamental barrier to efficient interaction. This is an example of a fundamental mismatch between computers and the people who use them and in most instances it is the user who is expected to compromise, to adapt to the way the computer works, not the other way round. This paper reports an attempt to address

the problem from the other direction – making the computer store and present data in a manner that closely resembles the way people think and therefore facilitating data retrieval.

Human long-term memory has episodic, semantic and procedural components, and generally relies on multiple links to any piece of information. The goal of the current research is to improve the capability of computer systems to take advantage of these rich associations. We focused on the file-opening dialog as a point of contact between the user's personal document collection and their current activity.

LITERATURE REVIEW

According to the model of human memory used by Baddeley [1] organizing and remembering are subjective and personal processes. Human beings try to impose meaning on what they observe, recalling previous experiences as a guide. Experts in particular are very good at structuring and making sense of complex information related to their field of expertise. This might be because information that has been structured previously is easier to remember and also because experts are very good at identifying patterns and organizing material into items with a higher information content.

There are motivations behind the construction of personal archives that go further than simply storing things for later retrieval [9]. Building them also pursues the goals of creating a legacy, sharing resources and reducing fear of loss. They play a big role in self-expression and self-identity. Practical considerations are important to an extent, but these values of legacy, sharing, anxiety and identity construction play a part in defining the archive's structure and should be taken into account while designing it and when evaluating its success.

Research has identified a number of consistent behaviors in the ways that users organize and find files in their computers [2, 8]. In the cited work, location-based filing and recovery of information was preferred in general, with search used very sparsely and mainly only as a last resort. Users arrange their personal workspace and feel more comfortable navigating through a structure that they have created themselves. Location-based filing also provides an important reminding function. Users work with three different types of information: ephemeral, working and archived. The way that information is used (or is expected to be) determines how it will be organized, stored and retrieved. Participants did not tend

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CHI 2011, May 7–12, 2011, Vancouver, BC, Canada.

Copyright 2011 ACM 978-1-4503-0267-8/11/05...\$10.00.

to expend great energy on archiving old information or in creating elaborate filing structures, yet most systems seem to be focused precisely on those two functionalities. Keeping order and proper organization in the workplace was a common problem and many users tended to accumulate a large amount of useless information that did not serve any practical purpose and that only added confusion. This tension between organization for current use and for later re-use has been confirmed and discussed in more recent studies [6].

Folder hierarchies contain information about the items and their relationships; the process of organizing items in folders might in itself help the user understand the information better [7]. Folders are used as a way of decomposing a problem or planning a project, although they have limitations in that they might obscure some information and their lack of flexibility may often force additional information to be squeezed into the hierarchy.

Some of the cited works were carried out at a time when search facilities were rather crude and lacking functionalities. Therefore, some could have predicted that improving the search tools should lead to a substantial increase in their usage and eventually to a preference for search over navigation. However, it was found that this might not be the case and that there is a strong preference for navigation that did not decrease significantly over the years when using more advanced search tools [3]¹. Again, search was mainly used as a last resort when navigation had failed to locate the required file. Possible explanations are that navigation is more consistent, exploits recognition (as opposed to recall), leverages procedural memory (as opposed to declarative), allows for cognitive automation and provides a location metaphor that mimics the natural world.

Different designers have proposed new approaches to personal information management; a significant one is Stuff I've Seen (SIS) [5]. Its goal is to provide unified access to information that a person has used, regardless of its origin, by providing a unified index with all the information that a person has seen and using contextual cues to enrich the search interface. An interesting use of those cues is the annotation of a time-line with personal and public landmarks in time to leverage episodic memory and significantly reduce search times [11]. This approach can be related with other research to create journal interfaces that provide a personal experience trace [10], a consolidated view of a user's events and activities. SIS's interface encourages the use of iterative refinement strategies following the concept of *orienteering* [12]: an information seeking strategy in which users take small steps towards their target using partial information and contextual knowledge as a guide. This strategy is contrasted with teleporting, where users try to jump directly to their target, e.g. by a single search query.

Rich search tools like SIS could minimize the need for explicit personal information management [4], notwithstand-

ing the support that organizational structures may provide for functions other than simply re-accessing information. Search is prevalent in the web nowadays, but searching in a vast unknown unstructured domain is different than doing so in one's personal archive, as people are already familiar with many different characteristics of their documents and the contexts in which they previously encountered them, and can benefit from the use of different cues. For searching to effectively replace the need to organize personal information, it needs to cut across the many possible sources of information and include all kinds of characteristics that describe the data, in order to leverage the rich associations that characterize human memory. Whereas conventional file systems record only the last-modified date, the SIS interface is designed to use dates that are more meaningful to users and more likely to be remembered. While for a document it may be the date modified, for an email it is the date sent and for a Web site it is when the page was viewed.

