Quill: A Narrative-Based Interface for Personal Document Retrieval

Daniel Gonçalves
Instituto Superior Técnico
Universidade Técnica de Lisboa
Av. Rovisco Pais, 49
1050-001 Lisboa, Portugal
Tel: +351-213100289
daniel.goncalves@inesc-id.pt

Joaquim A. Jorge
Instituto Superior Técnico
Universidade Técnica de Lisboa
Av. Rovisco Pais, 49
1050-001 Lisboa, Portugal
Tel: +351-213100289
jaj@inesc-id.pt

***CHI 2006,* April 22–27, 2006, Montréal, Québec, Canada.**
ACM 1-59593-298-4/06/0004.

Abstract
The ways to manage and retrieve documents have changed little in recent years. Browsing is increasingly unpractical and search still is fairly simple, relying mostly on keywords. The wide range of autobiographic information that users remember about their documents cannot be used. We present a new interaction paradigm, narrative-based interfaces, especially well suited for document retrieval. Stories make remembering information easier since it appears contextualized in a coherent whole. We describe the Quill system, a narrative-based query-formulation interface for personal document retrieval, explaining the user studies and results that led to its design in a sound and effective way. Its evaluation confirms that stories can be told naturally, containing the desired information about documents.

Keywords
Narrative-based Interfaces, Document Retrieval, Personal Information Management

ACM Classification Keywords
H.5.2 User Interfaces: User Centered Design, I.3.6 Methodology and Techniques: Interaction Techniques, H.3.3 Information Search and Retrieval: Query formulation.

Introduction

Hierarchical information organization is the predominant storage strategy for personal information. Documents are collected in hierarchically structured filesystems. Emails and bookmarks are often organized in a similar way. However, hierarchic classifications are fraught with problems. Organizing information this way leads to misclassifications, hindering its retrieval at a later time. It is not surprising that many users are reluctant to hierarchically organize their documents. [8, 10]. Despite this, in most existing document-organization approaches a document's location in the hierarchy coupled with its filename are the major, if not only, hints of its whereabouts. Not only are they fallible, but also probably not what first comes to mind when thinking about a document. Unlike the more general case of search in digital libraries or the Internet, users have an intimate knowledge of their documents. For instance, studies show that users overload their email inboxes as document storage and retrieval facilities, because email messages appear in a context, associated to all kinds of autobiographic information (sender, date, subject, etc.) that makes their retrieval easier and more natural [13].

Keyword-based search, as seen in Google Desktop or Spotlight, is not expressive enough to handle the wealth of autobiographical information that users might remember. Some property-based systems use this information [2, 7]. However, asking the users for isolated information tidbits is not effective, since remembering them is a difficult task. Some way of structuring that information is required to make the users recall it more easily. Our research shows that asking the users to tell stories about their documents can be a novel and effective way to elicit information from them. We listen to stories from an early age and tell them all our lives, regardless of location or culture [1]. They are a natural way to interact and convey information. Information in stories appears as a coherent whole, making it easier to remember. As such, stories are ideal to collect information about documents in order to retrieve them.

We developed a knowledge-based system, Quill, which is able to capture the users' document-describing stories and use them to retrieve their documents. It automatically gathers contextual information required to make sense of those stories. Below, we describe Quill, its underlying infrastructure, and the user studies that led to its creation. A brief discussion of its evaluation will ensue.

Designing the Interface

To understand what stories describing personal documents are like and verify that they contain enough information to identify and retrieve those documents, we conducted a series of semi-structured interviews. Twenty users were asked to remember documents they had handled in the past, and to tell a story describing them [3]. The users were specifically instructed to mention everything they remembered about the documents, reporting to the wider context in which they had been handled. We asked each user to tell three stories, describing: a Recent document, created by the user up to two weeks ago; an Old document, created by the user at least 6 months ago; and a document with Other authors. We collected 60 stories.

The Contents Analysis of the stories' transcripts allowed us to identify the different information elements that constitute them [6]. We found 951 story elements in

the 60 stories, that were, on average, 16 elements long. In descending order of frequency, they were: Time, Place, Co-Author, Purpose, Subject, Other Documents, Exchanges, Type, Tasks, Storage, Contents, Authors, Personal Events, World Events, Versions, Events, and Names. A Relational Analysis [6] yielded a set of transition probabilities between the elements, used us to train Hidden Markov Models and infer archetypical stories for each document type. Also, we created a set of qualitative guidelines for the design of narrative-based document retrieval interfaces [3]. The most important are:

- There is no need to consider different stories according to document age, since they are similar.
- Dialoguing with the users will prevent them from digressing and help to elicit more information. Most stories share similar structures, regardless of user, that can be used to guide the storytelling process.
- Much of the information in stories is context-dependant, requiring an understanding of the users' activities and environment to be useful.
- Stories about related documents often occur. We must not confuse them with the main story.
- The users rarely mention events in the world and occurring around then, even when prompted to do so.

