Workshop on Desktop Search

David Elsweiler
University of Erlangen
91058 Erlangen
Germany
david@elsweiler.co.uk

Gareth Jones
School of Computing,
Dublin City University,
Ireland
Gareth.Jones@computing.dcu.ie

Liadh Kelly
School of Computing,
Dublin City University,
Ireland
lkelly@computing.dcu.ie

Jaime Teevan
Microsoft Research,
Redmond
USA
teevan@microsoft.com

Abstract

The first SIGIR workshop on Desktop Search was held on 23rd July 2009 in Geneva, Switzerland. The workshop consisted of 2 industrial keynotes, 10 paper presentations in a combination of oral and poster format and several discussion sessions. This report presents an overview of the scope and contents of the workshop and outlines the major outcomes.

1 Introduction

Desktop search refers to the process of searching within one's personal space of information. The information searched during a desktop search can include content that resides on one's personal computer (e.g., documents, emails, visited Web pages, and multimedia files), and may extend to content on other personal devices, such as music players and mobile phones.

Desktop collections are distinct from other collections such as the web in several key ways. For one, most items have been explicitly viewed, created, or saved by the searcher. As such, these items are personal to the individual and are intertwined with personal experience and memories. Good support for re-finding is important. And while explicit connections between documents are rare (other than via file organization), these experiential connections create rich possibilities for search tools to exploit. Another unique aspect of desktop collections is that they are highly multimedia. Documents come in many forms, including text, image, and video, and often have unique structured meta-data associated with them (e.g., emails have senders and recipients, documents have authors). As the available storage for desktop collections becomes cheaper and more plentiful and new media types continue to appear, the size and types of items stored in personal collections is growing rapidly. The need for effective methods to integrate the interaction experience between different information types is becoming ever more pressing.

Being retrieval related, desktop search is of interest to the information retrieval (IR) community and papers on the topic are published at top IR venues. There are, however, challenges hindering research progress. In particular, evaluation of desktop search tools is difficult. There are no established or standardized baselines or evaluation metrics, and no commonly available test collections. Privacy concerns, the challenges of working with personal collections [Elsweiler and Ruthven, 2007], and the individual differences in behaviour between users [Gwizdka and Chignell, 2007] all must be addressed to advance research in this domain. This workshop acted as a focal point for researchers

and practitioners whose work is related to Desktop Search and had the goal of fostering collaborations and addressing the challenges faced in this area.

Papers submitted to the workshop were reviewed by an international programme committee and 10 short papers and position papers were accepted for the workshop proceedings. These papers reflect the diversity of interest and contributions to the area of Desktop Search, describing research on algorithm support, memory and visualisation, task switching and context detection, the relation of desktop search to other kinds of search such as social media search and evaluation issues, amongst other topics. The proceedings are available at:
http://www.cdvp.dcu.ie/DS2010/DesktopSearchProceedings.pdf

What follows is our personal interpretation of the workshop activities, including the keynote speeches, the presented papers and the various discussion sessions. We conclude by summarising the main points raised during the workshop, the main achievements during the day and our the open points for future investigation

2 Keynote #1 – Susan Dumais Microsoft Research

We were very fortunate to have two excellent keynotes, each of which provided a different perspective on the problem of Desktop Search based on the speaker's past experiences of the topic from within industry. The workshop began with a talk by Susan Dumais, a Principal Researcher in the Context, Learning, and User Experience for Search (CLUES) Group of Microsoft Research. Susan and her colleagues were the first group to publish material on the Desktop Search within the IR community [Dumais et al., 2003]. This publication has become corner stone for research on Desktop Search and Personal Information Management. In her keynote, Susan described the motivation for starting to work in this area, the various desktop-search related projects that have since been performed at Microsoft Research, what has been learned from these projects, how Desktop Search has changed since the original paper and what the future challenges are for the field. The main take-home points of the talk were the differences between Desktop Search and other kinds of search; the different user cognitive processes involved and insight into how these can and should be supported; the range of evaluation methods that can be applied to understand Desktop Search Behaviour and the fact that we have much to investigate in this domain.