In experimental evaluations, orienteering appeared to reduce the cognitive burden imposed on users by saving them from having to express exactly their information needs and by allowing them to use established habits in order to get to the proximity of the information they were looking for. Taking small steps allowed the participants to retain a sense of location, of knowing where they were, which helped in making them feel in control. They were able to get reinforcement that they were moving in the right direction, could backtrack to a previous step and felt certain that they had explored the information space completely when they were not able to find what they were looking for. The iterative process of orienteering also provided a context within which to understand the results obtained. Different ways could be used to find information, of which people and time were the most common.

INITIAL DESIGN

In the current study it was decided to investigate the efficacy of a computer file system which organized more along the lines of human memory. If successful, the whole file system might be structured in this way, but for the purposes of this experiment it was easier to concentrate on just one central aspect of the file system – the file-open dialog. A low-fidelity paper prototype was designed based on the human memory model. This was then evaluated using a Cognitive Walkthrough. This led to an improved design which was implemented as a high-fidelity prototype which was then evaluated through user testing.

People use a variety of strategies to store and retrieve their documents. Most current environments offer a standardized way to open documents from within an application. This is compatible with offering ways to browse the file system and automatically select the best application to open a document, although this is not always offered by default (e.g. iOS). File-opening dialogs are a point of contact between the personal archive and the current activity, so they provide a single point where document management can be improved. Furthermore, they are often implemented as a shared library, so a change there would be automatically propagated

¹There are reasons to be wary of the results of this study, though. In particular, it might be argued that a questionnaire might not be the best way to gauge the frequency of use of search tools.

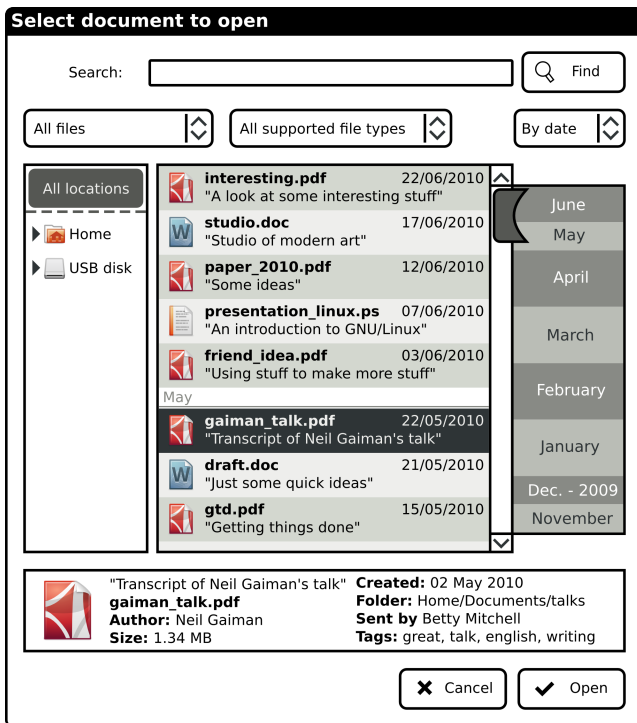


Figure 1. First iteration of the design. The search field is at the top, with selectors for file type, origin and sorting criteria below it. The file list occupies the center of the dialog, with the location selector at its left and the overview at its right. At the bottom there is the preview panel.

throughout the system. They can leverage implicit information about the user’s current task gathered from the application that is being used, allowing them to adapt their interface to the characteristics of the current application.

As described in [?], positive information practices contain two intertwined activities: horizontal exploration, where information spaces are explored at a high level, and vertical immersion, where the user becomes deeply engaged at a more detailed level. Explorability is presented as an overarching goal, complemented by orientation to provide overview and direction, high visual momentum to ease the shift between horizontal exploration and vertical immersion, and opportunities for serendipitous encounters of information. In this sense, there is an opportunity for improving the information management experience by providing a means to switch from the current task to a faceted overview of the document collection, tailored to the task at hand and which allows the use of additional facets to recall and discover relevant information, facilitating immediate vertical immersion once a relevant piece of information has been found.

The problem that we are trying to tackle is the design of a system for finding and opening files from a running application that supports the combination of multiple search strategies while remaining simple and focused. Such a solution would leverage spatial, episodic and semantic memory [1] to ease the retrieval of information, allowing for multiple possible paths and strategies to a given information item.

