What Makes Re-finding Information Difficult? A Study of Email Re-finding

David Elsweiler, Mark Baillie, and Ian Ruthven

[1] Dept. Computer Science 8 (AI), Univeristy of Erlangen-Nuremberg, Germany.
[2] Dept. Computer and Information Sciences, University of Strathclyde, Scotland.
david@elsweiler.co.uk, bailliem@gmail.com, ir@cis.strath.ac.uk

Abstract. Re-finding information that has been seen or accessed before is a task which can be relatively straight-forward, but often it can be extremely challenging, time-consuming and frustrating. Little is known, however, about what makes one re-finding task harder or easier than another. We performed a user study to learn about the contextual factors that influence users' perception of task difficulty in the context of re-finding email messages. 21 participants were issued re-finding tasks to perform on their own personal collections. The participants' responses to questions about the tasks combined with demographic data and collection statistics for the experimental population provide a rich basis to investigate the variables that can influence the perception of difficulty. A logistic regression model was developed to examine the relationships between variables and determine whether any factors were associated with perceived task difficulty. The model reveals strong relationships between difficulty and the time lapsed since a message was read, remembering *when* the sought-after email was sent, remembering *other recipients* of the email, the experience of the user and the user's filing strategy. We discuss what these findings mean for the design of re-finding interfaces and future re-finding research.

1 Introduction

Personal information management (PIM) as a research field covers efforts to understand PIM behaviour, the information strategies that people employ and attempts to develop systems that help people manage and re-find their information effectively. One PIM activity that has received growing research attention in recent times is information re-finding, where there has been a specific emphasis placed on improving the tools available to assist people with locating information that they have previously seen or accessed. In IR this has involved developing interfaces for Desktop Search [16] and creating testbeds to simulate re-finding behaviour and evaluate the performance of Desktop Search algorithms [12, 25].

In this paper we argue that while these aspects are important, before focusing too heavily on the development of improved tools, and before we can adequately simulate re-finding behaviour with test collections, we need an improved understanding of real behaviour. This includes learning about the re-finding tasks

that people perform, the user's perception of these tasks, the factors that affect this perception and how these factors influence the behaviour the user exhibits. We add to this understanding by analysing data collected from a study of email re-finding behaviour to determine the factors that influence the user's perception of task difficulty. Understanding what makes a re-finding task difficult offers researchers insight into which aspects of behaviour are important to study and where and how to provide appropriate user support. Our findings have implications for the design of re-finding tools and for the design of experiments used to understand behaviour and evaluate new tools.

2   Background Literature

This section describes related work that hints at factors that may influence user behaviour. Section 2.1 provides a summary of research relating to organising and management behaviours. In Section 2.2, the focus switches to information re-finding and research that has uncovered factors that may influence this activity. The highlighted factors form the basis of our data analyses described below.

2.1   Organising and Managing Behaviour

Much of the research on organising and managing personal information stems from Malone's seminal study of office information management behaviour [28]. Malone examined how office workers organised paper-based information and uncovered two principle strategies – filing and piling – which Malone related to the tidiness or messiness of the individual. The approach has since been replicated several times to study organisational habits for different kinds of information objects, including email messages, where user strategies include frequent filing, not filing and spring-cleaning [34], as well as web bookmarks [1] and computer files [7]. Boardman and Sasse [7] investigated the similarities and differences between the organisational techniques used to manage email messages, files and web bookmarks. Although they found some examples of overlap, particularly that pro-organising participants tended to be pro-organising for all of their collections and the opposite finding for non-organising participants, the way that collections were organised were, in the main, fundamentally different for different types of information object. This finding reveals that the organisational strategy applied depends not only on characteristics of the individual user, but also on the type of information being organised. Studies of classification decisions in offices have provided similar findings. Cole [13] and Case [11] both uncovered examples where information objects were classified not primarily by their semantic content, but by the physical form of the object e.g. journals were often organised differently to books. Kwasnik [26] demonstrated that in addition to these document attributes (topic, author and form), there are situational factors, such as the document's source or intended use that influence how it may be classified.