Susan outlined the main challenges for the future as 1) understanding how to support universal or specialized search – e.g. different ways to find emails and music. 2) Dealing with federation – people often have multiple desktops (PCs, mobiles and other devices). How can we best integrate search across these? Susan believes that mobile context search is especially interesting because the tasks performed are different. A related challenge concerns Desktop activities moving to to cloud-based services – e.g. twitter, facebook mail, google documents etc. This move brings both benefits and several challenges including the different interaction modes associated with different kinds of systems (e.g. ephemeral twitter content etc.). A third challenge regards evaluation methods. In her talk Susan described several methods for evaluation stating that the best way is “to give users the tool and watch how they use it”, but conceded that controlled, repeatable evaluation is extremely difficult.

3 Presented Papers

The next session consisted of short paper presentations. We selected 3 papers from those appearing in the proceedings that reflected the main themes within the papers submitted for the workshop.

The first paper "Detecting Contexts on the Desktop Using Bayesian Networks" was presented by *Tereza Iofciu* on behalf of her colleagues *Stefania Costache, Julien Gaugaz, Ekaterini Ioannou, Claudia Niederée, Wolfgang Nejdl.* from the L3S Research Centre in Hannover. The presented work investigated how Bayesian Network's can be used to combine multiple sources of evidence (e.g. concurrently open files, document content analysis, usage analysis and folder hierarchy) to detect specific user contexts that could be used to relate objects to usage contexts in order to enhance personalized desktop or web search results. Preliminary experiments show promising results regarding the choice of evidences with respect to manual user assessments. This kind of work may seed future Desktop Search algorithms.

Karl Gyllstrom presented "Not gone, but forgotten: Helping users re-find web pages by identifying those which are most likely to be lost" co-authored with Elin Rønby Pedersen. This work suggested various ranking methods for Desktop Search based on applying users' temporal document access patterns to determine the documents that are both important and have not been recently accessed (indicating greater potential for loss), understanding users' topical access patterns to determine the topics that are more unfamiliar and hence more difficult to re-find documents within, and assessing users' difficulties in originally finding documents in order to predict future difficulties in re-finding them.

The final paper presentation was "deskWeb2.0: Combining Desktop and Social Search", presented by *Sergej Zerr,* and co-authored *by Elena Demidova, Sergey Chernov.* This work explored the relationship between Desktop Search and Search over Social-networks, presenting example tasks which would require information from each of these sources. The work introduced a ranking function that combined results and from various sources and a small user study found improved performance for this new function against individual sources. This work demonstrates the interesting relationship between desktop search and related cloud-based services.

4 Initial Discussion Session

The first keynote and presented papers provided an excellent platform for discussion. An initial round of debate served to raise important themes in the work presented so far and discuss how this influences the field. The discussion revolved around creating a working definition of Desktop Search and scoping research aims for the field. The first point concerned what is exactly included in Desktop Search. The presentations had mentioned search contexts outwith the desktop including on mobile devices and in the cloud. Several examples were given of information that researchers should be interested in, but key definitions that participants agreed were useful "Stuff the user has seen before" and "Stuff related to an individual that he should see". These definitions broaden the scope of Desktop Search beyond the desktop.

A second point of debate related to Search vs Re-finding. Many of the search interfaces that have been proposed in the community have gone beyond search to include other kinds of tactic e.g. browsing, reordering the information space, navigation etc. The participants largely agreed that while search is important, perhaps it is better to consider information access in general.

A third important point that came up in the first round of discussion related to tasks for Desktop Search. It was suggested that a small pool of tasks should be developed and studied specifically across researchers. There was a general consensus that certain tasks were currently being ignored. For instance, recall-oriented tasks are often ignored e.g. finding all the files to put in a folder. There is

the well-recognised fact that people often forget to look for something because they forget that it exists. This is a task that is problem in real-life but difficult to simulate and study in the lab.