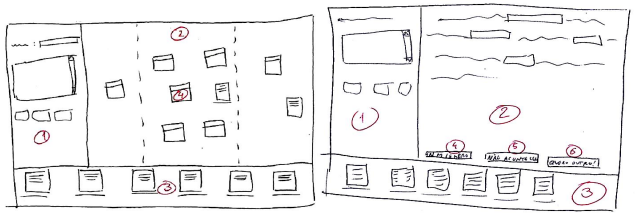

Figure 1 - Sketches of Low-Fidelity Prototypes

Choosing the interface

To choose the actual shape the interface should take, we developed and evaluated two low fidelity interface prototypes [4]. Both follow the aforementioned guidelines. Prototype A was based on the direct manipulation of story elements using a point-and-click interface (Figure 1 - left). In Prototype B, stories were represented textually (Figure 1 - right). Twenty users participated in the evaluation. Thirty stories were collected with each prototype, 20 for each document kind (Recent, Old, and Other). Those told using Prototype B were more similar to those previously analyzed than the ones told using Prototype A, considering both their contents and their structures. A questionnaire used to assess the users' subjective opinions showed Prototype B to be simpler, more novel, and easier to understand than Prototype A. For these reasons, we chose Prototype B, where the story is displayed textually, as the basis for Quill, as it better replicates the experience of telling stories to humans.

The Quill System

Quill's interface closely mimics Prototype B (Figure 2). In the large area at the top-right corner of the interface the story is incrementally written. Each element is represented by a sentence, initially incomplete, that will change to reflect the information entered by the users.

The information is entered with the help of specialized dialogues, one for each element, at the left of the story area. The story elements are suggested to the users in the order inferred from stories told to human interviewers. Under the dialogues, three buttons allow the user to enter additional information into the story. The "Done" button just commits the information in the dialogue. The "Didn't Happen" button can be used to

state that something didn't occur. Finally, the "Can't Remember" button should be pressed if the user cannot remember if something took place or not. Not knowing something is different from knowing something not to have happened. The dialogues' title can be clicked causing a list to pop up, from which the users can choose the element to be mentioned next. The system is continuously looking for the target document, based on the story told that far. The most likely candidates are displayed at the bottom of the interface window. The document's visual aspects are easy to remember. Scanning the thumbnails is, thus, easy.

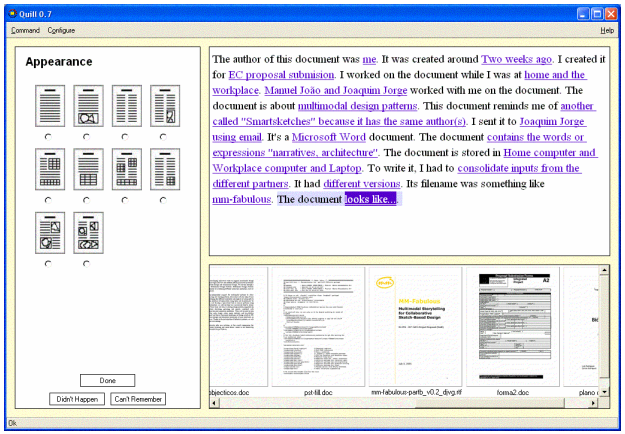

Figure 2 – The Quill Interface

The Quill Infrastructure

Central to the system is a Knowledge Base (KB). All autobiographical information required to understand the stories and find the documents is stored therein. By using a KB we can represent all sorts of common-sense-related knowledge that we'll be able to interrelate with the autobiographic information in meaningful ways. The KB uses RDF and RDF Schema as knowledge representation formalisms. We implemented a library, Scroll, which provides an abstraction layer allowing complex constructs to be handled with ease. Apart from path- and node-based inference, Scroll can use a set of case-frames defined in the `iQuill` schema to perform inference. `iQuill`'s expressiveness is similar to First-Order Logic, without the ability to represent negation and existential quantifiers (removed for efficiency reasons). Procedural attachment is allowed. A second schema, `Quill`, defines all entities required to store information about documents, their properties, and the events involving them. Common-sense knowledge is also used, to help understand the stories.

We constrained what the users can mention in their stories within reasonable limits based on the contents of stories gathered in the interviews. Free-form text entry is allowed for some elements. The dialogue in which the text is entered is a clue to what its meaning might be. For instance, in the Time dialogue any text entered by the user is likely to describe an instant in time. We parse the text with the help of context-free grammars and a chart parser, used since it provides adequate performance [11]. We used augmented grammars to automatically derive the phrases' semantics during the parsing process. The resulting information is used in conjunction with that in the KB to understand what the users are referring to.

Searching for a Document

Whenever the user enters a new element into the story, a new set of inference rules is created. Each of the documents identified as a result of those rules is assigned a score. The sum of those scores provides an overall ranking score of the document. Those with higher scores are shown to the user.

THE MONITORING SYSTEM

All relevant information is fed into the KB by a monitoring system that gathers it from different sources. It continuously observes what is going on in the users' computers, selects the relevant information, and updates the KB accordingly. It is plugin-based and currently able to monitor the web pages visited by the user (and files downloaded from the web), all emails sent and received, and all documents in the users' computer. Although each plugin works independently from the rest, provisions were made so that they produce a consistent body of knowledge. The users' privacy is guaranteed, as the KB is stored locally.