An important feature of previous research is that it demonstrates that different factors influence the way that people behave when managing their information. As pointed out by Gwizdka and Chignell[22], these factors can be at the

level of the individual user, e.g., their "tidiness" or "messiness"; at the level of users groups, e.g., people who have different types of job or apply different management strategies; or factors relating to the context or task being performed, e.g. the type of media being managed or whether in a paper-based or digital environment.

2.2 Re-finding Behaviour

There is an increasing body of evidence suggesting that people regularly need to re-access and re-use information they have accessed in the past [16, 14, 32]. Investigations have revealed that people perform three main types of re-finding tasks: lookup tasks, item tasks and multi-item tasks [20]. Lookup tasks involve searching for specific information from within a resource, for example finding a password from within an email. Item tasks involve looking for a particular item, perhaps to pass on to someone else or when the entire contents are needed to complete the task. Multi-item tasks involve finding more than one item and often require the user to process or collate information in order to solve the task. The evidence suggests that re-finding tasks can often be difficult, time consuming and frustrating [2, 7, 8, 32]. No-one to date has compared the different factors which could make a task difficult in a single experiment. There is fragmentary evidence in the literature, however, as to which factors could play a role.

Task-Context-level factors: In a study of the factors that affect web page re-finding, Capra and Perez-Quinones [10] discovered that both the frequency with which a task is performed and the familiarity with the information source used to solve the task had an influence on method by which the participants chose to re-find. Elsweiler and Ruthven [20] discovered a link between the time period that had elapsed since the sought-after information had been last been accessed and how difficult the re-finding task was perceived to be. Barreau and Nardi [6] showed that people prefer to re-find older information in different ways to information that is new or used regularly. Their participants generally used search-based strategies to re-find older, archived information and browse-based strategies to re-find working or ephemeral information. Other studies suggest that time may have less of an impact on re-finding strategies, finding that the path originally taken to get to the information target is much more important [31, 9].

Group-level factors: The way people organise and store their information has been shown to influence peoples' behaviour. Elsweiler et al. [18] demonstrated that the filing strategy that a participant employed influenced what they tended to remember about the email messages they were asked to find. Similarly, in studies of file re-finding [6] and web page re-finding [32], the participants who employed different filing strategies tended to use different retrieval strategies in different contexts.

Individual-level factors. Baelter [5] found that there was a link between the number of emails in a collection and the time taken to file and retrieve them. The re-finding experience of the user may also be important. Studies of search behaviour have noted differing behaviour and differing success rates for novice

and experienced users [24, 29] and similar outcomes have been established for re-finding [2].

Therefore, according to previous research, there are many potential factors that may influence how difficult re-finding information will be, although the quality of evidence varies for different factors. To determine which of these variables have the largest influence and to understand the relationships between these variables, we describe a study of email re-finding, which takes into account many of variables outlined above. We analyse the data with the aim of establishing the factors, which influence the perception of email re-finding task difficulty.

3   Materials and methods

The decision to focus on email re-finding tasks was taken for a number of reasons. First, log-based studies of re-finding behaviour have shown that people have to re-find email messages more than any other kind of information object [16, 14]. Second, although email is an asynchronous communication tool, it is, in practice, used for collaborative working [15], data archiving [34, 5] and for task [34] and contact management [33]. Conflicting strategies for different uses of email suggests that re-finding emails could be particularly challenging and it is of interest to uncover the factors that influence this.

3.1   Experimental design

We performed a controlled user study which examined the participants' perception of the difficulty of tasks when they were required to re-find email messages. The study population included 21 participants from a well-known British university, consisting of a mix of academic and research staff, undergraduate computer science students and a post-graduate class with a variety of undergraduate academic backgrounds, including former business, geography, modern-languages and philosophy students. The participants had been using their collections for varying time periods, with the post-grads having relatively new collections (the average age of collection was less than 3 months) and the academic staff comparatively older collections (avg. age ~3years). Reflecting this, some of the collections contained few messages (min = 95) and others several thousand (max =8954, median = 5132). The participants also reported using email for different purposes. While the students tended to use email mainly for class announcements and collaborative working, the academics used email for a wide range of purposes, including task and contact management, data storage, version control, collaborative authoring, as well as simple communication.