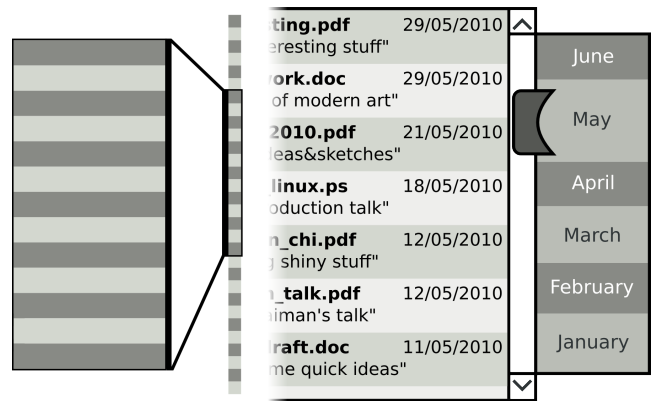


Figure 2. Metaphor (left) and wireframe (right) of the overview widget which displays a visual summary of the complete list of files; this puts the currently displayed subset in context and allows for quick scrolling to any place of the list.

The conceptual model of the proposed solution is based on the use of one big list containing all the possible files in the user’s collection that the current application can open. The user can browse this list to find the desired document, optionally using filters and different sorting criteria on it. An overview of the current contents of the list offers a high-level view and also eases browsing by providing shortcuts to any position. A preview pane after the list allows the user to confirm his selection before opening the document. The default behavior of the list is reminiscent of the *lifestreams* model proposed in [?]: a chronological stream of documents, where subsets are created by filters and searches, and where an overview aids orientation and sense-making.

The *file list* is the main element of the interface and it is located in the center. The list is ordered by one of the following sorting criteria: *useful* date [4], file name and title. The selection of this set follows research on ranking criteria for desktop search tools [?]. In this case, sorting leads quite naturally to the creation of subsets of data which can be used to structure the list; for instance the first letter in the case of file name and title sorting, or suitable periods (e.g. month) for date sorting. A fourth possibility was proposed, which would turn the list into a “Journal” that displayed the files used on each day by the user and where a given file could appear more than once; this idea was challenged by the results of the cognitive walkthrough.

The file list can be extremely long, so some kind of *visual summary* (figure 2) is needed to ease orientation by giving the user a visual indication of the part of the list that is currently being displayed. The visual metaphor chosen for this is a fish-eye widget that links the currently displayed subset of results with their position on the list. In order to increase the sense of location, this widget is annotated using the same subsets as the list. This provides a quick graphical summary of the structure of the list and allows for shortcuts: clicking on a point of the overview will scroll the list to that position. Some of the literature points that users tend to prefer to navigate through long lists of results rather than refine the search terms. This widget provides a way to accelerate these

strategies and put the results in context.

Users search for their documents by refining the list using filters and scanning it until they find the desired file, in incremental steps reminiscent of orienteering. The effect of these filters on the list is immediate, as other desktop search tools [3] and even web search engines [?] already do. Experimental evaluation of similar tools shows that the queries generated by the participants were typically short and almost half of them were followed by iterative refining and sorting of results [4].

Almost all modern desktop operating systems use folder navigation as the standard way of organizing and retrieving documents. Current solutions for file opener dialogs are mainly based on having the user navigate through the folder hierarchy until the desired target is reached, which has a number of benefits [3]. One such benefit is that the user exercises recognition rather than recall: each step down the folder hierarchy provides feedback that might confirm that this is indeed the right path. Moreover, one can recognize these intermediate steps as part of a way or path to be followed in order to arrive at the intended destination, thus leveraging procedural memory. The proposed design includes a *Location selector* that allows the user to restrict the results to a certain folder and its subfolders. The default setting is to display files from all available folders: this is indicated with the label “All locations”. Selecting a particular folder makes the list show only the files contained in that folder and all of its subfolders. It has to be noted that this solution has potential for mistakes if the user does not realize that only a subset of the data is being shown or if there is more than one file with the same name in different locations.

The *search panel* leverages semantic memory by allowing the user to recall characteristics of the document, such as its title, author, type, filename, tags and even content. Each memory fragment might be incomplete, but their combination allows incremental steps towards the goal.

The use of episodic memory to identify items and tasks in time is also documented in the literature. To leverage this kind of memory, the user can order items by the moment when they were used or created. Another way is to focus on the history of the document itself and the current context (e.g. the application that launched the dialog); thus, an *origin filter* in the search panel allows the user to filter the results according to the origin of the files. The initial possibilities for this filter are “created with this application”, “downloaded” and “sent by a contact”, although this list does not try to be exhaustive and more possibilities might appear as a result of further research.

The last step before opening the selected document is confirming that it is indeed the correct one. A *preview panel* between the file list and the “Open” and “Cancel” buttons shows more information about the selected file. The information displayed could depend on the file type; for instance, it could include a preview in the case of multimedia files.

