

Ubiquitous Access to Documents: Using Storytelling to Alleviate Cognitive Problems

Daniel J. Gonçalves, Joaquim A. Jorge
Computer Science Department
Instituto Superior Técnico
Av. Rovisco Pais, 1049-001 Lisboa, Portugal
+351-218417269, +351-213100363
djvg@gia.ist.utl.pt, jorgej@acm.org

Abstract

With the dissemination of computers and the advent of ubiquitous computing, elder citizens and other people with cognitive problems face increasingly large barriers to the interaction with their electronic documents. The average user now handles increasingly large numbers of documents of ever increasing variety both in type and in content. Moreover, those documents can now be accessed, at times and places where such interactions were thus far impossible. Current approaches require documents to be classified prior to storage, which is a significant source of cognitive overload, retrieving documents stored days, months or years ago is even more difficult, especially for collections of large document numbers spread across many different places. Hence, new, innovative ways for storing and retrieving documents are sorely needed.

We argue that classification in itself is not necessary. By taking into account the context in which documents are manipulated, narrative descriptions can be interactively constructed and used to later retrieve them. Ubiquitous computing will allow the capture of all relevant contextual information and allow a multimodal narrative access to supplement users' cognitive shortcomings.

1 Introduction

In recent years, we have witnessed a dissemination of computing devices. They appear not only in the form of desktop computers but also as laptops, Personal Digital Assistants and mobile phones, to name a few. Given the amount of computing capabilities of most of those devices, it is now common for users to create and access documents on a wide range of platforms, and at times and places where this was not possible before. The advent of ubiquitous, pervasive computing will only increase the number of devices available to work on or access documents for any given user. However, the physic limitations of most such devices (small screens, direct manipulation on small surfaces) means their accessibility for elder and other cognitively, sensorial or motor impaired users is reduced. Unfortunately, the use of such devices is becoming a necessity on daily life.

Despite this new reality, the ways in which a user can deal with his documents remain essentially the same. Most common file systems are hierarchic in nature, requiring documents to be classified somewhere into the hierarchy. This hierarchical structure has long ago crossed the borders between file systems and other tools users must deal with, such as email and bookmarks. Users' Personal Document Spaces (PDS) are spreading to span multiple places, while growing both in number and diversity (Gonçalves, 2002). In fact, the existence of several places means documents

are no longer organized as a single hierarchy, but rather, as a *polyarchy*, for which traditional ways of document handling are not particularly effective.

We expect that users will be able to retrieve their documents by remembering how they were classified. While this could be feasible for small document numbers of one computer, it most certainly is not suitable for more complex PDS. The classification process is, in itself, fraught with problems. Users are aware that their choice will be crucial for the later recovery of documents and will try to make the best possible choice. However, there is no guarantee that an apparently good classification will be the one the user will remember and use later on. A classification can only be considered good in light of the appropriate context, and context tends to change. This leads to severe difficulties while trying to retrieve the documents. First of all, once a document is stored in the hierarchy it is no longer visible to the user. So, it is necessary for the user to remember the document's existence and search it. In that search, apart from the classification, no additional information (contents, when and why it was created, etc.) can easily be used, even if it is readily remembered by users. While some tools can use that information, they tend not to do so in an integrated, efficient, way. The whole process unnecessarily taxes the user's cognitive abilities. This leads to serious difficulties in that use by a significant percentage of the population: persons with cognitive problems. It is the case of most senior citizens. It is well known that they often face a gradual loss of motor, sensorial and cognitive skills. A system that relies on such skills for successful interaction is automatically a discriminating one, preventing them from taking full advantage of new technologies and applications. This might lead to exclusion problems, as interaction with computing devices becomes increasingly pervasive. With elderly citizens accounting for over 10% of the European population, new ways of handling documents are an imperative necessity. We should thus look into different ways of handling documents, both with respect to their manipulation and in terms of tools and techniques for organizing them. We argue that narrative access is more flexible and amenable to enhanced interaction modalities that can help overcome cognitive impairments in a flexible manner.

In the following section, we'll provide a more thorough view of the problems we just mentioned. Then a brief discussion of how some works are trying to cope with these problems will be provided. We'll then see how the computer can take advantage of a wide range of information, collectable by ubiquitous computing devices, to facilitate both the retrieval and the suggestion of documents, by building a narrative structure for each document, which can be interactively explored.