We went to great lengths to establish realistic re-finding tasks for participants that could be performed on their own personal collections without invading individual privacy. This was achieved following the methodology proposed by [20] and involved performing a number of preliminary studies with the participants

and their peers, including interviews, collection tours and diary studies of the re-finding tasks people in these groups perform. This work allowed us to establish a pool of experimental tasks suitable for different groups of users, reflecting the contents of their collections and simulating the kinds of re-finding tasks they may perform in a naturalistic setting. Example tasks involved discovering when the security code for the computer labs was last changed, finding specific details from class announcements and finding emails regarding department events or activities.[3]

Each participant was allocated 9 tasks from these pools including 3 lookup, 3 item and 3 multi-item tasks, which were rotated appropriately. After each task was issued the participant was asked to subjectively rate the task in terms of difficulty on a 5-point scale (very easy, easy, okay, difficult, very difficult), before performing the task on his own personal collection. There was a good mix of task difficulties (median = 2, interquartile range (IQR) =1). Some of the tasks were perceived to be quite easy, while others were considered challenging. These ratings form the basis of our analyses in this paper. To determine the variables which influenced the perceived difficulty of a task we also asked a range of further questions including how long it had been since the participant had previously accessed the email to be found and the attributes of the email that the participant could remember. The recorded task data combined with the demographic data and collection statistics for the experimental population provided a rich basis to investigate the variables influencing the perception of difficulty of email re-finding tasks[4].

3.2   Analysing the data

A logistic regression model was developed to determine whether any factors were associated with perceived difficulty score assigned by the participants. All available factors (24 in total) collated from the user study were analysed initially using a stepwise procedure in order to isolate any significant relationships. The stepwise procedure automatically enters and removes factors at each step assessing the overall goodness of fit of the linear regression model ([30] provide an overview on generalized linear models and the stepwise procedure]).

Following this initial investigation, we considered specific factors of interest (see Table 1) alongside those factors included in the model during the stepwise procedure. The reasoning behind this decision was that automatic model building procedures can be erroneous as it does not consider the real world importance of each factor. Therefore, the final model presented also included those factors believed to be important based on previous research (described in Section 2). These factors were entered into the model to assess both their effect (if any) on perceived difficulty and also the relationship between these factors and the

---

[3] [20] provides detailed descriptions of the process we followed and the experimental tasks can be found in [17]

[4] The full questionnaires can be found in [17]. Further details of the experimental design and user population can be found in our previous publications [18, 19].

Task or context factors:
The type of re-finding task (look-up vs item vs multi-item [20])
The period of time since the information was last accessed (<week (hot), <month (warm) >month(cold))
How frequently the information is accessed by the participant
If the participant remembered when the information was sent
If the participant remembered other people associated with the information
If the participant remembered what the email was about or why it was sent
If the participant remembered other people associated with the email
The number of memory attributes remembered by the participant
Group factors
Filing strategy (filers vs no-filers vs spring-cleaners)
Re-finding strategy (preference for browsing vs preference for searching vs no preference)
User group (Undergraduate Students vs Postgraduate Students vs Research Staff)
Individual factors
Number of emails in collection
Re-finding experience
Age of oldest message in collection

Table 1. Specific Factors of Interest

other remaining variables found to be significant. The final model is presented in Table 2 and the effect plots for the models are shown in Figure 1, with the regression coefficient, the standard error, Z value and p-value presented for all factors included. In the effect plots the solid line represents the mean perceived difficulty value and the probability of the sample assigning a value higher than the mean irrespective of the factor. Therefore, if the factor level is above the line this represents a positive effect and vice versa.