The intended workflow for this interface is that the user would launch the dialog from a running application, examine the list (which would be faceted and adapted to the particular kinds of file desired), optionally apply filters to it, select a candidate, validate it in the preview panel and confirm the selection. Thereupon the dialog would be closed and the application would open the selected file. Thus, to support this workflow in a top to bottom manner, the filters are placed on the top-right side of the window. The file list is in the center, the preview panel is beneath it and the “Open” and “Cancel” buttons are at the bottom-right corner. This organization is intended to give a strong hint about the steps needed to open a file by establishing a visual hierarchy.

COGNITIVE WALKTHROUGH

Introduction

Method

In order to evaluate the design ideas previously outlined, three experts were asked to perform a *cognitive walkthrough* [?] on a detailed paper prototype that described the typical scenarios of the system step by step. Each one was given a description on paper of the interaction for each of the subtasks, along with a set of sub-questions. The experts were gathered together to receive an overview of the interface and then each completed a questionnaire on their own.

The evaluation covered the following tasks:

- T1 : *start*, use an “Open” button to launch the dialog;
- T2 : *sort by title*, reorder the file list using title as the sorting criteria;
- T3 : *journal*, display the files used on each day;
- T4 : *only filetype*, display only files of a certain type;
- T5 : *select folder*, display only files that are located in a certain folder and its subfolders;
- T6 : *only created with this application*, display only files that were created with the same application from which the dialog was launched;
- T7 : *sent by contact*, display only files that the user has received from a contact (e.g. through email)
- T8 : *search*, display only files that match the keywords entered
- T9 : *use overview*, understand the overview widget as a representation and shortcut to the current file list
- T10 : *preview*, display more information about the selected file
- T11 : *open file*, open the file in the application that was being used

A typical success scenario would be formed by step 1, optionally followed by zero or more of steps 2–9, followed by step 10, followed by step 11.

For each subtask, the experts were asked to answer the following questions:

	Q1	E%	Q2	E%	Q3	E%	Q4	E%
T 1	3/3	8	2/3	8	3/3	0	3/3	0
T 2	2/3	33	2/3	33	3/3	16	2/3	33
T 3	2/3	42	0/3	58	2/3	50	2/3	8
T 4	2/3	8	2/3	16	3/3	0	3/3	0
T 5	3/3	0	3/3	0	3/3	8	3/3	0
T 6	2/3	17	2/3	33	3/3	8	3/3	0
T 7	2/3	25	2/3	33	3/3	8	3/3	0
T 8	3/3	0	3/3	0	3/3	0	2/3	8
T 9	2/3	25	2/3	16	3/3	16	2/3	8
T 10	3/3	0	3/3	8	3/3	0	3/3	0
T 11	3/3	0	3/3	0	3/3	0	3/3	0

Table 1. Results of the cognitive walkthrough for each of the eleven subtasks. Columns Q1, Q2, Q3 and Q4 show the fraction of experts that answered affirmatively each question. Next to each of them we show the average of the expert’s estimations of the likelihood of a problem related to that question.

Q1 : Will the user understand that this step/subtask is needed to reach their goal?

Q2 : Will the user notice that the correct action is available?

Q3 : Will the user understand that the required step/subtask can be achieved by the action?

Q4 : Does the user get appropriate feedback if they make the appropriate action?

Results and redesign

The results of the cognitive walkthrough are summarized in Table 1. The evaluation showed that the main part of the design could work well in practice but it raised doubts about some of the elements.

The experts felt that there could be potential problems in *subtask 2* because the combobox is far from the overview widget, so the user might not discover the association between both; likewise, sorting by title while leaving the file name in bold could be confusing. The solution for this was to place the combobox right above the overview widget. The labels in the combobox were changed to “Order by X”, to make explicit the functionality of this element. Also, in order to make the field that is used for sorting more noticeable, it will be displayed in bold characters.

The “Journal” in *subtask 3* was found to be potentially confusing, as the user might not be able to understand its meaning or foretell its behavior, specially since a given file might appear more than once. For these reasons, this sorting criteria was changed to “Order by last used” in an effort to keep the episodic component while making it easier to understand.