5 Poster Session

The seven workshop papers not presented orally in the morning presentation session were presented in a lively, engaging poster session. The posters covered a wide spectrum of the desktop search space, covering topics from interfaces, to backend search, task creation, metadata. The list of papers covered in this session follow:

- Memory Support for Desktop Search, *Yi Chen, Liadh Kelly, Gareth J.F. Jones*
- Using Personal Information Visualization for Document Retrieval, *Paulo Gomes, Sandra Gama, Daniel Gonçalves*
- Towards Task-Organised Desktop Collections, *Yukun Li, David Elsweiler, Xiaofeng Meng*
- History Structure for Exploring Desktop Data, *Harumi Murakami*
- Improving Re-Finding upon Work Resumption, *Thorsten Prante, Jens Sauer, Albrecht Schmidt*
- Can Maps Provide the Answer to Desktop "Search"?, *Roy A. Ruddle*
- Metasearch Tools for Desktop Search, *Paul Thomas, David Hawking*

6 Keynote #2 – Gregory Greffenstette, Exalead

The afternoon session started with our second keynote, which was given by Gregory Greffenstette, the chief-scientific officer for Exalead. Gregory talked about Desktop Search from the perspective of Enterprise Search – making content available from various sources within a company including individual employee's desktops. Enterprise Search was compared against Websearch, contrasted by talking about issues such as relevance rather than links, the importance of rich semantics and the challenges involved in working with structured data in varied formats e.g. databases, lotus notes etc.. Gregory also described the challenges in interfacing with other enterprise applications. Through presenting Exalead's desktop search tool, he demonstrated how his company has been attempting to deal with these problems. Talking about the future, Gregory provocatively questioned whether Desktop Search will still be an issue in the future - "won't everything be online anyway?". Gregory suggested that the cloud could be an alternative to the desktop i.e. the desktop could be stored in the cloud. However, this does not take away from the fact that people still need to re-access and re-use information they have seen before or personal to them. With respect to future challenges, Gregory cited merging local data on desktop with global enterprise data as a problem. He also raised issues such as security and privacy when moving towards a cloud-based model. This concluded a talk that provided the participants with two very different industry-based perspectives on Desktop search.

7 Breakout Sessions

The first of the afternoon's discussion sessions afforded workshop attendees the opportunity to discuss topics of interest to them and of importance to the desktop search domain of research. Attendees voted to split into two discussion groups, one exploring the challenges, important issues, and possibilities for future exploration for HCI in the desktop search domain, the other discussed algorithm challenges, important issues, and possibilities for future exploration for backend desktop search algorithms. Following breakout discussions, workshop attendees regrouped and a representative from each group presented a summary of their breakout discussion, leading to further group discussion on the topics.

From the HCI breakout group issues of the best means to present desktop search results were raised. It was felt that ranking approaches should be combined with browsing, in other words search and browsing should be combined. Users could then browse through, or work with a re-ranked set using an appropriate interface. The HCI breakout group also raised the point that individuals are more likely to remember atypical behaviour and considered how this might be factored in desktop search interfaces. It was suggested that a successful application might first present subjects with the most common metadata and then present them with less common metadata. So faceted search might be extended, allowing people to search based on recollection of such things as the required item contains images or contains an index, or the required item is a scientific report, for example.

The algorithms group focused their discussion on search aspects rather than other forms of interaction. They established that their main goal is to develop tools that return the sought after item purely from a textual query and without the need for other forms of interaction. They discussed how algorithms can take advantage of what we currently know about re-finding behaviour i.e. that memory plays an important role and that different factors influence available memories, and how search algorithms can exploit user interaction over time to improve performance, i.e., through (better feature value estimation, personalization of ranking method).

8 Evaluation Challenges Discussion Session

The final discussion session of the workshop was dedicated to the evaluation challenges faced in the desktop search domain. Task creation approaches, the difficulties of working with personal collections, difficulties in gaining experiment subjects, difficulties in building up personal collections, lack of established or standardized baselines and evaluation metrics, and lack of commonly available test collections were all discussed.