4   Results

	Coefficient estimate	Std. Error	Z value	$Pr(> |Z|)$
(Intercept)	2.70	0.36	7.49	< 0.01
Experience: Low	-0.25	0.15	-1.67	0.10
Filing.group: No filers	**-0.54**	**0.21**	**-2.50**	**0.01***
Filing.group: Spring cleaners	**-0.61**	**0.21**	**-2.89**	**<0.01***
Other recipients: Yes	**0.43**	**0.13**	**3.30**	**<0.01***
Temperature: Warm	**-0.42**	**0.19**	**-2.26**	**0.03***
Temperature: Hot	**-0.79**	**0.17**	**-4.56**	**<0.01***
Temperature: Range	-0.09	0.19	-0.50	0.62
When: Yes	-0.26	0.14	-1.88	0.06

Table 2. Regression model for perceived difficulty (significant factors in bold)

The model demonstrates the main factors that influenced the perception of difficulty associated with the allocated tasks. Three of the five variables that feature in the model are at the task-context level. The temperature of the task, i.e. how long it had been since the required information had been accessed, was
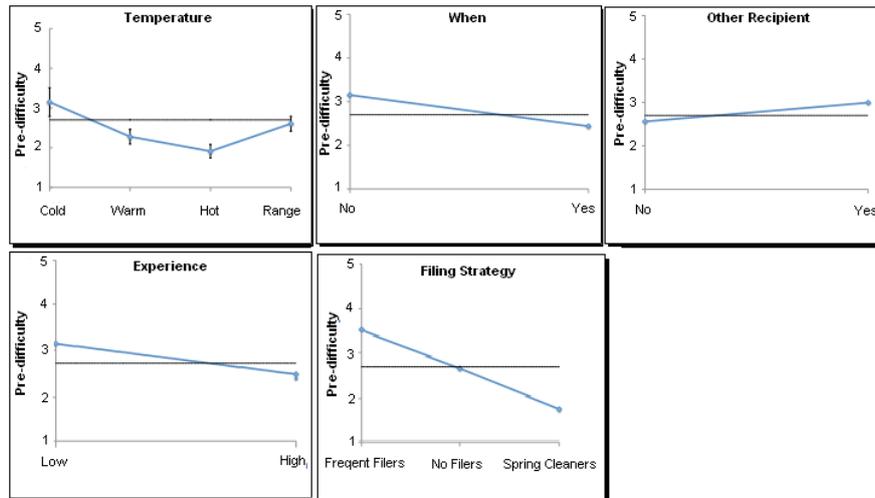
Fig. 1. Main effect plots for regression model for perceived difficulty

a significant factor in the model. Hot tasks involved finding information that had been accessed in the previous week, warm tasks information from the previous month and cold, information not accessed for longer than a month. The category range means that multiple messages were re-found and more than one temperature category was appropriate. Intuitively, the data show that longer the time period between accessing and re-finding, the more difficult the task was perceived to be. Other important factors at the task-context-level include two memory variables. When the participant remembered *when* the required email was sent, the task tended to be perceived as being easier. This also makes sense given that this information could be used directly as part of a re-finding strategy. A less intuitive feature of the model was that when the participants thought they were able to remember *other recipients* of an email, the task tended to be perceived as more difficult. It seems to make little sense that remembering extra information about a email – information that could be used to re-find the mail – could lead to increased task difficulty. However, there are possible explanations for this outcome. It could be, for example, that when other recipients are remembered, such as remembering that a mail was sent to every member of the department or company, that this is an indicator of a less personal connection with the email, leading to a poorer overall recollection. Secondly, if emails are regularly sent to these recipients, as was often the case in our tasks, the remembered information is of little benefit in the re-finding process.

The one user group-level factor that featured in the developed model was filing strategy. The participants who employed a frequent-filing strategy tended to rate re-finding tasks as harder than the user groups who utilised other filing strategies. The participants who rated tasks as easiest were the spring-cleaners, who generally do not file, but periodically, when the inbox becomes too large,

tidy up messages into appropriate folders. Again, initially, this finding makes little sense – the people who make a conscious effort to organise their mails for future retrieval in reality find tasks to be more difficult. Further, one could imagine that a spring-cleaning strategy could lead to confusion or doubt when remembering if an email was stored in a folder or not. Nevertheless, the finding aligns with previous work, which found those who file frequently to have the poorest recollection for information they need to re-find whereas spring-cleaners had the best recollection [18].