There could be problems in *subtasks 6 and 9* because it might not be easy to discover that the systems allows the user to select the origin of the files that will be shown in the list. The problem was that the labels used to select the origin (“Created with this application” and “Sent by a contact”) were considered clear enough, but the default state of

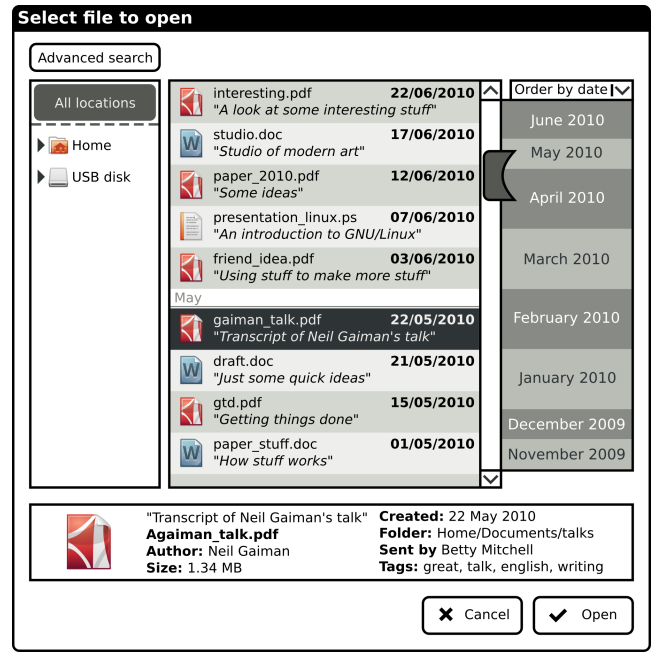


Figure 3. Redesigned UI. Search field, type filter and origin filter have been moved into the “advanced search” panel. The selector for sorting criteria is now more integrated with the overview widget. The list uses bold characters to display the field that is being used for sorting.

the combobox (“All files”) was not. As a solution, the label for the default state of the combobox was changed to “All origins”.

The overview widget created some doubts, as the experts were not sure that its usage would be obvious to the user in *subtask 9*. In this case, we took the decision to carry on with it and evaluate its usability experimentally.

Regarding the visual appearance of the dialog, the experts considered that the list of files was being overwhelmed by the filters surrounding it, making it look too small instead of taking the central place in the interface. There were also concerns about the filters and search elements located at the top of the window, as they made the whole interface look too complicated and might be a bit scary for some users.

The experts estimated that the users would have trouble understanding the dialog because the search elements might be too overwhelming; they suggested hiding them behind an “Advanced search” label, so that they would only be displayed when needed. The literature showed that in general people prefer to navigate and scroll through long lists, leaving search for when other methods have failed.

EXPERIMENT

Prototype

The redesigned version of the interface was used as a guide to build a high-fidelity prototype to be used for experimental evaluation. It was written in Python using the GTK+ toolkit, and ran on GNU/Linux. Two new widgets were developed: a cell renderer for the items in the main list, and an overview




Figure 4. High-fidelity prototype.

widget that draws a visual representation of the current state of the list and allows the user to scroll to a desired position within it.

The prototype's interface was high-fidelity but the functionality was simulated. The list of files contained 355 generated entries, 4 correct solutions to the tasks in the experiment and 51 confound items for those.

The experiment's script was contained in a small application that gave instructions to the participants in order to complete each of the four tasks. Each one of those consisted on giving the participants a description of a file and asking them to locate it using the "Open file" dialog. This involved clicking on the button below each description to open the dialog, locating the file, selecting it and clicking the "Open" button in the dialog.

Experimental design

The aim of this experiment was to validate the design decisions outlined in the previous sections; to this end, we performed an experimental evaluation of the prototype. The experiment consisted of a set of four tasks that the participants would carry out; each task asked the participant to locate a specific file after being given a very vague description of it.

We were interested in measuring three variables:

- *success rate* : whether the participant could effectively locate the target file;
- *time* : the time spent in completing each of the tasks;
- *strategy* : the way chosen to locate the file.

As the time in completing a task was one of the independent variables that we were interested in measuring, it was considered that the usage of a think-aloud protocol would interfere with their performance. For this reason, the prototype logged each action in order to create a record of the strategies followed by the participants. The resulting outcome of this part of the experiment was one log file per participant, to be analyzed later on.

Participants

Ten participants were recruited for the experiment, with the following characteristics: 60% men, 40% women; 70% post-graduate students, 30% SW engineers; 5 between 16 and 24 years old, 5 were between 25 and 34; all of them use computers almost everyday.

Procedure

The experiments were performed individually. Each participant was briefed about the purpose of the experiment and sign a statement of consent. After an initial questionnaire, participants were given a short briefing about the interface, as we were more interested in the strategies that they used than on whether they could correctly identify the interface elements on the very first time that they used it.

Participants were asked to complete four tasks individually:

Task 1: open a text document

WORD PROCESSOR

A few days ago I began writing a document about Yorkshire and its literature using this application. I would like to open it again so I can continue working on it. I think that I saved it somewhere in home/Documents, or maybe home/Writings? Your task is to locate this document. It was called lit.doc.