2 The Problems of Hierarchical Classification

Hierarchies are the most common type of information organization in computing systems nowadays. However, are they the most suited for a natural, straightforward interaction? One of the first works to address that concern was that of Thomas Malone. He identified the most common document organization approaches in office settings (Malone 1983): *files*, where documents are classified according to their subject and similar documents are placed together; and *piles*, sets of unclassified documents just grouped together without an apparent criterion. Piles were the more common, exactly due to problems associated with the document classification. Users were able to find their documents because the positions of the pile in the office and of the document itself in the pile were powerful enough clues for its location and contents to be guessed. Later on, Deborah Barreau and Bonnie Nardi performed an empirical study that confirmed these results on a computer setting (Barreau & Nardi 1995). The classification of documents in the hierarchy was found to be something users tend to avoid whenever possible.

In the past years, the problem has grown larger. In a set of interviews, several users actually complained that computers do not provide them with the more natural organization mechanisms they use in the physical world (Peters 1998). Those include spatial memory, as found by Malone, but also other kinds of contextual information regarding their documents. For instance, a user might not remember exactly where a document was stored in the file system but it does remember its subject, who the co-authors were and roughly when it was written. Today's tools provide little or no support for those more natural interaction modes.

Furthermore, the need for a hierarchical classification has become pervasive, since new applications tend to use it. Such is the case of bookmarks and email. It is common for users to receive several dozen email messages daily that require some kind of organization to be dealt with. However, it was found (Bälter & Sidner 2002) that users tend to avoid organizing their email. This was again explained by the fact that the most email organization tools are hierarchy-based. Another interesting aspect of email usage is the fact the most users *overload* their email tools, using them for things they were not designed for (Whittaker & Sidner 1996). Some people use their inboxes as To Do Lists or to maintain a set of unread documents. Despite all its shortcomings, email provides additional autobiographical information about the messages, such as when and by whom they were produced, other messages received roughly at the same time, etc.

All problems inherent to hierarchical classifications are made evident when the need to retrieve of a document arises. Since the way that document was classified is the primary way of finding it, it must be remembered by the user. After several weeks or even months, this information might not be easy to remember. There is no guarantee that what the user thought was a good classification back then is the same one he'd choose now. Hence, both storing and retrieving a document impose a significant and undesirable cognitive load on the user. This becomes even more serious when dealing with users with cognitive problems.

3 Some Different Approaches

It is evident that different approaches to the organization of documents are sorely needed. Several researchers have tried to develop such approaches, to abbreviate the need for a classification to be provided by the user. Spatial-based approaches try to take advantage of the spatial memory humans excel in. They allow documents to be displayed in some kind of desktop or landscape in hope that their positions help users in retrieving them. It is the case of the Data Mountain system (Robertson 1998). However, this approach suffers from a serious problem: the number of documents that can be displayed and spatially organized is limited, and far too small for today's reality. Furthermore, spatial skills are not equally developed across all users.

The temporal-based approach recognizes the importance of time when remembering a document (as seen for email overloading). It imposes a chronological order to the documents, allowing users to 'navigate through time', relating a document with others created or manipulated at roughly the same time. One of the best-known works in this area is the Lifestreams System (Freeman & Gelernter 1996). However, other ways of relating and remembering documents are often needed.

Finally, there are works based in the automatic association of properties or meta-information to documents. The groundbreaking work in the area was Gifford's Semantic File System (Gifford et al 1991), where properties such as the author and owner of documents were used to establish virtual directories (actually, database queries) for easy retrieval. This work inspired other more

sophisticated approaches such as Baeza-Yates et al's work (Baeza-Yates, Jones & Rawlins 1996). Also noteworthy is Paul Dourish et al's Placeless Documents, where an infrastructure for gathering and maintaining those properties independently from applications was developed (Dourish et al. 2000). Finally, the Haystack system (Karger, Huynh & Quan 2002) tries to bring some structure into the properties in order to more effectively relate documents.

4 Our approach: Narrative Structure

A major fault of some of approaches just mentioned is that they require extra information to be associated with documents as name-value property pairs. This is not very expressive and does not help users describe all relevant information about a document. Worse, specifying properties *actually increases the cognitive load* when classifying documents, requiring users to remember both property names and their possible values. There is need for a more structured way for users to provide that information, both in a natural and straightforward way. In short, a user must be able to *tell the story* of a document. Thus, what an efficient organizing and retrieving approach should consider is the narrative structure associated to documents. This is not limited to the usual properties other systems already gather. What is needed is an efficient way to structure metadata about a document into a coherent narrative description. Inter-relations not only with other documents but also with other users ("the document I sent to John..."), places ("... while working at home...") and events ("... shortly after that conference on HCI") must be correctly handled.