The only factor in the model that was at the level of the individual user was experience. As would be expected, participants with more experience with email tended to rate tasks as less difficult. However, although featuring in the model, this factor was not significant (p=0.1)

There are some variables that are notably absent from the model. For example, despite the numerous differences between the user groups in the experiment and the vast differences in terms of collection sizes, we found no evidence for a correlation between either of these variables and the perceived difficulty of the task. This shows that despite many of the participants having new collections with relatively few messages, there were still tasks which they perceived as difficult. It is possible that experience - a variable that did feature in the model - played a role here, with experience and developed strategies compensating against increased collection sizes. Another point of note was that there was no evidence that our participants found any type of task, i.e., lookup, item or multi-item to be especially difficult.

5   Implications

We believe the findings have important implications for future re-finding research. First, they provide indicators as to where researchers should focus their attention, highlighting particular situations or people requiring support. Second, they provide insight into how research should be performed, with implications for the design of laboratory-based user studies. We continue to outline our thoughts on these issues below.

5.1   Situations to support and study

The findings suggest that specific support is required when the user is looking for older information. Aligning with the findings of past studies [6, 20], our data show a clear correlation between the length of time that has passed since the participant previously accessed the message and the perception of the difficulty of the task. We know from Desktop Search Engine query logs that people seldom query based on date [16] and there are few other alternatives for the user to indicate that he is looking for information not seen for some time. If the user was able to communicate this information to the system, specialised support could be provided to ease the re-finding of older information. There is a need for

such support and researchers should be investigating ways of providing it, either through specific interface features or search algorithms.

Our data also indicate that particular groups of users would benefit from specific re-finding support. Inexperienced participants, despite having much smaller collections, still tended to perceive tasks as more difficult than experienced participants who had collections several times larger. Observing the participants during the completion of the tasks revealed that inexperienced participants lacked the honed strategies that more experienced participants had developed over time. This is a problem not limited to re-finding and has been noted many times in the IR literature, e.g., [2, 3, 23]. The key question that needs addressing both in IR and PIM is how do we design interfaces that both support the user to use the system with his current abilities and actively encourage and progress the development of advanced user strategies?

A second group of users who, according to our findings, would benefit from additional support are those who file frequently. People file emails for many reasons [34] so rather than encouraging these users to adopt different behaviours, what is needed is an understanding of the reasons for Frequent-filers perceiving tasks as more difficult so that appropriate support can be provided. Elsweiler et al. [18] proposed that the reason this group of users tend to remember less about their mails is that act of filing removes the message from its context and the message is thereafter seldom interacted with. This means that there is no reminder of the message and its content through natural interaction with the system. One means of addressing this problem would be to provide filers with the ability to toggle between different views of their collection. One view could show mails in their folder structure and another could show mails in a temporally ordered stream, perhaps colour-coded to indicate the folder it would normally reside in. Such a view could be filtered by various attributes as in [14] and therefore allow increased interaction with filed messages without losing the benefits offered by folders.

This point and the fact that two memory variables were shown to be significant factors in our model emphasizes the importance of memory and recollections to the process of re-finding. Memory has been studied in the PIM community e.g. [27, 18, 21]. Nevertheless, there is still much to learn regarding the role of memory and how it influences user behaviour and indeed how behaviour influences memory. Our discovery in this study that remembering other mail recipients was associated with more difficult tasks shows that remembering more is not necessarily always a positive sign. It would be interesting to learn about the relationship between memory and re-finding behaviour, e.g., the queries a user submits to a Desktop Search Engine. Is there an obvious mapping between recollection and behaviour and can this relationship be improved?

5.2   Performing Re-finding Experiments

One problem when performing lab-based user studies to investigate re-finding is the number of variables that need to be controlled or managed in the design and data analysis phases. Our findings suggest that some variables are more

important than others to consider. To achieve a balanced set of tasks in terms of task difficulty, for example, it may be more effective to control the task temperature (hot vs warm vs cold) rather than the type of task (lookup, item and multi-item) as we did. The wording of task descriptions is also important as the information provided will determine how much the participants remember when they go to complete the task. As remembering *when* an email was sent tended to be associated with easier tasks, to help create a set of tasks with a range of difficulties, we suggest that experimenters omit temporal cues from their task descriptions. This will help ensure difficult tasks are included, as it is unlikely that all participants will remember this information organically [18].