Task 2: open a text document

DOCUMENT VIEWER

Tonight I have a dinner with some of my friends. When I was thinking about what to cook, I remembered that another friend Louise Miller, had sent me her lasagna recipe around Fall last year. I don't remember where I saved it, but I know that I had a look at it not long ago (maybe sometime during this last spring)? Your task is to locate Louise's recipe.

Task 3: open music file

MUSIC PLAYER

I just found myself humming the song "Torn and Frayed" by the Rolling Stones and I realized that I have not heard it in a very long time. Your task is to locate this song so I can listen to it again.

Task 4: open a file

EMAIL CLIENT

Some days ago I downloaded a presentation in




Figure 5. Histograms of completion times for each task.

PDF with a lot of advice for taking care of your cat. I would like to send that presentation to my aunt because she is quite fond of cats. I will do this by attaching it to an e-mail that I am currently writing to her.

Your task is to locate this PDF document.

After the experiment, the participants answered a set of questions to measure the six dimensions of the NASA TLX [?]: Mental Demand, Physical Demand, Temporal Demand, Performance, Effort and Frustration. Questions and results are shown in table 3.

Results

Probably the most important result is the following: the participants had 100% accuracy in identifying and locating the correct files for each task. All of them were able to use the interface to retrieve the requested items. Table 2 shows the average completion time for each task. Figure 5 shows the

Task	Average (sec.)
Task 1	36.60
Task 2	58.80
Task 3	19.00
Task 4	32.30

Table 2. Average of the time devoted to complete each task.

Question	Avg	Median
Q1: Did you feel that the interface required a high level of mental effort?	2.50	2
Q2: Did you feel that the interface required a high level of physical effort?	1.30	1
Q3: Did you feel that the interface required a lot of time to use?	2.40	2
Q4: Did you feel that the interface let you successfully locate the files that you were looking for?	6.40	7
Q5: In general, did you feel that using the interface required a big effort?	1.90	1
Q6: Did you feel frustrated by the interface at any point?	1.90	1
Q7: How easy to use did you find this interface compared to the solution that you usually use?	3.00	3

Table 3. Results of the post-experiment questionnaire. Scores are on the TLX scale which ranges from 1 to 7, with 1 meaning “Not at all” (“Very easy” for Q7) and 7 meaning “Very much so” (“Very hard” for Q7).

four histograms for the completion times for each task.

All participants did comparatively better in some tasks than in others, but it seems that some of them are consistently above or under the average. Participants 1, 6, 8 and 9 performed each task in under 30 seconds. An interesting learning effect happened with participant 4, who was among the slowest for the two first tasks but managed to perform each of the last two in under 20 seconds. The opposite case, where someone performed very well at the beginning but badly at the end, did not happen.

It should be mentioned that the results of the four tasks can not be compared in aggregate. One of the reasons is that the set of possible files was not the same for each task, as they were supposed to be started from different applications that could open different kinds of files. Only document files were displayed for tasks 1 and 2, and only music files for task 3. Task 4 allowed the selection of any kind of file to be attached, so the whole set was used. Additionally, we would have to run the experiment with more participants to be able to perform a reliable statistical analysis.

Figure 6 pairs the strategies followed by each participant in each task with the time it took them to complete the task. It has to be noted that two participants found a mode prob-




Figure 6. Time per task and strategies used, for each participant and task.

lem when searching after having set a location: the search applied to files in that location and not to all files, but the participants did not notice this. These two cases are labelled as “multiple” in the figure because the participants had to try a different strategy after encountering the problem.

The analysis of figure 6 shows that the participants’ performance was faster when they could figure out a good strategy for finding the target file. When this was not the case, they tried out different strategies (e.g. different locations) and reverted to scanning the list, thus taking longer to complete the task.

The best strategies were different for each task. Interestingly, we can see that the fastest participants in each task got similarly short times despite using different strategies.

In Task 1, the fastest participants either simply scanned the default list, set the location or chose to see only files created with the current application. In Task 2, the fastest participants used search or chose to see only files sent by a contact. In Task 3, the fastest searched or set the location and then sorted the results by title. Finally, in Task 4 the fastest participants combined search, file type and/or location to find the requested item.

What does this plurality of best strategies tell us? There might not be a single “best” strategy and different people may find different solutions to the same problem. Additionally, the interface seems to be good at allowing the use of different strategies without a major penalization. This is also a confirmation that the tasks succeeded in their goal of not being so specific that they would lead the participants into using only one strategy.

It is noteworthy that only 1/5 of the task runs made any use of location. A likely explanation for this is that the wording of the tasks did not mention specific locations and the participants could not take advantage of their own spatial memory because they were not browsing their own files.