The construction of narratives requires information that might not be readily available in the user's computer but, rather, on the user's environment. Everything a user remembers should be used, and most of it will include not only his interactions with the computer but also what as being said at the time, loud noises, phone calls, and so on. Hence, a multimodal information gathering must take place. The advent of ubiquitous computing provides a straightforward solution to the gathering of that multimodal contextual information.) When we start carrying computing devices with us at all time, always on, that gathering becomes a trivial matter. Most of those devices will serve to perform relevant tasks (as occurs today with mobile phones). The communication between such devices will help relating that information and build the multimodal narrative.

Some way for the user to tell the story is also necessary, as is a mechanism for comparing stories to retrieve documents. Again, multimodality plays an important role in allowing users to present their stories in natural ways. This is especially important for elderly users, where a streamlined interaction is of the utmost importance. By not imposing rigid structures based solely on users' rote memory, both regular users and those with cognitive problems will be able to manage their documents in more easily and flexible manners. The computer can resort to several multimodal interactions to start a dialogue with the user helping him to tell the story and to remember relevant facts about the desired documents, thus solving existing ambiguities in order to find it.

5 Conclusions

The need for the hierarchical classification of documents imposes unnecessary difficulties to the storing and retrieval of those documents, especially for users with cognitive difficulties. The advent of ubiquitous computing worsens matters by allowing a larger number of documents to be manipulated in the most varied contexts. However, that same increase in number of computing devices can be of capital importance towards solving the problem. The best way to retrieve documents is resorting to what users really remember about them: their narrative structure. With new ubiquitous computing devices we'll be able to collect relevant contextual information from

the user's environment at all times, allowing the construction of multimodal narratives for a more natural and straightforward document retrieval.

The study of how those narratives can be constructed, their constituent elements and how they inter-relate, is work in progress and will be opportunely presented.

6 References

Baeza-Yates, R., Jones, T. and Rawlins G. (1996). A New Data Model: Persistent Attribute-Centric Objects. Technical Report, University of Chile.

Bälter, O and Sidner, C. (2002). Bifrost inbox organizer: Giving users control over the Inbox. *Proceedings of the Second Nordic Conference on Human-Computer Interaction*, pp 111-118, ACM Press.

Barreau, D. and Nardi, B. (1995). Finding and Reminding: File Organization from the Desktop. *ACM SIGCHI Bulletin*, 27(3), pp 39-43.

Karger, D. and Huynh, D. and Quan, D. (2002). Haystack: A platform for creating, organizing and visualizing information using RDF. *Semantic Web Workshop, The Eleventh World Wide Web Conference 2002 (WWW2002)*.

Dourish, P. *et al.* (2000). Extending Document Management Systems with User-Specific Active Properties. *ACM Transactions on Information Systems*, 18(2), pp 140-170, ACM Press.

Freeman, E. and Gelernter, D. (1996). Lifestreams: A Storage Model for Personal Data. *ACM SIGMOD Record*, 25(1), pp 80-86, ACM Press.

Gifford, D., Jouvelot, P., Sheldon, M. and O'Toole, J. (1991). Semantic File Systems. *13th ACM Symposium on Principles of Programming Languages*.

Gonçalves, D. (2002). Users and Their Documents. Technical Report, Instituto Superior Técnico. Available online at http://www.gia.ist.utl.pt/~djvg/phd/files/users_and_docs.pdf.

Malone, T. (1983). How do People Organize their Desks? Implications for the Design of Office Information Systems, *ACM Transactions on Office Information Systems*, 1(1), pp 99-112, ACM Press.

Peters, R. (1998). Organize This! Investigating Personal Information Management Practices. *Conference Proceedings on Designing Interactive Systems 2002*. ACM Press.

Robertson, G. *et al.* (1998). Data Mountain: Using Spatial Memory for Document Management. *Proceedings of the 11th Annual ACM Symposium on User Interface Software and Technology*, p153-162, ACM Press.

Whittaker, S. and Sidner, C. (1996). Email Overload: Exploring Personal Information Management of email. *Proceedings on Human Factors in Computing Systems*, pp 276-283, ACM Press.