A key outcome of this session was a need to move towards standardised repeatable evaluation techniques in the desktop search space. However, there are a number of significant challenges associated with this, not the least of which is the fact that the data associated with this domain is personal to the individual, multimedia in nature, and different users will have different forms of collections, differing information needs and different memories of required information. Some possibilities raised here were the sharing of developed retrieval algorithms and testing of these algorithms by individuals on their personal indexed collections, and the possibility of using pseudo collections. However it was contended that it might not be possible to ever fully move towards pseudo collections. The notion of generic tasks was also raised. It was suggested that these could be implemented in form of blackbox, remote testing. Users would run algorithms on their system and evaluate the rankings, but the writer of the algorithm would not have access to the personal data.

9 Wrap-up Session

During the wrap-up session a definition for desktop search was discussed. The earlier points regarding scope and presentation of work relating to mobile and cloud stores, as well as discussion of behavioural strategies for relocating documents led to the decision that Desktop Search is not the best name for the field. It was decided that a more suitable title for this space of research would be *Personal Information Access (PIA)*.

Given the current evaluation challenges associated with PIA research, the wrap-up session also discussed the challenges associated with getting PIA papers published at leading forums. Key points raised to consider when targeting such forums included the need to be upfront and describe the space one's work is carving out, the need to make clear the scope of the presented study, the need to ensure

one makes clear that all possible avenues to remove bias were considered, and the importance of using appropriate keywords to describe one's work, such as 'user studies' for example.

Evaluation is one of the key areas hindering research in this domain. Means to allow for repeatable standardized evaluation approaches is becoming ever more pressing. The need to move towards standardized evaluation approaches was raised again in the wrap-up session. Moving towards TREC style evaluation approaches was proposed as the most viable option, with first steps of creating collections and tasks, and establishing what should be measured. Use of collections such as tweets was proposed as possible realistic starting point to address this challenge.

10 Concluding Remarks

Following an exciting, lively workshop there were several important outcomes from the day. Of particular note, the term *desktop search* was replaced at the workshop with *personal information access* (PIA) to more accurately describe the scope of work covered in this domain. PIA refers to the process of accessing information within one's personal space of digital information, e.g., searching one's desktop or mobile phone for required data items or information.

While progress has been made in this domain to date, as was highlighted at the workshop as personal collections grow and incorporate more media types and devices, including PDAs, the cloud, etc, many new exciting challenges are raised for the community. As was highlighted in the workshop the research community can complement industry in this regard, through the generation of new forms of metadata for example.

A key challenge raised in the workshop was the pressing need to move towards means of standardised repeatable evaluation in this domain. This point was raised for post workshop follow up.

11 References

Dumais, S.; Cutrell, E.; Cadiz, J.; Jancke, G.; Sarin, R. & Robbins, D. (2003), Stuff I've seen: a system for personal information retrieval and re-use, *in* Proc. 26th ACM SIGIR conference on Research and development in information retrieval', ACM Press, New York, NY, USA, pp. 72--79.

Elsweiler, D. & Ruthven, I. (2007), Towards task-based personal information management evaluations, *in* 'SIGIR '07: Proc. of the 30th annual international ACM SIGIR conference on Research and development in information retrieval', ACM Press, New York, NY, USA, pp. 23--30.

Gwizdka, J. & Chignell, M.Jones, W. & Teevan, J., ed., (2007), *Personal Information Management*, Seattle: University of Washington Press., chapter Individual differences, pp. 206-220.

12 Acknowledgements

We would like to thank ACM and SIGIR for hosting the workshop. Thanks also go to the program committee (Leif Azzopardi (University of Glasgow, UK), Ofer Bergman (Sheffield University, UK), Daragh Byrne (Dublin City University, Ireland), Rob Capra (University of North Carolina, USA), Yi Chen (Dublin City University, Ireland), Sergey Chernov (University of Hanover, Germany), Bruce Croft (University of Massachusetts, USA), Ed Cutrell (Microsoft Research, India), Susan Dumais (Microsoft Research, USA), Karl Gyllstrom (Katholieke Universiteit Leuven, Belgium), Diane Kelly

(University of North Carolina, USA), Yukun Li (Renmin University, China) and Ian Ruthven (University of Strathclyde, UK)), invited speakers (Susan Dumais (Microsoft Research Redmond, USA) and Gregory Grefenstette (Exalead, France)), paper authors and workshop attendees, without whom the workshop wouldn't have been the success it was.