Keywords are extracted from text files and the meta-information of non-textual documents. All text is tokenized and the different words stemmed using the Porter algorithm [9]. This will make matching the elements in the users' stories to the keywords easier, as tense, plurals, and other inflections cease to be a problem. Finally, the tfidf algorithm [12] is used to select the keywords that best represent a document. All these features were built modularly to easily encompass new document types or different languages.

Evaluating Quill

To verify our hypothesis that stories can be told without any human intervention, using Quill, we collected a new set of stories and compared them with those told to a human interviewer [5]. We interviewed ten users and collected 30 stories, 10 for each document kind (Recent, Old, Other). We visited the users either at work or at home, since it was necessary to install Quill on their computers to deal with real documents.

The stories were very similar to those told using the low-fidelity prototype. The only statistically significant difference we found was that the Name of the file in which a document is stored is mentioned slightly more frequently (25 times vs. 18 using the PBF). Again, only one user, once for each story, chose a different element than the one suggested by the interface. We can conclude that stories are being told using Quill in a way similar to the way they are told to humans.

We also evaluated the trustworthiness of stories, important since incorrect information would hinder the retrieval process. After each story had been told, we asked the users to locate the document they were describing. This document was then directly compared to the information in the story. Several strategies were used in that comparison, for each of the story elements. They can be found in the study's technical report [5]. We considered two accuracy levels: elements whose accuracy was verified beyond any reasonable doubt and those that, while believed to be accurate, were not directly verified. We found that, for the most part, we can believe what users say in their stories. Up to 91% of the information is accurate, 81% of which being so beyond all doubt. This corresponds to 1 to 3 inaccurate elements per story. Given that Quill deals with story elements not as absolute truths but, rather, by weighing them in search of a document, it can properly deal with low inaccuracy rates.

Conclusions

The amount of information that must be dealt with continuously increases but little has been done to provide users with new tools and strategies to cope with this. We presented a new interaction paradigm, Narrative-Based interfaces that helps them to formulate

queries to retrieve personal documents. We've shown how those interfaces are able to elicit a wealth of information from the users that might otherwise remain unmentioned. Furthermore information is told in a natural and effortless way, as if telling a story to a human. The Quill system is an example of such an interface. Stories similar to the ones told to humans can be told using it, and we've shown the information in those stories to be accurate.

In the near future, we plan to expand the number of sources the monitoring system can handle, and, consequently, enrich the KB. An extended, long term, user study is being prepared with two goals in mind. First, we will test Quill in a real-life usage scenario, over an extended period of time. Apart from learning-curve and user-acceptance estimation issues, the analysis of the data collected at that time will allow us to answer an important research question still unaddressed in our work: are stories discriminative enough to correctly identify a single document? In second place, we will compare its performance directly with other retrieval systems. This study will provide the ultimate validation of the Quill system in particular and of narrative-based personal document retrieval paradigm in general.

Acknowledgements

This research was sponsored in part by Portuguese Science Foundation grant BIRD-POSI/EIA/59022/2004.

References

1. Brown, D. E. Human Universals. New York: McGraw-Hill 1991.
2. Dourish, P. et al. Extending Document Management Systems with User-Specific Active Properties. *ACM Transactions on Information Systems, 18(2)*, pp 140-170, ACM Press 2000.
3. Gonçalves, D. and Jorge, J., "Tell Me a Story": Issues on the Design of Document Retrieval Systems. In *Proceedings DSV-IS'04*, Lecture Notes on Computer Science, Springer-Verlag, July 2004, Hamburg, Germany.
4. Gonçalves, D. and Jorge, J. Telling Stories to Computers. *In Proceedings CHI2004*, ACM Press, 27-29 April 2004, Vienna, Austria.
5. Gonçalves, D. Real Stories about Real Documents: Evaluating the Trustworthiness of Document-Describing Stories. *Tech.Rep.*, *IST/UTL*. March 2005, (http://immi.inesc.pt/~djvg/phd/files/real_stories.zip)
6. Huberman, M. and Miles, M. Analyse des données qualitatives. Recueil de nouvelles méthodes. Bruxelles, De Boeck, 1991.
7. Karger, D. et al. Haystack: A platform for creating, organizing and visualizing information using RDF. In Proc. *Semantic Web Workshop, The Eleventh World Wide Web Conference 2002 (WWW2002).* 2002.
8. Malone, T. 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 1983.
9. Porter, M. F. 1980. An algorithm for suffix stripping. *Program* 14, 130-137.
10. Rodden, K.. How do people organize their photographs. In *Proc. of the BCS IRSG 21st Annual Colloquium on Information Retrieval Research*. 1999
11. Russel, S. and Norvig, P. Artificial Intelligence: A Modern Approach, 2nd. Ed. Prentice Hall, USA. 2003.
12. Salton, G. Automatic text processing, Addison-Wesley Longman Pub. Co.,Inc., Boston MA, 1988.
13. Whittaker, S., Sidner, C. Email overload exploring personal information management of email. In *Conference proc. on Human factors in computing systems*, pages 276-283, ACM Press, 1996.