Table 3 shows a summary of the answers to the post-experiment questionnaire. The participants reported that the system had a very low physical demand, which was to be expected because of the low physical activity required to perform the experiment and its short duration of just 5–10 minutes. They felt that the interface succeeded in letting them find the documents that they were looking for, which is consistent with the overall 100% success rate. Also, they did not think that the interface required a big effort or that it was frustrating.

On the other hand, the answers were not so clear when asked about the time and mental effort required. On average, the participants estimated that the proposed interface was easier than the one that they currently use. However, some of them considered that there was not a great difference.

Discussion

The prototype was good at testing the user interface, although the use of a generated list of files and not the participant’s personal archive limited the scope of the things that we could evaluate.

All of the participants were able to complete the tasks successfully, which is a big success for the design of the interface.

The time spent in completing the tasks was generally short, but there were many cases where the participants took a long time to complete the task. This might be explained by the limitations of the prototype, which did not use the participant’s own real files; another factor could be simply that some of the participants would need some learning and practice before being able to use the interface efficiently. On the other hand, this might point to problems in the design of the interface that made it too confusing and hard to use for the users.

There were no single best strategies for each task, as each one could be completed following different paths. The participants who could plan a good strategy were the fastest in retrieving the files; the ones who relied on browsing the file list or had to try different strategies were slower. The evaluation detected a mode problem when combining filters for location and searching. This might be related to inexperience on the part of the participants, but it would be good to address this issue in the next iteration of the interface.

FURTHER WORK

The experimental evaluation has a relatively narrow scope. Specifically, the dialog was not integrated in an existing application or operating system. Furthermore the files used were not the participant's personal files (they were not even real files) and the experiment was focused on the performance of predetermined tasks rather than on the evaluation of long term usage.

Follow-up experiments should fix these problems. The dialog should be integrated in real applications so it can become part of the users' real work flows. It should work with the user's personal archive and use the real data about the files. Such a high-fidelity prototype would be suitable for long term usage in realistic settings.

Further research could implement a more realistic prototype and use it on far more participants. Then we would be able to use more advanced statistical tests to find factors that might affect the user's performance when using the interface.

This design project tries to ease the problem of file managing by providing an user interface that supported different search strategies and that could be extended to better support specific kinds of files. Further work could study ways to adapt the interface to the kind of file that a specific application intends to use. The main idea is that the dialog should look different when launched from different applications. Contextual information would be used to support the more appropriate search strategies for the type of items that the current application intends to manipulate. The representation of items in the file list should be adapted to the kind of file being displayed in each case. For instance, in the case of images a thumbnail and relevant EXIF tags could be shown.

For some kinds of documents where the location of the file might not be that important (e.g. music, photographs), other criteria might be used instead; for instance, actual physical location or tags might be used for photographs. Four sorting criteria have been discussed (file name, title, creation date and last use); other possible criteria could depend on the particular kind of file, so for instance music files could be sorted by artist, album or genre. The search panel could be extended to include more options that would allow users to further refine their searches; for instance, they could select whether the search would include the files' contents or not. The list of possible origins of the file could be expanded to include e.g. "downloaded from the Internet". Participants reported that the search is not visible enough, so we could evaluate the possibility of having it displayed all the time.

A couple of participants ran into trouble because the location selector added a mode to the interface that they were not aware of: successive iterations of the design should look for ways to remove or alleviate this problem.

The preview panel should adapt to the kind of file being preview; modern systems already do that, some even going so far as to e.g. allowing the user to listen to a selected audio file from the file chooser dialog.

SUMMARY AND CONCLUSIONS

We did not get the information from real would-be users of our system; instead of this, we used the wide corpus of existing scientific literature. Document retrieval is a well researched topic, so we were able to use that information to inform the design. The study of human memory describes different kinds of memories and different kinds of getting back to them through recognition or recall. These plurality of strategies for retrieval points towards the need to be able to combine them in order to increase their effectiveness.

The personal archive contains three different types of documents: ephemeral, current and archived. Most modern systems structure these documents in a hierarchy of folders; research shows that these structures may contain relevant information about, for instance, the decomposition of a problem or the planning of a project. One cannot simply do away with folders because there is an important value in structuring information.

Research in personal information management contrasts orienteering (take small steps towards an information goal, using partial feedback to orient oneself) and teleporting (reach the goal in one jump). The literature discusses many different interfaces for managing one's personal information. Journals are still largely experimental, but they show promising capabilities for self-reflection of past activities.