This work emphasizes that the outcomes of such studies are influenced by several variables, many of which will be confounded. We would recommend, based on our findings, that when analysing the results of any system evaluation, i.e., testing whether one system performs better than another in a user evaluation, researchers should specifically examine the data for any possible influences of the variables featured in our model, i.e., task temperature, memory variables, user filing strategy and user experience.

6   Conclusions and Future Work

To summarise, in this paper we have described a user study investigating the factors which lead to re-finding tasks being perceived as difficult, finding clear relationships between particular contextual variables and perceived task difficulty. While this work only provides understanding regarding one aspect of tasks, it opens up several important questions for further study. We have outlined our thoughts on how our findings should influence future work in terms of what researchers should be investigating and also in terms of how experimental designs and data analyses can be optimised.

In terms of our own future work, we are currently working on some of the ideas we discussed in Section 5.2. We are looking to design interfaces to support the re-finding of older information and for users who file. We are also building on these findings to improve user simulations for automated Desktop Search evaluation. Given the relationships between particular variables and task difficulty ratings, it is possible that different behaviour will be exhibited in these situations as has been shown previously for web search engine behaviour. In [3] when users were faced with a difficult task, they started to formulate more diverse queries, used advanced operators more, and spent longer on the result page. If similar behavioural changes occur when performing difficult re-finding tasks then it would important consequences for the way behaviour is simulated in automated evaluation approaches for Desktop Search. Current approaches, e.g., [4, 25] apply generalised techniques to generating queries. If we can discover different query characteristics in different situations, e.g., for difficult tasks, we can use the findings to seed improved simulations and allow the performance of algorithms to be tested in different situations, such as the re-finding of older in-

formation problem outlined above. We believe, therefore, our findings here have many important implications for future re-finding research.

References

1. D. Abrams, R. Baecker, and M. Chignell. Information archiving with bookmarks: Personal web space construction and organization. In *Proceedings of Human Factors in Computing Systems (CHI 98)*, pages 41–48. ACM Press, 1998.
2. A. Aula, N. Jhaveri, and M. Kki. Information search and reaccess strategies of experienced web users. In *Proc. 14th Intl Conf. World Wide Web*, pages 583–592. ACM Press, 2005.
3. A. Aula, R. M. Khan, and Z. Guan. How does search behavior change as search becomes more difficult? In *CHI '10: Proceedings of the 28th international conference on Human factors in computing systems*, pages 35–44, New York, NY, USA, 2010. ACM.
4. L. Azzopardi, M. de Rijke, and K. Balog. Building simulated queries for known-item topics: an analysis using six european languages. In *Proceedings of ACM SIGIR conference on Research and development in information retrieval*, pages 455–462, New York, NY, USA, 2007. ACM.
5. O. Bälter. Keystroke level analysis of email message organization. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 105–112, New York, NY, USA, 2000. ACM Press.
6. D. K. Barreau and B. Nardi. Finding and reminding: File organization from the desktop. *ACM SIGCHI Bulletin*, 27(3):39–43, 1995.
7. R. Boardman and M. A. Sasse. "stuff goes into the computer and doesn't come out": a cross-tool study of personal information management. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 583–590. ACM Press, 2004.
8. H. Bruce, W. Jones, and S. Dumais. Keeping and re-finding information on the web: What do people do and what do they need? In *Proceedings of the 67th ASIST annual meeting*, October 2004.
9. R. G. Capra and M. A. Perez-Quinones. Re-finding found things: An exploratory study of how users re-find information. Technical report, Computer Science Dept., Virginia Tech, 2003.
10. R. G. Capra and M. A. Perez-Quinones. Using web search engines to find and refind information. *Computer*, 38(10):36–42, 2005.
11. D. O. Case. Conceptual organization and retrieval of text by historians: The role of memory and metaphor. *JASIST*, 42(9):657–668, 1991.
12. S. Chernov, P. Serdyukov, P. Chirita, G. Demartini, and W. Nejdl. Building a desktop search test-bed. In *Proceedings of the 29th European conference on IR research*, pages 686–690, 2007.
13. I. Cole. Human aspects of office filing: Implications for the electronic office. In *Proceedings Human Factors Society*, 1982.
14. E. Cutrell, D. Robbins, S. Dumais, and R. Sarin. Fast, flexible filtering with phlat. In *Proceedings of the SIGCHI conference on Human Factors in computing systems*, pages 261–270, New York, NY, USA, 2006. ACM Press.
15. N. Ducheneaut and V. Bellotti. E-mail as habitat: an exploration of embedded personal information management. *interactions*, 8(5):30–38, 2001.