We designed a file chooser that would allow for the use and combination of different retrieval strategies while remaining as simple and focused as possible. This initial design was evaluated through a cognitive walkthrough. Preparing the materials for this was a very good way to go over the different aspects of the interface. The walkthrough was a good tool for eliciting opinions about the design and identifying possible problems. Apart from the formal evaluation, the informal interaction among the experts opinions was very interesting and productive. A redesign of the interface was carried out in order to address the problems identified by the cognitive walkthrough. This new design was used as the base for implementing a prototype using Python and GTK+. These tools proved very appropriate for fast prototyping. The implemented dialog had a high-fidelity UI but the underlying functionality was absolutely simulated. This narrowed the scope of the things that we could test, but nevertheless it served as a good tool for evaluating the main ideas behind the design.

We used this prototype to carry out an experiment where ten participants were asked to carry out four retrieval tasks

using the file chooser dialog. It is interesting to note that the experiment's script was contained in a small application that presented the tasks and instructions one by one to the participant. The application logged all of the participant's interactions, so a speak-aloud technique was not necessary. Additionally, this automatization of the experiment would have made it suitable to remote evaluations.

The experiment was successful in that the accuracy was 100%: all participants were able to accurately use the interface to complete the tasks. For each task there were several different good techniques; participants who used one of those were faster than those who had to try different strategies or simply resorted to scanning the file list. We observed learning effects, with some participants becoming comparatively faster in the latter tasks.

This experiment had a very defined scope and it did not cover the usage of real files from the participant's personal archive, nor did it use real applications or was suitable for studying long-term usage of this system. However, the experiment was reasonably successful and it sets the direction for further research that might evaluate a more complete implementation of this design.

A file-open dialog based on file characteristics pertinent to human memory was highly successful in a small-scale evaluation. Users were able to locate target files with 100% accuracy while expending very little effort.

REFERENCES

1. Baddeley, A. *Your memory: a user's guide*. Carlton Books, 2004.
2. Barreau, D., and Nardi, B. A. Finding and reminding: file organization from the desktop. *SIGCHI Bull.* 27, 3 (1995), 39–43.
3. Bergman, O., Beyth-Marom, R., Nachmias, R., Gradovitch, N., and Whittaker, S. Improved search engines and navigation preference in personal information management. *ACM Trans. Inf. Syst.* 26, 4 (2008), 1–24.
4. Cutrell, E., Dumais, S. T., and Teevan, J. Searching to eliminate personal information management. *Commun. ACM* 49, 1 (2006), 58–64.
5. Dumais, S., Cutrell, E., Cadiz, J., Jancke, G., Sarin, R., and Robbins, D. C. Stuff i've seen: a system for personal information retrieval and re-use. In *SIGIR '03: Proceedings of the 26th annual international ACM SIGIR conference on Research and development in information retrieval*, ACM (New York, NY, USA, 2003), 72–79.
6. Hicks, B. J., Dong, A., Palmer, R., and McAlpine, H. C. Organizing and managing personal electronic files: A mechanical engineer's perspective. *ACM Trans. Inf. Syst.* 26 (October 2008), 23:1–23:40.
7. Jones, W., Phuwanartnurak, A. J., Gill, R., and Bruce, H. Don't take my folders away!: organizing personal information to get things done. In *CHI '05: CHI '05 extended abstracts on Human factors in computing systems*, ACM (New York, NY, USA, 2005), 1505–1508.
8. Kaptelinin, V. Creating computer-based work environments: an empirical study of macintosh users. In *SIGCPR '96: Proceedings of the 1996 ACM SIGCPR/SIGMIS conference on Computer personnel research*, ACM (New York, NY, USA, 1996), 360–366.
9. Kaye, J. J., Vertesi, J., Avery, S., Dafoe, A., David, S., Onaga, L., Rosero, I., and Pinch, T. To have and to hold: exploring the personal archive. In *CHI '06: Proceedings of the SIGCHI conference on Human Factors in computing systems*, ACM (New York, NY, USA, 2006), 275–284.
10. Prante, T., Sauer, J., Lotfy, S., and Schmidt, A. Personal experience trace: Orienting oneself in one's activities and experiences. In *CHI '10 Workshop – Know Thyself: Monitoring and Reflecting on Facets of One's Life* (2010).
11. Ringel, M., Cutrell, E., Dumais, S., and Horvitz, E. Milestones in time: The value of landmarks in retrieving information from personal stores. In *Human-Computer Interaction INTERACT '03*, IOS Press (2003).
12. Teevan, J., Alvarado, C., Ackerman, M. S., and Karger, D. R. The perfect search engine is not enough: a study of orienteering behavior in directed search. In *CHI '04: Proceedings of the SIGCHI conference on Human factors in computing systems*, ACM (New York, NY, USA, 2004), 415–422.