16. S. Dumais, E. Cutrell, J. Cadiz, G. Jancke, R. Sarin, and D.C. Robbins. Stuff i've seen: a system for personal information retrieval and re-use. In *Proceedings of ACM SIGIR conference on Research and development in information retrieval*, pages 72–79, 2003.

17. D. Elsweiler. *Supporting Human Memory in Personal Information Management*. PhD thesis, The University of Strathclyde, 2007.

18. D. Elsweiler, M. Baillie, and I. Ruthven. Memory and email re-finding. *ACM TOIS CFP special issue on Keeping, Re-finding, and Sharing Personal Information*, 2008.

19. D. Elsweiler, M. Baillie, and I. Ruthven. On understanding the relationship between recollection and refinding. *Journal of Digital Information*, 10(5), 2009.

20. D. Elsweiler and I. Ruthven. Towards task-based personal information management evaluations. In *Proceedings of ACM SIGIR conference on Research and development in information retrieval*, pages 23–30. ACM Press, 2007.

21. D. Gonçalves and J. A. Jorge. Describing documents: what can users tell us? In *Proceedings of the 9th international conference on Intelligent user interfaces*, pages 247–249, New York, NY, USA, 2004. ACM.

22. J. Gwizdka and M. Chignell. *Personal Information Management*, chapter Individual differences, pages 206–220. Seattle: University of Washington Press., 2007.

23. C. Hölscher and G. Strube. Web search behavior of internet experts and newbies. In *Proceedings of the International WWW Conference*, pages 337–346, 2000.

24. I. Hsieh-Yee. Effects of search experience and subject knowledge on the search tactics of novice and experienced searchers. *JASIST*, 44(3):161174, 1993.

25. J. Kim and W. B. Croft. Retrieval experiments using pseudo-desktop collections. In *Proceeding of the 18th ACM conference on Information and knowledge management*, pages 1297–1306, New York, NY, USA, 2009. ACM.

26. B. H. Kwasnik. How a personal document's intended use or purpose affects its classification in an office. *SIGIR Forum*, 23(SI):207–210, 1989.

27. M. Lansdale and E. Edmonds. Using memory for events in the design of personal filing systems. *International Journal of Man-Machine Studies*, 36(1):97126, 1992.

28. T. W. Malone. How do people organize their desks?: Implications for the design of office information systems. *ACM Trans. Inf. Syst.*, 1(1):99–112, 1983.

29. G. Marchionini, X. Lin, and S. Dwiggins. Effects of search and subject expertise on information seeking in a hypertext environment. In *Proceedings of the 53rd Annual Meeting of the American Society for Information Science*, 1990.

30. P. McCullagh and J. A. Nelder. *Generalized Linear Models*. Chapman & Hall/CRC, second edition, 1989.

31. J. Teevan. How people refind information when the web changes. Technical report, MIT AI, June 2004.

32. J. Teevan, C. Alvarado, M. S. Ackerman, and D. R. Karger. The perfect search engine is not enough: a study of orienteering behavior in directed search. In *Proceedings of SIGCHI conference on Human factors in computing systems*, pages 415–422, New York, NY, USA, 2004.

33. S. Whittaker, Q. Jones, and L. Terveen. Contact management: Identifying contacts to support long term communication. In *Proceedings of Conference on Computer Supported Cooperative Work*, pages 216–225. New York: ACM Press., 2002.

34. S. Whittaker and C. Sidner. Email overload: exploring personal information management of email. In M. J. Tauber, editor, *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 276–283, New York, NY, USA, 1996. ACM Press.